

Hierarchical Sliced Wasserstein Distance

Khai Nguyen[†] Tongzheng Ren[†] Huy Nguyen[†] Litu Rout[†] Tan Nguyen[◊] Nhat Ho[†]

University of Texas, Austin[†]; University of California, Los Angeles[◊]
September 28, 2022

Abstract

Sliced Wasserstein (SW) distance has been widely used in different application scenarios since it can be scaled to a large number of supports without suffering from the curse of dimensionality. The value of sliced Wasserstein distance is the average of transportation cost between one-dimensional representations (projections) of original measures that are obtained by Radon Transform (RT). Despite its efficiency in the number of supports, estimating the sliced Wasserstein requires a relatively large number of projections in high-dimensional settings. Therefore, for applications where the number of supports is relatively small compared with the dimension, e.g., several deep learning applications where the mini-batch approaches are utilized, the complexities from matrix multiplication of Radon Transform become the main computational bottleneck. To address this issue, we propose to derive projections by linearly and randomly combining a smaller number of projections which are named *bottleneck projections*. We explain the usage of these projections by introducing *Hierarchical Radon Transform* (HRT) which is constructed by applying Radon Transform variants recursively. We then formulate the approach into a new metric between measures, named *Hierarchical Sliced Wasserstein* (HSW) distance. By proving the injectivity of HRT, we derive the metricity of HSW. Moreover, we investigate the theoretical properties of HSW including its connection to SW variants and its computational and sample complexities. Finally, we compare the computational cost and generative quality of HSW with the conventional SW on the task of deep generative modeling using various benchmark datasets including CIFAR10, CelebA, and Tiny ImageNet.¹

1 Introduction

Wasserstein distance [65, 54] has been widely used in applications, such as generative modeling on images [3, 64, 56, 46], domain adaptation to transfer knowledge from source to target domains [10, 5, 47, 48], clustering problems [22], and various other applications [30, 69, 70]. Despite the increasing importance of Wasserstein distance in applications, prior works have alluded to the concerns surrounding the high computational complexity of that distance. When the probability measures have at most n supports, its computational complexity scales with the order of $\mathcal{O}(n^3 \log n)$ [52]. Additionally, it suffers from the curse of dimensionality, i.e., its sample complexity is of the order of $\mathcal{O}(n^{-1/d})$ [16], where n is the sample size and d is the number of dimensions.

Over the years, numerous attempts have been made to improve the computational and sample complexities of the Wasserstein distance. One primal line of research focuses on using entropic regularization [11]. This variant is known as entropic regularized optimal transport (or in short entropic regularized Wasserstein). By using the entropic version, one can approximate the Wasserstein distance with the computational complexities $\mathcal{O}(n^2)$ [1, 36, 37, 35] (up to some polynomial orders of

¹Code for the paper will be published at <https://github.com/UT-Austin-Data-Science-Group/HSW>.

approximation errors). Furthermore, the sample complexity of the entropic version had also been shown to be at the order of $\mathcal{O}(n^{-1/2})$ [39], which indicates that it does not suffer from the curse of dimensionality.

Another line of work builds upon the closed-form solution of optimal transport in one dimension. A notable distance metric along this direction is sliced Wasserstein (SW) distance [8]. SW is defined between two probability measures that have supports belonging to a vector space, e.g., \mathbb{R}^d . As defined in [8], SW is written as the expectation of one-dimensional Wasserstein distance between two projected measures over the uniform distribution on the unit sphere. Due to the intractability of the expectation, Monte Carlo samples from the uniform distribution over the unit sphere are used to approximate SW distance. The number of samples is often called the number of projections that is denoted as L . On the computational side, the projection directions matrix of size $d \times L$ is sampled and then multiplied by the two data matrices of size $n \times d$ resulting in two matrices of size $n \times L$ that represent L one-dimensional projected probability measures. Thereafter, L one-dimensional Wasserstein distances are computed between the two corresponding projected measures with the same projecting direction. Finally, the average of those distances yields an approximation of the value of the sliced Wasserstein distance.

Prior works [27, 14, 13, 46, 49] show that the number of projections L should be large enough compared to the dimension d for a good performance of the SW. Despite the large L , SW has lots of benefits in practice. It can be computed in $\mathcal{O}(n \log_2 n)$ time, with the statistical rate $\mathcal{O}(n^{-1/2})$ that does not suffer from the curse of dimensionality, while becoming more *memory efficient*² compared with the vanilla Wasserstein distance. For these reasons, it has been successfully applied in several applications, such as (deep) generative modeling [67, 14, 27], domain adaptation [32], and clustering [28]. Nevertheless, it also suffers from certain limitations in, e.g., deep learning applications where the mini-batch approaches [15] are utilized. Here, the number of supports n is often much smaller than the number of dimensions. Therefore, the computational complexity of solving L one-dimensional Wasserstein distance, $\Theta(Ln \log_2 n)$ is small compared to the computational complexity of matrix multiplication $\Theta(Ldn)$. This indicates that almost all computation is for the projection step. The situation is ubiquitous since there are several deep learning applications involving processing high-dimensional data, including images [17], videos [67], and text [60].

Motivated by the low-rank decomposition of matrices, we propose a more efficient approach to project original measures to their one-dimensional projected measures. In particular, two original measures are first projected into k one-dimensional projected measures via Radon transform where $k < L$. For convenience, we call these projected measures as *bottleneck projections*. Then, new L one-dimensional projected measures are created as random linear combinations of the bottleneck projections. The linear mixing step can be seen as applying Radon transform on the joint distribution of k one-dimensional projected measures. From the computational point of view, the projecting step consists of two consecutive matrix multiplications. The first multiplication is between the data matrix of size $n \times d$ and the bottleneck projection directions matrix of size $d \times k$, and the second multiplication is between the bottleneck projection directions matrix and the linear mixing matrix of size $k \times L$. Columns of both the bottleneck projection directions matrix and the linear mixing matrix are sampled randomly from the uniform distribution over the corresponding unit-hypersphere. For the same value of L , we show that the bottleneck projection approach is faster than the conventional approach when the values of L and d are large.

²SW does not need to store the cost matrix between supports.

Contribution: In summary, our main contributions are two-fold:

1. We formulate the usage of bottleneck projections into a novel integral transform, *Hierarchical Radon Transform* (HRT) which is the composition of Partial Radon Transform [34] (PRT) and Overparameterized Radon Transform (ORT). By using HRT, we derive a new sliced Wasserstein variant, named *Hierarchical sliced Wasserstein* distance (HSW). By showing that HRT is injective, we prove that HSW is a valid metric in the space of probability measures. Furthermore, we discuss the computational complexity and the projection complexity of the proposed HSW distance. Finally, we derive the connection between HSW and other sliced Wasserstein variants, and the sample complexity of HSW.
2. We design experiments focusing on comparing HSW to the conventional SW in generative modeling on standard image datasets, including CIFAR10, CelebA, and Tiny ImageNet. We show that for approximately the same amount of computation, HSW provides better generative performance than SW and helps generative models converge faster. In addition, we compare generated images qualitatively to reinforce the favorable quality of HSW.

Organization. The remainder of the paper is organized as follows. We first provide background about Radon Transform, Partial Radon Transform, and sliced Wasserstein distance in Section 2. In Section 3, we propose Overparameterized Radon Transform, Hierarchical Radon Transform, Hierarchical sliced Wasserstein distance, and analyze their relevant theoretical properties. Section 4 contains the application of HSW to generative models, qualitative results, and quantitative results on standard benchmarks. We then draw concluding remarks in Section 5. Finally, we defer the proofs of key results, supplementary materials, and discussion on related works to Appendices.

Notation. For $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, 2, \dots, n\}$. For any $d \geq 2$, $\mathbb{S}^{d-1} := \{\theta \in \mathbb{R}^d \mid \|\theta\|_2 = 1\}$ denotes the d dimensional unit sphere, and $\mathcal{U}(\mathbb{S}^{d-1})$ is the uniform measure over \mathbb{S}^{d-1} . We denote the set of absolutely integrable function on \mathbb{R}^d as $\mathbb{L}^1(\mathbb{R}^d) := \{f : \mathbb{R}^d \rightarrow \mathbb{R} \mid \int_{\mathbb{R}^d} |f(x)| dx < \infty\}$. For $p \geq 1$, we denote the set of all probability measures on \mathbb{R}^d that have finite p -moments as $\mathcal{P}_p(\mathbb{R}^d) := \{f \in \Delta(\mathbb{R}^d) \mid \mathbb{E}_f[X^p] < \infty\}$. For $m \geq 1$, we denote $\mu^{\otimes m}$ as the product measure of m random variables that follow μ . The Dirac delta function is denoted by δ . For a vector $X \in \mathbb{R}^{dm}$, $X := (x_1, \dots, x_m)$, P_X denotes the empirical measures $\frac{1}{m} \sum_{i=1}^m \delta(x - x_i)$. For any two sequences a_n and b_n , the notation $a_n = \mathcal{O}(b_n)$ means that $a_n \leq C b_n$ for all $n \geq 1$, where C is some universal constant.

2 Background

We first review the definition of the Radon Transform. We then review the sliced Wasserstein distance, and discuss its limitation in high-dimensional setting with relatively small number of supports.

Definition 1 (Radon Transform [20]). *The Radon Transform $\mathcal{R} : \mathbb{L}_1(\mathbb{R}^d) \rightarrow (\mathbb{R} \times \mathbb{S}^{d-1} \rightarrow \mathbb{R})$ is defined as: $(\mathcal{R}f)(t, \theta) = \int_{\mathbb{R}^d} f(x) \delta(t - \langle x, \theta \rangle) dx$. Note that, the Radon Transform defines a linear bijection.*

Sliced Wasserstein distance: From the definition of the Radon Transform, we can define the sliced Wasserstein distance as follows.

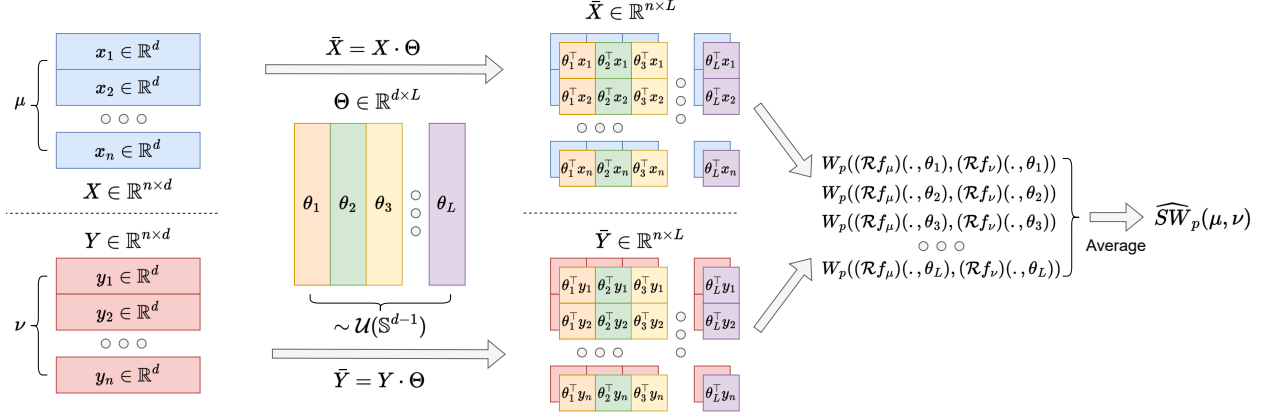


Figure 1: The Monte Carlo estimation of conventional sliced Wasserstein distance with L projections.

Definition 2 (Sliced Wasserstein Distance [8]). *For any $p \geq 1$ and dimension $d \geq 1$, the sliced Wasserstein- p distance between two probability measures $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ and $\nu \in \mathcal{P}_p(\mathbb{R}^d)$ is given by:*

$$SW_p(\mu, \nu) = \left(\mathbb{E}_{\theta \sim \mathcal{U}(\mathbb{S}^{d-1})} W_p^p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta)) \right)^{\frac{1}{p}}, \quad (1)$$

where $f_\mu(\cdot), f_\nu(\cdot)$ are the probability density functions of μ, ν respectively, and

$W_p(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|_p^p d\pi(x, y) \right)^{\frac{1}{p}}$ is the Wasserstein distance of order p [65, 53]. With slightly abuse of notation, we use $W_p(\mu, \nu)$ and $W_p(f_\mu, f_\nu)$ interchangeably.

The main benefit of sliced Wasserstein is that the one-dimensional Wasserstein distance $W_p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta))$ has a closed-form $(\int_0^1 |F_{((\mathcal{R}f_\mu)(\cdot, \theta))}^{-1}(z)} - F_{((\mathcal{R}f_\nu)(\cdot, \theta))}^{-1}(z)}|^p dz)^{1/p}$, where $F_{((\mathcal{R}f_\mu)(\cdot, \theta))}^{-1}(z)$ is the inverse cumulative distribution function of the random variable that has the density $(\mathcal{R}f_\mu)(\cdot, \theta)$ (similarly for $F_{((\mathcal{R}f_\nu)(\cdot, \theta))}^{-1}(z)$). We denote the one-dimensional pushforward measures with the density $(\mathcal{R}f_\mu)(\cdot, \theta)$ as $\theta_\# \mu$.

Sliced Wasserstein distance between discrete measures: In several applications, such as deep generative modeling [14], deep domain adaptation [32], point cloud reconstruction [50], sliced Wasserstein distance had been used for discrete measures. Let two measures μ and ν that have the pdfs $f_\mu(x) = \frac{1}{n} \sum_{i=1}^n \alpha_i \delta(x - x_i)$, $f_\nu(y) = \frac{1}{m} \sum_{j=1}^m \beta_j \delta(y - y_j)$ ($\alpha_i, \beta_j > 0 \forall i, j$), and a given projecting direction $\theta \in \mathbb{S}^{d-1}$, the corresponding Radon Transform are $(\mathcal{R}f_\mu)(z, \theta) = \frac{1}{n} \sum_{i=1}^n \alpha_i \delta(z - \theta^\top x_i)$ and $(\mathcal{R}f_\nu)(z, \theta) = \frac{1}{m} \sum_{j=1}^m \beta_j \delta(z - \theta^\top y_j)$. Since the expectation in Definition 2 is intractable, Monte Carlo estimation is used with L projecting directions $\theta_1, \dots, \theta_L \sim \mathcal{U}(\mathbb{S}^{d-1})$:

$$\widehat{SW}_p(\mu, \nu) = \left(\frac{1}{L} \sum_{i=1}^L W_p^p((\mathcal{R}f_\mu)(\cdot, \theta_i), (\mathcal{R}f_\nu)(\cdot, \theta_i)) \right)^{\frac{1}{p}}. \quad (2)$$

We denote X as the matrix of size $n \times d$ which has rows being supports of μ , $[x_1, \dots, x_n]^\top$, Y as the matrix of size $m \times d$ which has rows being supports of ν , $[y_1, \dots, y_m]^\top$, and Θ as the matrix of size $d \times L$ which has columns being sampled projecting directions $[\theta_1, \dots, \theta_L]$. The supports of the Radon Transform measures from μ and ν are the results of matrix multiplication $\tilde{X} = X \cdot \Theta$ (with

the shape $n \times L$) and $\bar{Y} = Y \cdot \Theta$ (with the shape $m \times L$). Columns of \bar{X} and \bar{Y} are supports of projected measures. Therefore, L one-dimensional Wasserstein distances are computed by evaluating the quantile functions which are based on sorting columns. A visualization is given in Figure 1.

Computational and projection complexities of sliced Wasserstein distance: Without the loss of generality, we assume that $n \geq m$, the time complexity of sliced Wasserstein is $\mathcal{O}(Ldn + Ln \log_2 n)$ where Ldn is because of matrix multiplication ($X \cdot \Theta$) and $Ln \log_2 n$ is because of the sorting algorithm. The projection complexity of SW is the memory complexity for storing the projecting directions Θ which is $\mathcal{O}(Ld)$. We would like to remark that, the value of number of projections L should be comparable to the number of dimension d for a good performance in applications [14, 46, 49].

Computational issue of sliced Wasserstein when $d \gg n$: In deep learning applications [14, 32, 27] where mini-batch approaches are used, the number of dimension d is normally much larger than the number of supports n , e.g., $d = 8192$ and $n = 128$. Therefore, $\log_2 n \ll d$ that leads to the fact that the main computation of sliced Wasserstein is for doing projecting measures $\mathcal{O}(Ldn)$. To the best of our knowledge, prior works have not adequately addressed this limitation of sliced Wasserstein.

3 Hierarchical Sliced Wasserstein distance

In this section, we propose an efficient way to improve the projecting step of sliced Wasserstein distance. In particular, we first project measures into a relatively small number of projections (k), named *bottleneck projections* ($k < L$). After that, L projections are created by a random linear combination of the bottleneck projections. To explain the usage of the bottleneck projections, we first introduce *Hierarchical Radon Transform* (HRT) in Section 3.1. We then define the *Hierarchical Sliced Wasserstein* (HSW) distance and investigate its theoretical properties in Section 3.2. We show that the usage of bottleneck projections appears in an efficient estimation of HSW.

3.1 Hierarchical Radon Transform

To define the Hierarchical Radon Transform, we first need to review an extension of Radon Transform which is *Partial Radon Transform* (PRT). After that, we propose a novel extension of Radon Transform which is named *Overparameterized Radon Transform* (ORT).

Definition 3 (Partial Radon Transform [34]). *The Partial Radon Transform $\mathcal{PR} : \mathbb{L}_1(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}) \rightarrow (\mathbb{R} \times \mathbb{S}^{d-1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R})$ is defined as: $(\mathcal{PR}f)(t, \theta, y) = \int_{\mathbb{R}^{d_1}} f(x, y) \delta(t - \langle x, \theta \rangle) dx$. Given a fixed y , the Partial Radon Transform is the Radon Transform of $f(\cdot, y)$.*

Definition 4 (Overparameterized Radon Transform). *The Overparameterized Radon Transform (ORT) of a function $f \in \mathbb{L}^1(\mathbb{R}^d)$ is defined as:*

$$(\mathcal{OR}f)(t_{1:k}, \theta_{1:k}) = \int_{\mathbb{R}^d} f(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) dx, \quad (3)$$

where $t_{1:k} := (t_1, \dots, t_k) \in \mathbb{R}^{\otimes k}$ and $\theta_{1:k} := (\theta_1, \dots, \theta_k) \in (\mathbb{S}^{d-1})^{\otimes k}$.

Definition 4 is called “overparameterized” since the dimension of the transformed function’s arguments is higher than the original dimension. Our motivation for ORT comes from the success of overparametrization in deep neural networks.

Proposition 1. *The Overparameterized Radon Transform (ORT) is injective, i.e., for any functions $f, g \in \mathbb{L}^1(\mathbb{R}^d)$, $(\mathcal{OR}f)(t_{1:k}, \theta_{1:k}) = (\mathcal{OR}g)(t_{1:k}, \theta_{1:k})$ implies that $f = g$.*

Since ORT is an extension of RT, the injectivity of ORT is derived from the injectivity of RT. The proof of Proposition 1 is in Appendix C.1.

We now define the *Hierarchical Radon Transform* (HRT).

Definition 5 (Hierarchical Radon Transform). *Hierarchical Radon Transform (HRT) of a function $f \in \mathbb{L}^1(\mathbb{R}^d)$ is defined as:*

$$(\mathcal{HR}f)(v, \theta_{1:k}, \psi) = \int_{\mathbb{R}^d} f(x) \delta \left(v - \sum_{i=1}^k \langle x, \theta_i \rangle \psi_i \right) dx, \quad (4)$$

where $v \in \mathbb{R}$, $\psi = (\psi_1, \dots, \psi_k) \in \mathbb{S}^{k-1}$, and $\theta_{1:k} = (\theta_1, \dots, \theta_k) \in (\mathbb{S}^{d-1})^{\otimes k}$.

Definition 5 is called “hierarchical” since it is the composition of Partial Radon Transform and Overparameterized Radon Transform. We can verify that $(\mathcal{HR}f)(v, \theta_{1:k}, \psi) = (\mathcal{PR}(\mathcal{OR}f))(v, \theta_{1:k}, \psi)$ (we refer the reader to Appendix B for the derivation). To the best of our knowledge, ORT and HRT have not been proposed in the literature.

Proposition 2. *The Hierarchical Radon Transform (HRT) is injective, i.e., for any functions $f, g \in \mathbb{L}^1(\mathbb{R}^d)$, $(\mathcal{HR}f)(v, \theta_{1:k}, \psi) = (\mathcal{HR}g)(v, \theta_{1:k}, \psi)$ implies that $f = g$.*

Since HRT is the composition of PRT and ORT, the injectivity of HRT is derived from the injectivity of ORT and PRT. The proof of Proposition 2 is in Appendix C.2.

Hierarchical Radon Transform of discrete measures: Let $f(x) = \frac{1}{n} \sum_{j=1}^n \alpha_j \delta(x - x_j)$, we have $(\mathcal{HR}f)(v, \theta_{1:k}, \psi) = \frac{1}{n} \sum_{i=1}^n \alpha_i \delta \left(v - \sum_{j=1}^k \langle x_i, \theta_j \rangle \psi_j \right) = \frac{1}{n} \sum_{i=1}^n \alpha_i \delta \left(v - \psi^\top \Theta^\top x_i \right)$, where Θ is the matrix which has columns are $\theta_{1:k}$. Here, $\theta_{1:k}$ are bottleneck projecting directions, ψ is the mixing direction, $\psi^\top \Theta$ is the final projecting direction, $\theta_j^\top x_i, \dots, \theta_k^\top x_i$ for any $j \in [k]$ are the bottleneck projections, and $\psi^\top \Theta^\top x_i, \dots, \psi^\top \Theta^\top x_n$ is the final projection. We consider that the bottleneck projections are the spanning set of a subspace that has the rank at most k and the final projection belongs to that subspace. Based on the linearity of Gaussian distributions, we provide the result of HRT on multivariate Gaussian and mixture of multivariate Gaussians in Appendix B.

Applications of Hierarchical Radon Transform: In this paper, we focus on showing the benefit of the HRT in the sliced Wasserstein settings. However, similar to the Radon Transform, the HRT can also be applied to multiple other applications such as sliced Gromov Wasserstein [63], sliced mutual information [18], sliced score matching (sliced Fisher divergence) [62], sliced Cramer distance [25], and other tasks that need to project probability measures.

3.2 Hierarchical Sliced Wasserstein distance

By using Hierarchical Radon Transform, we define a novel variant of sliced Wasserstein distance which is named *Hierarchical Sliced Wasserstein* (HSW) distance.

Definition 6. For any $p \geq 1$, $k \geq 1$, and dimension $d \geq 1$, the hierarchical sliced Wasserstein distance of order p between two probability measures $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ and $\nu \in \mathcal{P}_p(\mathbb{R}^d)$ is given by:

$$HSW_{p,k}(\mu, \nu) = \left(\mathbb{E}_{\theta_{1:k}, \psi} W_p^p((\mathcal{H}\mathcal{R}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_\nu)(\cdot, \theta_{1:k}, \psi)) \right)^{\frac{1}{p}}, \quad (5)$$

where $\theta_1, \dots, \theta_k \sim \mathcal{U}(\mathbb{S}^{d-1})$ and $\psi \sim \mathcal{U}(\mathbb{S}^{k-1})$.

Properties of hierarchical sliced Wasserstein distance: First, we have the following result for the metricity of HSW.

Theorem 1. For any $p \geq 1$ and $k \geq 1$, the hierarchical sliced Wasserstein $HSW_{p,k}(\cdot, \cdot)$ is a metric on the space of probability measures on \mathbb{R}^d .

Proof of Theorem 1 is given in Appendix C.3. Our next result establishes the connection between the HSW, max hierarchical sliced Wasserstein (Max-HSW) (see Definition 8 in Appendix B), max sliced Wasserstein (Max-SW) (see Definition 7 in Appendix B), and Wasserstein distance. We refer the reader to Appendix B for more theoretical properties of the Max-HSW. The role of Max-HSW is to connect HSW with Max-SW that further allows us to derive the sample complexity of HSW.

Proposition 3. For any $p \geq 1$ and $k \geq 1$, we find that

- (a) $\frac{1}{k} HSW_{p,k}(\mu, \nu) \leq \frac{1}{k} \text{Max-HSW}_{p,k}(\mu, \nu) \leq \text{Max-SW}_p(\mu, \nu) \leq W_p(\mu, \nu)$,
- (b) $SW_p(\mu, \nu) \leq \text{Max-SW}_p(\mu, \nu) \leq \text{Max-HSW}_{p,k}(\mu, \nu)$,

where we define

$$\begin{aligned} \text{Max-HSW}_{p,k}(\mu, \nu) &:= \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{H}\mathcal{R}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_\nu)(\cdot, \theta_{1:k}, \psi)), \\ \text{Max-SW}_p(\mu, \nu) &:= \max_{\theta \in \mathbb{S}^{d-1}} W_p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta)) \end{aligned}$$

as max hierarchical sliced p -Wasserstein, and max sliced p -Wasserstein, respectively.

Proof of Proposition 3 is given in Appendix C.4. Given the bounds in Proposition 3, we demonstrate that the hierarchical sliced Wasserstein does not suffer from the curse of dimensionality for the inference purpose, namely, the sample complexity for the empirical distribution from i.i.d. samples to approximate their underlying distribution is at the order of $\mathcal{O}(n^{-1/2})$ where n is the sample size.

Proposition 4. Assume that P is a probability measure supported on compact set of \mathbb{R}^d . Let X_1, X_2, \dots, X_n be i.i.d. samples from P and we denote $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ as the empirical measure of these data. Then, for any $p \geq 1$, there exists a universal constant $C > 0$ such that

$$\mathbb{E}[HSW_{p,k}(P_n, P)] \leq Ck\sqrt{(d+1)\log n/n},$$

where the outer expectation is taken with respect to the data X_1, X_2, \dots, X_n .

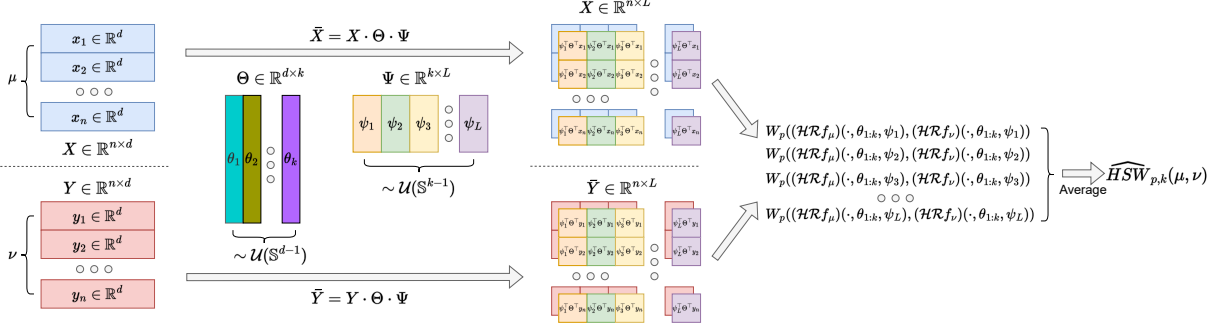


Figure 2: The Monte Carlo estimation of hierarchical sliced Wasserstein distance (HSW) with $H = 1$, k bottleneck projections, and L final projections.

Proof of Proposition 4 is given in Appendix C.5.

Monte Carlo estimation: Similar to SW, the expectation in Definition 6 is intractable. Therefore, Monte Carlo samples $\psi_1, \dots, \psi_L \sim \mathcal{U}(\mathbb{S}^{k-1})$ (by abuse of notations) and $\theta_{1,1}, \dots, \theta_{1,H}, \dots, \theta_{k,1}, \dots, \theta_{k,H} \sim \mathcal{U}(\mathbb{S}^{d-1})$ are used to approximate the HSW, which leads to the following approximation:

$$\widehat{HSW}_{p,k}(\mu, \nu) = \left(\frac{1}{HL} \sum_{h=1}^H \sum_{l=1}^L W_p^p((\mathcal{H}Rf_\mu)(\cdot, \theta_{1:k,h}, \psi_l), (\mathcal{H}Rf_\nu)(\cdot, \theta_{1:k,h}, \psi_l)) \right)^{\frac{1}{p}}. \quad (6)$$

Computational and projection complexities on discrete measures: It is clear that the time complexity of the Monte Carlo estimation of HSW with discrete probability measures of n supports is $\mathcal{O}(Hkdn + HLkn + HLn \log_2 n)$. The projection complexity of HSW is $\mathcal{O}(Hdk + kL)$. For a fast computational complexity and a low projection complexity, we simply choose $H = 1$. In this case, the computational complexity is $\mathcal{O}(kdn + Lkn + Ln \log_2 n)$ and the projection complexity is $\mathcal{O}(dk + kL)$. Recall that $k < L$, which implies the estimation of HSW is faster and more efficient in memory than the estimation of SW. This fast estimator is unbiased, though its variance might be high.

Benefit of HSW when $d \gg n$: For the same computation complexity, HSW can have a higher number of final projections L than SW. For example, when $d = 8192$ and $n = 128$ (the setting that we will use in experiments), SW with $L = 100$ has the computational complexity proportion to 104.94×10^6 and the projection complexity proportion to 0.82×10^6 . In the same setting of d and n , HSW with $H = 1, k = 50$ and $L = 500$ has the computational complexity proportion to 89.45×10^6 and the projection complexity proportion to 0.66×10^6 .

Implementation of HSW on discrete measures: The slicing process of HSW contains H matrices of size $d \times k$ in the first level and a matrix of size $k \times L$ in the second level. The projected measures are obtained by carrying out matrix multiplication between the support matrices with the projection matrices in the two layers in turn. When $H = 1$, we observe the composition of projection matrices in the two layers as a two-layer neural network with linear activation. However, the weights of the neural network have a constraint (spherical constraint) and are sampled instead of being optimized. In the paper, since we focus on the efficiency of the estimation, we consider only settings where $H = 1$. A visualization of the process when $H = 1$ is given in Figure 2.

Beyond single hierarchy with standard Radon transform: In HRT, PRT is applied on the arguments (t_1, \dots, t_k) . By applying many PRTs on subsets of arguments e.g., $(t_1, t_2), \dots, (t_{k-1}, t_k)$ and then applying PRT again on the arguments of the output function, we derive a more hierarchical version of HRT. Similarly, we could also apply ORT on the arguments multiple times to make the transform hierarchical. Despite the hierarchy, Hierarchical Sliced Wasserstein distance still uses linear projections. By changing from Radon Transform variants to Generalized Radon Transform [4] variants, we easily extend Hierarchical Radon Transform to Hierarchical Generalized Radon Transform (HGRT). Furthermore, we derive the Hierarchical Generalized Sliced Wasserstein (HGSW) distance. Since HGSW has more than one non-linear transform layer, it has the provision of using a more complex non-linear transform than the conventional Generalized sliced Wasserstein (GSW) [26]. Compared to the neural network defining function of GRT [26], HGSW preserves the metricity, i.e., HGSW satisfies the identity property due to the injectivity of HGRT. Since the defining functions in [26] e.g., circular functions and homogeneous polynomials with an odd degree are not scalable in high-dimension, we defer the investigation of HGSW and finding a more scalable function for future work.

Distributions of final projecting directions in HRT: We recall that the final projecting directions of HSW are $\psi_1^\top \Theta, \dots, \psi_L^\top \Theta$, where $\Theta = (\theta_{1:k}), \theta_{1:k} \sim \mathcal{U}(\mathbb{S}^{d-1})$, and $\psi_1, \dots, \psi_L \sim \mathcal{U}(\mathbb{S}^{k-1})$. It is clear that the final projecting directions are distributed uniformly from the manifold $\mathcal{S} := \{x \in \mathbb{R}^{d-1} \mid x = \sum_{i=1}^k \psi_i^\top \theta_i, (\psi_1, \dots, \psi_k) \in \mathbb{S}^{k-1}, \theta_i \in \mathbb{S}^{d-1}, \forall i = 1, \dots, k\}$. Therefore, HSW can be considered as the sliced Wasserstein with the projecting directions on a special manifold. This is different from the conventional unit hyper-sphere. To our knowledge, the manifold \mathcal{S} has not been explored sufficiently in previous works, which may be a potential direction of future research.

On the choices of k and L in HSW: We consider the HSW as an alternative option for applications that use the SW. For a given value of L in the SW, we can choose a faster setting of HSW by selecting $k \leq \frac{Ld}{L+d}$ based on the analysis of their computational complexities. Similarly, the HSW with L_2 final projections can be faster than the SW with L_1 projections by choosing $k \leq \frac{L_1 d - (L_2 - L_1) \log_2 n}{d + L_2}$. Later, we use this rule for the choices of k in the experiments while comparing the HSW with the SW.e experiments while comparing the HSW with the SW.

4 Experiments

In this section, we compare HSW with SW on benchmark datasets: CIFAR10 (with image size 32x32) [29], CelebA (with image size 64x64), and Tiny ImageNet (with image size 64x64) [31]. To this end, we consider deep generative modeling with the standard framework of the sliced Wasserstein generator [14, 13, 46, 42]. We provide a detailed discussion of this framework including training losses and their interpretation in Appendix D.1. The SW variants are used in the feature space with dimension $d = 8192$ and the mini-batch size 128 for all datasets. The main evaluation metrics are FID score [21] and Inception score (IS) [58]. We do not report the IS score on CelebA since it poorly captures the perceptual quality of face images [21]. The detailed settings about architectures, hyper-parameters, and evaluation of FID and IS are provided in Appendix E.

Our experiments aim to answer the following questions:

1. *For approximately the same computation, is the HSW comparable in terms of perceptual quality while achieving better convergence speed?*

Table 1: Summary of FID scores, IS scores, computational complexity, and memory complexity of different estimations of SW and HSW on CIFAR10 (32x32), CelebA (64x64), and Tiny ImageNet (64x64).

Method	Com (\downarrow)	Proj (\downarrow)	CIFAR10		CelebA	Tiny ImageNet	
			FID (\downarrow)	IS (\uparrow)	FID (\downarrow)	FID (\downarrow)	IS (\uparrow)
SW (L=100)	104.95	0.82	51.62 \pm 3.69	5.74 \pm 0.28	17.54\pm1.85	96.03 \pm 3.17	5.38 \pm 0.29
HSW (k=70, L=2000)	93.11	0.71	47.64\pm5.20	5.98\pm0.22	17.59 \pm 2.12	89.77\pm3.56	5.83\pm0.31
SW (L=1000)	1049.47	8.19	42.26 \pm 3.52	6.30 \pm 0.19	17.35 \pm 2.56	84.67 \pm 3.93	5.98 \pm 0.17
HSW (k=400, L=6000)	732.01	5.68	41.80\pm1.08	6.38\pm0.15	15.89\pm2.19	82.52\pm4.40	6.00\pm0.19
SW (L=10000)	10494.72	81.92	38.60 \pm 2.23	6.54 \pm 0.18	16.05 \pm 1.64	84.37 \pm 3.68	6.06 \pm 0.21
HSW (k=3000, L=18000)	10073.85	78.58	38.22\pm4.96	6.57\pm0.32	15.74\pm1.46	80.69\pm5.87	6.07\pm0.25




Figure 3: The FID scores over epochs of different training losses on datasets. We observe that HSW helps the generative models converge faster.

2. For the same number of final projections L and a relatively low number of bottleneck projections $k < L$, how is the performance of the HSW compared to the SW?
3. For a fixed value of bottleneck projections k , does increasing the number of final projections L improve the performance of the HSW?
4. For approximately the same computation, does reducing the value of k and raising the value of L lead to a better result?

After running each experiment 5 times, we report the mean and the standard deviation of evaluation metrics.

The HSW is usually better than the SW with a lower computation: We report the FID scores, IS scores, the computational complexity ($\times 10^6$), and the projection complexity ($\times 10^6$) for the SW with $L \in \{100, 1000, 10000\}$ and the HSW with $(k, L) = \{(70, 2000), (400, 6000), (3000, 18000)\}$ respectively in Table 1. According to the table, the HSW usually yields comparable FID and IS scores while having a **lower computational complexity** and a **lower projection complexity** on benchmark datasets. We show random generated images on CelebA in Figure 4, on CIFAR10 in Figure 6 in Appendix D.2, and on Tiny ImageNet in Figure 7 in Appendix D.2 as qualitative comparison. Those generated images are consistent with the quantitative scores in Table 1.

The HSW leads to faster convergence than the SW with a lower computation: We plot the FID scores over training epochs of the SW and the HSW with the same setting as in Table 1 in Figure 3. We observe that FID scores from models trained by the HSW (with a better computation)

Table 2: FID scores, IS scores, computational complexity, and memory complexity for ablation studies of k and L of the SW on CIFAR10 (32x32), CelebA (64x64), and Tiny ImageNet (64x64).

Method	Com (\downarrow)	Proj (\downarrow)	CIFAR10		CelebA	Tiny ImageNet	
			FID (\downarrow)	IS (\uparrow)	FID (\downarrow)	FID (\downarrow)	IS (\uparrow)
SW ($L=1000$)	1049.47	8.19	42.26 \pm 3.52	6.30 \pm 0.19	17.35 \pm 2.56	84.67 \pm 3.93	5.98 \pm 0.17
HSW ($k=500, L=1000$)	589.18	4.59	43.58 \pm 4.01	6.25 \pm 0.31	17.50 \pm 2.25	89.02 \pm 2.32	5.92 \pm 0.19
HSW ($k=500, L=4000$)	783.87	6.09	41.85 \pm 4.42	6.36 \pm 0.23	16.80 \pm 1.23	86.57 \pm 3.80	5.94 \pm 0.45
HSW ($k=400, L=6000$)	732.01	5.68	41.80 \pm 1.08	6.38 \pm 0.15	15.89 \pm 2.19	82.52 \pm 4.40	6.00 \pm 0.19
HSW ($k=100, L=50000$)	789.65	5.81	44.70 \pm 3.19	6.09 \pm 0.15	17.41 \pm 2.12	89.01 \pm 2.73	5.90 \pm 0.21

reduce faster than ones trained from the SW. The same phenomenon happens with the IS scores in Figure 5 in Appendix D.2 where IS scores from the HSW increase earlier. The reason is that the HSW has a higher number of final projections than the SW, hence, it is a more discriminative signal than the SW.

Ablation studies on k and L in the HSW: In the Table 2, we reports the additional FID scores, IS scores, the computational complexity ($\times 10^6$), and the projection complexity ($\times 10^6$) for the HSW with $(k, L) = \{(500, 1000), (500, 4000), (100, 50000)\}$. First, we see that given $k = 5000$, increasing L from 1000 to 4000 improves the generative performance on all three datasets. Moreover, we observe that the HSW with the same number of final projections as the SW ($L = 1000$) gives comparable scores while the complexities are only about half of the complexities of the SW. When decreasing k from 500 to 400 and increasing L from 4000 to 6000, the generative performance of the HSW is enhanced and the complexities are also lighter. However, choosing a too small k and a too large L does not lead to a better result. For example, the HSW with $k = 100, L = 50000$ has high complexities compared to $k = 500, L = 4000$ and $k = 400, L = 6000$, however, its FID scores and IS scores are worse. The reason is because of the linearity of the HRT. In particular, the k bottleneck projections form a subspace that has a rank at most k and the L final projections still lie in that subspace. This fact suggests that the value of k should be also chosen to be sufficiently large compared to the true rank of the supports. In practice, data often lie on a low dimensional manifold, hence, k can be chosen to be much smaller than d and L can be chosen to be large for a good estimation of discrepancy. This is an interpretable advantage of the HSW compared to the SW since it can separate between the assumption of the ranking of supports and the estimation of the discrepancy between measures.

Max hierarchical sliced Wasserstein: We also compare Max-HSW (see Definition 7 in Appendix B) with the conventional Max-SW in generative modeling in Table 3 in Appendix D.2. We observe that the overparameterization from HRT could also improve the optimization of finding good projections. We would like to recall that the Max-HSW is the generalization of the Max-SW, namely, Max-HSW with $k = 1$ is equivalent to the Max-SW. Therefore, the performance of the Max-HSW is at least the same as the Max-SW.

5 Conclusion

In this paper, we proposed a hierarchical approach to efficiently estimate the Wasserstein distance with provable benefits in terms of computation and memory. It formed final projections by combining linearly and randomly from a smaller set of bottleneck projections. We justified the main idea




Figure 4: Random generated images of the SW and the HSW on CelebA.

by introducing Hierarchical Radon Transform (HRT) and hierarchical sliced Wasserstein distance (HSW). We proved the injectivity of the HRT, the metricity of the HSW, and investigated its theoretical properties including computational complexity, sample complexity, projection complexity, and its connection to other sliced Wasserstein variants. Finally, we conducted experiments on deep generative modeling where the main computational burden due to the projection was highlighted. In this setting, HSW performed favorably in both generative quality and computational efficiency.

References

- [1] J. Altschuler, J. Niles-Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in Neural Information Processing Systems*, pages 1964–1974, 2017. (Cited on page 1.)
- [2] B. Amos. Tutorial on amortized optimization for learning to optimize over continuous domains. *arXiv preprint arXiv:2202.00665*, 2022. (Cited on page 18.)
- [3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017. (Cited on page 1.)
- [4] G. Beylkin. The inversion problem and applications of the generalized radon transform. *Communications on pure and applied mathematics*, 37(5):579–599, 1984. (Cited on page 9.)

- [5] 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. (Cited on page 1.)
- [6] C. Bonet, P. Berg, N. Courty, F. Septier, L. Drumetz, and M.-T. Pham. Spherical sliced-wasserstein. *arXiv preprint arXiv:2206.08780*, 2022. (Cited on page 18.)
- [7] C. Bonet, N. Courty, F. Septier, and L. Drumetz. Sliced-Wasserstein gradient flows. *arXiv preprint arXiv:2110.10972*, 2021. (Cited on page 18.)
- [8] 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. (Cited on pages 2 and 4.)
- [9] X. Chen, Y. Yang, and Y. Li. Augmented sliced Wasserstein distances. *International Conference on Learning Representations*, 2022. (Cited on page 18.)
- [10] 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. (Cited on page 1.)
- [11] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013. (Cited on page 1.)
- [12] B. Dai and U. Seljak. Sliced iterative normalizing flows. In *International Conference on Machine Learning*, pages 2352–2364. PMLR, 2021. (Cited on page 18.)
- [13] 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. (Cited on pages 2, 9, and 19.)
- [14] 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. (Cited on pages 2, 4, 5, 9, and 26.)
- [15] K. Fatras, Y. Zine, R. Flamary, R. Gribonval, and N. Courty. Learning with minibatch Wasserstein: asymptotic and gradient properties. In *AISTATS 2020-23rd International Conference on Artificial Intelligence and Statistics*, volume 108, pages 1–20, 2020. (Cited on pages 2 and 26.)
- [16] 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. (Cited on page 1.)
- [17] A. Genevay, G. Peyré, and M. Cuturi. Learning generative models with Sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR, 2018. (Cited on page 2.)
- [18] Z. Goldfeld and K. Greenewald. Sliced mutual information: A scalable measure of statistical dependence. *Advances in Neural Information Processing Systems*, 34, 2021. (Cited on page 6.)

- [19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. (Cited on page 26.)
- [20] S. Helgason. The radon transform on \mathbb{R}^n . In *Integral Geometry and Radon Transforms*, pages 1–62. Springer, 2011. (Cited on page 3.)
- [21] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. (Cited on page 9.)
- [22] N. Ho, X. Nguyen, M. Yurochkin, H. H. Bui, V. Huynh, and D. Phung. Multilevel clustering via Wasserstein means. In *International Conference on Machine Learning*, pages 1501–1509, 2017. (Cited on page 1.)
- [23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. (Cited on page 29.)
- [24] L. B. Klebanov, V. Beneš, and I. Saxl. *N-distances and their applications*. Charles University in Prague, the Karolinum Press Prague, Czech Republic, 2005. (Cited on page 26.)
- [25] S. Kolouri, N. A. Ketz, A. Soltoggio, and P. K. Pilly. Sliced cramer synaptic consolidation for preserving deeply learned representations. In *International Conference on Learning Representations*, 2019. (Cited on page 6.)
- [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. (Cited on pages 9 and 19.)
- [27] S. Kolouri, P. E. Pope, C. E. Martin, and G. K. Rohde. Sliced Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018. (Cited on pages 2 and 5.)
- [28] 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. (Cited on pages 2 and 19.)
- [29] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009. (Cited on page 9.)
- [30] T. Le, T. Nguyen, N. Ho, H. Bui, and D. Phung. Lamda: Label matching deep domain adaptation. In *International Conference on Machine Learning*, pages 6043–6054. PMLR, 2021. (Cited on page 1.)
- [31] Y. Le and X. Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. (Cited on page 9.)
- [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. (Cited on pages 2, 4, and 5.)

- [33] J. Lezama, W. Chen, and Q. Qiu. Run-sort-rerun: Escaping batch size limitations in sliced Wasserstein generative models. In *International Conference on Machine Learning*, pages 6275–6285. PMLR, 2021. (Cited on page 18.)
- [34] Z.-P. Liang and D. C. Munson. Partial radon transforms. *IEEE transactions on image processing*, 6(10):1467–1469, 1997. (Cited on pages 3 and 5.)
- [35] T. Lin, N. Ho, X. Chen, M. Cuturi, and M. I. Jordan. Fixed-support Wasserstein barycenters: Computational hardness and fast algorithm. In *NeurIPS*, pages 5368–5380, 2020. (Cited on page 1.)
- [36] T. Lin, N. Ho, and M. Jordan. On efficient optimal transport: An analysis of greedy and accelerated mirror descent algorithms. In *International Conference on Machine Learning*, pages 3982–3991, 2019. (Cited on page 1.)
- [37] T. Lin, N. Ho, and M. I. Jordan. On the efficiency of the Sinkhorn and Greenkhorn algorithms and their acceleration for optimal transport. *ArXiv Preprint: 1906.01437*, 2019. (Cited on page 1.)
- [38] 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. (Cited on page 18.)
- [39] G. Mena and J. Weed. Statistical bounds for entropic optimal transport: sample complexity and the central limit theorem. In *Advances in Neural Information Processing Systems*, 2019. (Cited on page 2.)
- [40] N. Naderializadeh, J. Comer, R. Andrews, H. Hoffmann, and S. Kolouri. Pooling by sliced-Wasserstein embedding. *Advances in Neural Information Processing Systems*, 34, 2021. (Cited on page 18.)
- [41] 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. (Cited on page 18.)
- [42] K. Nadjahi, A. Durmus, P. E. Jacob, R. Badeau, and U. Simsekli. Fast approximation of the sliced-Wasserstein distance using concentration of random projections. *Advances in Neural Information Processing Systems*, 34, 2021. (Cited on pages 9 and 18.)
- [43] K. Nadjahi, A. Durmus, U. Simsekli, and R. Badeau. Asymptotic guarantees for learning generative models with the sliced-Wasserstein distance. In *Advances in Neural Information Processing Systems*, pages 250–260, 2019. (Cited on page 18.)
- [44] K. Nguyen and N. Ho. Amortized projection optimization for sliced Wasserstein generative models. *Advances in Neural Information Processing Systems*, 2022. (Cited on page 18.)
- [45] K. Nguyen and N. Ho. Revisiting sliced Wasserstein on images: From vectorization to convolution. *Advances in Neural Information Processing Systems*, 2022. (Cited on page 18.)

- [46] K. Nguyen, N. Ho, T. Pham, and H. Bui. Distributional sliced-Wasserstein and applications to generative modeling. In *International Conference on Learning Representations*, 2021. (Cited on pages 1, 2, 5, 9, and 18.)
- [47] K. Nguyen, D. Nguyen, Q. Nguyen, T. Pham, H. Bui, D. Phung, T. Le, and N. Ho. On transportation of mini-batches: A hierarchical approach. In *Proceedings of the 39th International Conference on Machine Learning*, 2022. (Cited on pages 1 and 26.)
- [48] K. Nguyen, D. Nguyen, T. Pham, and N. Ho. Improving mini-batch optimal transport via partial transportation. In *Proceedings of the 39th International Conference on Machine Learning*, 2022. (Cited on pages 1 and 26.)
- [49] K. Nguyen, S. Nguyen, N. Ho, T. Pham, and H. Bui. Improving relational regularized autoencoders with spherical sliced fused Gromov-Wasserstein. In *International Conference on Learning Representations*, 2021. (Cited on pages 2, 5, and 18.)
- [50] T. Nguyen, Q.-H. Pham, T. Le, T. Pham, N. Ho, and B.-S. Hua. Point-set distances for learning representations of 3D point clouds. In *ICCV*, 2021. (Cited on page 4.)
- [51] S. Nietert, Z. Goldfeld, R. Sadhu, and K. Kato. Statistical, robustness, and computational guarantees for sliced wasserstein distances. *Advances in Neural Information Processing Systems*, 2022. (Cited on page 18.)
- [52] 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. (Cited on page 1.)
- [53] G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. (Cited on page 4.)
- [54] G. Peyré and M. Cuturi. Computational optimal transport, 2020. (Cited on page 1.)
- [55] A. Rakotomamonjy and R. Liva. Differentially private sliced Wasserstein distance. In *International Conference on Machine Learning*, pages 8810–8820. PMLR, 2021. (Cited on page 18.)
- [56] L. Rout, A. Korotin, and E. Burnaev. Generative modeling with optimal transport maps. In *International Conference on Learning Representations*, 2022. (Cited on page 1.)
- [57] M. Rowland, J. Hron, Y. Tang, K. Choromanski, T. Sarlos, and A. Weller. Orthogonal estimation of Wasserstein distances. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 186–195. PMLR, 2019. (Cited on page 18.)
- [58] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. *Advances in Neural Information Processing Systems*, 29, 2016. (Cited on page 9.)
- [59] T. Salimans, H. Zhang, A. Radford, and D. Metaxas. Improving GANs using optimal transport. In *International Conference on Learning Representations*, 2018. (Cited on page 26.)

- [60] M. A. Schmitz, M. Heitz, N. Bonneel, F. Ngole, D. Coeurjolly, M. Cuturi, G. Peyré, and J.-L. Starck. Wasserstein dictionary learning: Optimal transport-based unsupervised nonlinear dictionary learning. *SIAM Journal on Imaging Sciences*, 11(1):643–678, 2018. (Cited on page 2.)
- [61] R. Shu. Amortized optimization <http://ruishu.io/2017/11/07/amortized-optimization/>, 2017. (Cited on page 18.)
- [62] Y. Song, S. Garg, J. Shi, and S. Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in Artificial Intelligence*, pages 574–584. PMLR, 2020. (Cited on page 6.)
- [63] V. Titouan, R. Flamary, N. Courty, R. Tavenard, and L. Chapel. Sliced Gromov-Wasserstein. *Advances in Neural Information Processing Systems*, 32, 2019. (Cited on page 6.)
- [64] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018. (Cited on page 1.)
- [65] C. Villani. *Optimal transport: Old and New*. Springer, 2008. (Cited on pages 1 and 4.)
- [66] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019. (Cited on page 25.)
- [67] J. Wu, Z. Huang, D. Acharya, W. Li, J. Thoma, D. P. Paudel, and L. V. Gool. Sliced Wasserstein generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3713–3722, 2019. (Cited on page 2.)
- [68] J. Xi and J. Niles-Weed. Distributional convergence of the sliced wasserstein process. *Advances in Neural Information Processing Systems*, 2022. (Cited on page 18.)
- [69] 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. (Cited on page 1.)
- [70] K. D. Yang, K. Damodaran, S. Venkatachalapathy, A. C. Soylemezoglu, G. Shivashankar, and C. Uhler. Predicting cell lineages using autoencoders and optimal transport. *PLoS computational biology*, 16(4):e1007828, 2020. (Cited on page 1.)
- [71] M. Yi and S. Liu. Sliced Wasserstein variational inference. In *Fourth Symposium on Advances in Approximate Bayesian Inference*, 2021. (Cited on page 18.)

Supplement to “Hierarchical Sliced Wasserstein Distance”

In this supplement, we first discuss some related works in Appendix A. We then present some additional materials including hierarchical Radon transform on multivariate Gaussian and mixture of multivariate Gaussians, max sliced Wasserstein distance, max hierarchical sliced Wasserstein (Max-HSW) distance, their computation, and their theoretical properties in Appendix B. After that, we collect proofs for key results in the paper in Appendix C. In Appendix D, we include the detailed training objectives of our generative modeling framework and some additional experiments including convergence in IS scores of generative models, generated images, and the comparison between the Max-HSW and the Max-SW. Moreover, we report experimental settings including parameter choices and neural network architectures in Appendix E.

A More Related Works

Authors in [46, 49] replace uniform distribution on the projecting directions on the unit hyper-sphere in SW with an estimated distribution that puts high probabilities for discriminative directions. A sliced Wasserstein variant between distributions that have supports on the hyper-sphere is introduced in [6]. One-dimensional transportation plans from orthogonal projecting directions are used for estimating Wasserstein distance in [57]. A fast biased approximation of the SW is proposed in [42]. The Augmented SW is introduced [9]. Statistical guarantees of training generative models with sliced Wasserstein is derived in [43]. A sliced Wasserstein variant between measures over tensors is defined in [45]. A differential private version of the SW is proposed in [55]. The SW is used in generative frameworks such as sliced iterative normalizing flows [12] and fine-tuning pre-trained model [33]. Amortized optimization [61, 2] is used for predicting the optimal slice in sliced Wasserstein generative models in [44]. The SW is used for set representations in [40]. The SW gradient flows are investigated in [38, 7]. Variational inference using the SW is carried out in [71]. Similarly, the SW is used for approximate Bayesian computation in [41].

On the theoretical side, statistical properties of sliced Wasserstein distance in learning generative models are investigated in [43]. Distributional convergence of the Sliced Wasserstein process is derived in [68]. Statistical, robustness, and computational guarantees for sliced Wasserstein distances are shown in [51]. A differential private version of the SW is proposed in [55].

B Additional Materials

Hierarchical Radon Transform as the composition of Partial Radon Transform and Overparametrized Radon Transform: From the definitions of ORT (Definition 4) and PRT

(Definition 3), we have:

$$\begin{aligned}
(\mathcal{PR}(\mathcal{OR}f))(v, \theta_{1:k}, \psi) &= \int_{\mathbb{R}^k} \int_{\mathbb{R}^d} f(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) dx \delta\left(v - \sum_{i=1}^k t_i \psi_i\right) dt \\
&= \int_{\mathbb{R}^k} \int_{\mathbb{R}^d} f(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) \delta\left(v - \sum_{i=1}^k t_i \psi_i\right) dx dt \\
&= \int_{\mathbb{R}^d} \int_{\mathbb{R}^k} f(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) \delta\left(v - \sum_{i=1}^k t_i \psi_i\right) dt dx \\
&= \int_{\mathbb{R}^k} \int_{\mathbb{R}^d} f(x) \delta\left(v - \sum_{i=1}^k \langle x, \theta_i \rangle \psi_i\right) dx \\
&:= (\mathcal{HR}f)(v, \theta_{1:k}, \psi),
\end{aligned}$$

where the second and the third equality are due to the Fubini's theorem.

Hierarchical Radon Transform of Multivariate Gaussian and Mixture of Multivariate Gaussians: Similar to the result of Radon Transform for Multivariate Gaussian and Mixture of Multivariate Gaussians [28], given projecting direction $\Theta = (\theta_{1:k}), \psi$, the one dimensional projected distributions from a Gaussian and a mixture of Gaussians with the HRT are also a Gaussian and a mixture of Gaussians respectively. These results are directly obtained from the linearity of Gaussian distributions. In more detail, let $f := \mathcal{N}(\mu, \Sigma)$, we have $(\mathcal{HR}f)(\cdot, \theta_{1:k}, \psi) = \mathcal{N}(\psi^\top \Theta^\top \mu, \psi^\top \Theta^\top \Sigma \Theta \psi)$. Similarly, let $f := \sum_{i=1}^k w_i \mathcal{N}(\mu_i, \Sigma_i)$, we have $(\mathcal{HR}f)(\cdot, \theta_{1:k}, \psi) = \sum_{i=1}^k w_i \mathcal{N}(\psi^\top \Theta^\top \mu_i, \psi^\top \Theta^\top \Sigma_i \Theta \psi)$.

Max sliced Wasserstein distance: We first review the definition of max sliced Wasserstein distance (Max-SW) [13].

Definition 7. For any $p \geq 1$ and dimension $d \geq 1$, the max sliced Wasserstein- p distance between two probability measures $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ and $\nu \in \mathcal{P}_p(\mathbb{R}^d)$ is given by:

$$\text{Max-SW}_p(\mu, \nu) = \max_{\theta \in \mathbb{S}^{d-1}} W_p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta)).$$

Max-SW is also a valid metric between probability measures [13]. The benefit of Max-SW is that it uses only one projecting direction instead of L projecting directions as SW. Therefore, its projection complexity is only $\mathcal{O}(d)$ compared to $\mathcal{O}(Ld)$ of SW. However, Max-SW requires an iterative optimization procedure to find the max projecting direction on the unit hyper-sphere e.g., projected gradient ascent [26]. We summarize the projected gradient ascent of the Max-SW in Algorithm 1.

Max hierarchical sliced Wasserstein distance: We now define the max hierarchical sliced Wasserstein distance.

Definition 8. For any $p \geq 1$, $k \geq 1$ and dimension $d \geq 1$, the max hierarchical sliced Wasserstein- p distance between two probability measures $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ and $\nu \in \mathcal{P}_p(\mathbb{R}^d)$ is given by:

$$\text{Max-HSW}_{p,k}(\mu, \nu) = \max_{\theta_{1:k} \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{HR}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_\nu)(\cdot, \theta_{1:k}, \psi)).$$

Algorithm 1 Max sliced Wasserstein distance

Input: Probability measures: μ, ν , learning rate η , max number of iterations T .

Initialize θ

while θ not converge or reach T **do**

$$\theta = \theta + \eta \cdot \nabla_{\theta} W_p((\mathcal{R}f_{\mu})(\cdot, \theta), (\mathcal{R}f_{\nu})(\cdot, \theta))$$

$$\theta = \frac{\theta}{\|\theta\|_2}$$

end while

Return: $\theta, W_p((\mathcal{R}f_{\mu})(\cdot, \theta), (\mathcal{R}f_{\nu})(\cdot, \theta))$

Similar to the idea of max sliced Wasserstein, Max-HSW improves the projection complexity of HSW by avoiding using Monte Carlo samples for projecting directions and mixing directions. In particular, Max-HSW finds the best set of bottleneck projections $\theta_{1:k}$ and the best mixing direction ψ in terms of maximizing the discrepancy between two interested measures.

Remark. We have that when $k = 1$, the max hierarchical sliced Wasserstein distance reverts into the max sliced Wasserstein distance.

$$\text{Max-HSW}_{p,1}(\mu, \nu) = \text{Max-SW}_p(\mu, \nu),$$

This is due to the fact that $\psi \in \mathbb{S}^0 = \{1\}$.

Computation of max hierarchical sliced Wasserstein distance: Similar to the Max-SW, Max-HSW can be computed via the projected gradient ascent algorithm. In practice, the gradient of $\theta_{1:k}$ and ψ can be computed by using backpropagation (chain rule) since the hierarchical Radon Transform can be seen as a two-layer neural network. We summarize the projected gradient ascent of the Max-HSW in Algorithm 2.

Algorithm 2 Max hierarchical sliced Wasserstein distance

Input: Probability measures: μ, ν , learning rate η , max number of iterations T .

Initialize θ

while $\theta_{1:k}, \psi$ not converge or reach T **do**

$$\psi = \psi + \eta \cdot \nabla_{\psi} W_p((\mathcal{H}\mathcal{R}f_{\mu})(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_{\nu})(\cdot, \theta_{1:k}, \psi))$$

$$\psi = \frac{\psi}{\|\psi\|_2}$$

for $i = 1$ to k **do**

$$\theta_i = \theta_i + \eta \cdot \nabla_{\theta_i} W_p((\mathcal{H}\mathcal{R}f_{\mu})(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_{\nu})(\cdot, \theta_{1:k}, \psi))$$

$$\theta_i = \frac{\theta_i}{\|\theta_i\|_2}$$

end for

end while

Return: $\theta_{1:k}, \psi, W_p((\mathcal{H}\mathcal{R}f_{\mu})(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_{\nu})(\cdot, \theta_{1:k}, \psi))$

Properties of max hierarchical sliced Wasserstein distance: We first have the following result for the metricity of Max-HSW.

Theorem 2. For any $p \geq 1$ and $k \geq 1$, the hierarchical sliced Wasserstein $\text{Max-HSW}_{p,k}(\cdot, \cdot)$ is a metric on the space of probability measures on \mathbb{R}^d .

Proof of Theorem 2 is given in Appendix C.6. We establish the connection between the HSW, max hierarchical sliced Wasserstein (Max-HSW), max sliced Wasserstein (Max-SW), and Wasserstein distance in Proposition 3. The proof is given in Appendix C.4.

We demonstrate that the max hierarchical sliced Wasserstein does not suffer from the curse of dimensionality for the inference purpose, namely, the sample complexity for the empirical distribution from i.i.d. samples to approximate their underlying distribution is at the order of $\mathcal{O}(n^{-1/2})$.

Proposition 5. *Assume that P is a probability measure supported on compact set of \mathbb{R}^d . Let X_1, X_2, \dots, X_n be i.i.d. samples from P and we denote $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ as the empirical measure of these data. Then, for any $p \geq 1$ and $k \geq 1$, there exists a universal constant $C > 0$ such that*

$$\mathbb{E}[\text{Max-HSW}_{p,k}(P_n, P)] \leq Ck\sqrt{(d+1)\log n/n},$$

where the outer expectation is taken with respect to the data X_1, X_2, \dots, X_n .

Proof of Proposition 5 is similar to the proof of Proposition 4. We refer the reader to Appendix C.5.

C Proofs

In this appendix, we provide proofs for key results in the main text and in Appendix B.

C.1 Proof of Proposition 1

Let us consider functions $f, g \in \mathbb{L}^1(\mathbb{R}^d)$ such that

$$(\mathcal{O}f)(t_{1:k}, \theta_{1:k}) = (\mathcal{O}g)(t_{1:k}, \theta_{1:k}),$$

It is clear from the Definition 4 that

$$\int_{\mathbb{R}^d} f(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) dx = \int_{\mathbb{R}^d} g(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) dx.$$

Taking the integral of both sides in the above equation with respect to $(k-1)$ variables t_2, \dots, t_k , we get

$$\int_{\mathbb{R}^{k-1}} \int_{\mathbb{R}^d} f(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) dx dt_2 \dots dt_k = \int_{\mathbb{R}^{k-1}} \int_{\mathbb{R}^d} g(x) \prod_{i=1}^k \delta(t_i - \langle x, \theta_i \rangle) dx dt_2 \dots dt_k.$$

Next, by applying the Fubini's theorem, we have

$$\int_{\mathbb{R}^d} f(x) \delta(t_1 - \langle x, \theta_1 \rangle) \left(\prod_{i=2}^k \int_{\mathbb{R}} \delta(t_i - \langle x, \theta_i \rangle) dt_i \right) dx = \int_{\mathbb{R}^d} g(x) \delta(t_1 - \langle x, \theta_1 \rangle) \left(\prod_{i=2}^k \int_{\mathbb{R}} \delta(t_i - \langle x, \theta_i \rangle) dt_i \right) dx.$$

Note that for any $i = 2, \dots, k-1$, by using change of variables $s_i = t_i - \langle x, \theta_i \rangle$, we have $\int_{\mathbb{R}} \delta(t_i - \langle x, \theta_i \rangle) dt_i = \int_{\mathbb{R}} \delta(s_i) ds_i = 1$. As a result,

$$\int_{\mathbb{R}^d} f(x) \delta(t_1 - \langle x, \theta_1 \rangle) dx = \int_{\mathbb{R}^d} g(x) \delta(t_1 - \langle x, \theta_1 \rangle) dx,$$

or equivalently, $(\mathcal{R})f(t_1, \theta_1) = \mathcal{R}g(t_1, \theta_1)$. Recall that the Radon transform \mathcal{R} is injective. Thus, we obtain that $f = g$, which completes the proof.

C.2 Proof of Proposition 2

Let us consider arbitrary functions $f, g \in L^1(\mathbb{R}^d)$ satisfying

$$(\mathcal{HR}f)(v, \theta_{1:k}, \psi) = (\mathcal{HR}g)(v, \theta_{1:k}, \psi),$$

where $v \in \mathbb{R}$, $\psi \in \mathbb{S}^{k-1}$ and $\theta_{1:k} \in \mathbb{S}^{d-1}$. It can be seen from the beginning of Appendix B that $(\mathcal{HR}f)(v, \theta_{1:k}, \psi) = (\mathcal{PR}(\mathcal{OR}f))(v, \psi, \theta_{1:k})$ and $(\mathcal{HR}g)(v, \theta_{1:k}, \psi) = (\mathcal{PR}(\mathcal{OR}g))(v, \psi, \theta_{1:k})$. Therefore, we get

$$(\mathcal{PR}(\mathcal{OR}f))(v, \psi, \theta_{1:k}) = (\mathcal{PR}(\mathcal{OR}g))(v, \psi, \theta_{1:k}).$$

Since Partial Radon Transform is injective, we obtain $h_f = h_g$, or equivalently,

$$(\mathcal{OR}f)(t_{1:k}, \theta_{1:k}) = (\mathcal{OR}g)(t_{1:k}, \theta_{1:k}),$$

which leads to $f = g$ due to the injectivity of Overparametrized Radon Transform. Hence, we reach the conclusion of the proposition.

C.3 Proof of Theorem 1

To prove that the hierarchical sliced Wasserstein $HSW_{p,k}(\cdot, \cdot)$ is a metric on the space of all probability measure on \mathbb{R}^d for any $p, k \geq 1$, we need to verify four following criteria:

Symmetry: For any $p, k \geq 1$, it is obvious that $HSW_{p,k}(\mu, \nu) = HSW_{p,k}(\nu, \mu)$ for any probability measures μ and ν .

Non-negativity: The non-negativity of $HSW_{p,k}(\cdot, \cdot)$ comes directly from the non-negativity of the Wasserstein metric.

Identity: For any $p, k \geq 1$, it is clear that when $\mu = \nu$, we have $HSW_{p,k}(\mu, \nu) = 0$. Now, assume that $HSW_{p,k}(\mu, \nu) = 0$, then $W_p((\mathcal{HR}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_\nu)(\cdot, \theta_{1:k}, \psi)) = 0$ for almost all $\psi \in \mathbb{S}^{k-1}$, $\theta_{1:k} \in \mathbb{S}^{d-1}$. By applying the identity property of the Wasserstein distance, we have $(\mathcal{HR}f_\mu)(\cdot, \theta_{1:k}, \psi) = (\mathcal{HR}f_\nu)(\cdot, \theta_{1:k}, \psi)$ for almost all $\psi \in \mathbb{S}^{k-1}$, $\theta_{1:k} \in \mathbb{S}^{d-1}$. Since the Hierarchical Radon Transform is injective, we obtain $f_\mu = f_\nu$, which implies that $\mu = \nu$.

Triangle Inequality: For any probability measures μ_1, μ_2, μ_3 , we find that

$$\begin{aligned} HSW_{p,k}(\mu_1, \mu_3) &= (\mathbb{E}_{\theta_{1:k}, \psi} W_p^p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}, \psi)))^{\frac{1}{p}} \\ &\leq (\mathbb{E}_{\theta_{1:k}, \psi} [W_p^p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi)) \\ &\quad + W_p^p((\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}, \psi))])^{\frac{1}{p}} \\ &\leq (\mathbb{E}_{\theta_{1:k}, \psi} W_p^p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi)))^{\frac{1}{p}} \\ &\quad + (\mathbb{E}_{\theta_{1:k}, \psi} W_p^p((\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}, \psi)))^{\frac{1}{p}} \\ &= HSW_{p,k}(\mu_1, \mu_2) + HSW_{p,k}(\mu_2, \mu_3), \end{aligned}$$

where the first inequality is due to the triangle inequality of Wasserstein metric, namely, we have

$$\begin{aligned} W_p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}, \psi)) &\leq W_p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi)) \\ &\quad + W_p((\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}, \psi)), \end{aligned}$$

while the second inequality is an application of the Minkowski inequality for integrals.

Hence, the hierarchical sliced Wasserstein $HSW_{p,k}(\cdot, \cdot)$ is a metric on the space of all probability measures on \mathbb{R}^d for any $p, k \geq 1$.

C.4 Proof of Proposition 3

The proof of this proposition is direct from the definition of the hierarchical sliced Wasserstein distance, the sliced Wasserstein distance, the max hierarchical sliced Wasserstein distance, and the max sliced Wasserstein distance. Here, we provide the proof for the completeness.

(a) We start with

$$\begin{aligned} HSW_{p,k}(\mu, \nu) &= \left(\mathbb{E}_{\theta_{1:k}, \psi} W_p^p((\mathcal{H}\mathcal{R}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_\nu)(\cdot, \theta_{1:k}, \psi)) \right)^{\frac{1}{p}} \\ &\leq \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{H}\mathcal{R}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_\nu)(\cdot, \theta_{1:k}, \psi)), \\ &:= \text{Max-}HSW_{p,k}(\mu, \nu). \end{aligned}$$

Subsequently, we have

$$\begin{aligned} \text{Max-}HSW_{p,k}(\mu, \nu) &= \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{H}\mathcal{R}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_\nu)(\cdot, \theta_1, \dots, \theta_k, \psi)) \\ &= \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \left| \psi^\top \Theta^\top x - \psi^\top \Theta^\top y \right|^p d\pi(x, y) \right)^{\frac{1}{p}} \\ &= \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \left| \sum_{j=1}^k \psi_j \theta_j^\top (x - y) \right|^p d\pi(x, y) \right)^{\frac{1}{p}} \\ &\leq \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\psi\|^p \left| \sum_{j=1}^k \theta_j^\top (x - y) \right|^p d\pi(x, y) \right)^{\frac{1}{p}} \\ &= \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \left| \sum_{j=1}^k \theta_j^\top (x - y) \right|^p d\pi(x, y) \right)^{\frac{1}{p}} \\ &\leq \max_{\theta \in \mathbb{S}^{d-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} k^p \left| \theta^\top x - \theta^\top y \right|^p d\pi(x, y) \right)^{\frac{1}{p}} \\ &= k \cdot \max_{\theta \in \mathbb{S}^{d-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \left| \theta^\top x - \theta^\top y \right|^p d\pi(x, y) \right)^{\frac{1}{p}} \\ &= k \cdot \max_{\theta \in \mathbb{S}^{d-1}} W_p(\theta \# \mu, \theta \# \nu) \\ &= k \cdot \max_{\theta \in \mathbb{S}^{d-1}} W_p((\mathcal{R})f_\mu(\cdot, \theta), (\mathcal{R})f_\nu(\cdot, \theta)) \\ &= k \cdot \text{Max-SW}_p(\mu, \nu). \end{aligned}$$

Finally, by applying the Cauchy-Schwartz inequality, we get

$$\begin{aligned}
\text{Max-SW}_p^p(\mu, \nu) &= \max_{\theta \in \mathbb{S}^{d-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d} |\theta^\top x - \theta^\top y|^p d\pi(x, y) \right) \\
&\leq \max_{\theta \in \mathbb{S}^{d-1}} \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\theta\|^p \|x - y\|^p d\pi(x, y) \right) \\
&= \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\theta\|^p \|x - y\|^p d\pi(x, y) \\
&= W_p^p(\mu, \nu).
\end{aligned}$$

Putting the above results together, we obtain the conclusion of the proposition.

(b) Firstly, it is obvious that

$$SW_p(\mu, \nu) = \left(\mathbb{E}_{\theta \sim \mathcal{U}(\mathbb{S}^{d-1})} W_p^p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta)) \right)^{\frac{1}{p}} \leq \max_{\theta \in \mathbb{S}^{d-1}} W_p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta)).$$

Subsequently, let $\theta'_1 = \theta^* = \arg \max_{\theta \in \mathbb{S}^{d-1}} W_p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta))$ and $\psi' = (1, 0, \dots, 0) \in \mathbb{S}^{k-1}$. Then, for any $\theta'_1 \in \mathbb{S}^{d-1}$, we have

$$\begin{aligned}
\text{Max-HSW}_{p,k}(\mu, \nu) &= \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{H}\mathcal{R}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{H}\mathcal{R}f_\nu)(\cdot, \theta_{1:k}, \psi)) \\
&\geq W_p((\mathcal{H}\mathcal{R}f_\mu)(\cdot, \theta'_1, \theta_2, \dots, \theta_k, \psi'), (\mathcal{H}\mathcal{R}f_\nu)(\cdot, \theta'_1, \theta_2, \dots, \theta_k, \psi')) \\
&= \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d} |\theta_1^\top x - \theta_1^\top y|^p d\pi(x, y) \right)^{\frac{1}{p}} \\
&= W_p((\mathcal{R}f_\mu)(\cdot, \theta'_1), (\mathcal{R}f_\nu)(\cdot, \theta'_1)).
\end{aligned}$$

Since the above result holds for all arbitrary $\theta'_1 \in \mathbb{S}^{d-1}$, we obtain

$$\text{Max-HSW}_{p,k}(\mu, \nu) \geq \max_{\theta \in \mathbb{S}^{d-1}} W_p((\mathcal{R}f_\mu)(\cdot, \theta), (\mathcal{R}f_\nu)(\cdot, \theta)) = \text{Max-SW}_p(\mu, \nu),$$

which completes the proof.

C.5 Proof of Proposition 4

From the assumption of Proposition 4, we denote $\Theta \subset \mathbb{R}^d$ as the compact set that the probability measure P is supported on. Based on the result of Proposition 3, we have

$$\mathbb{E}[HSW_{p,k}(P_n, P)] \leq \mathbb{E}[k \cdot \text{Max-SW}_p(P_n, P)],$$

where $\text{Max-SW}_p(P_n, P) := \max_{\theta \in \mathbb{S}^{d-1}} W_p((\mathcal{R}f_{P_n})(\cdot, \theta), (\mathcal{R}f_P)(\cdot, \theta)) := \max_{\theta \in \mathbb{S}^{d-1}} W_p(\theta^\top P_n, \theta^\top P)$. Therefore, to obtain the conclusion of the proposition, it is sufficient to demonstrate that $\mathbb{E}[\text{Max-SW}_p(P_n, P)] \leq C \sqrt{(d+1) \log_2 n/n}$ for some universal constant $C > 0$. Indeed, from the

closed-form expression of Wasserstein metric in one dimension, we have

$$\begin{aligned}
\text{Max-SW}_p^p(P_n, P) &= \max_{\theta \in \mathbb{S}^{d-1}} \int_0^1 |F_{n,\theta}^{-1}(u) - F_\theta^{-1}(u)|^p du \\
&= \max_{\theta \in \mathbb{R}^d: \|\theta\|=1} \int_0^1 |F_{n,\theta}^{-1}(u) - F_\theta^{-1}(u)|^p du \\
&\leq \max_{\theta \in \mathbb{R}^d: \|\theta\| \leq 1} \int_{\mathbb{R}} |F_{n,\theta}(x) - F_\theta(x)|^p dx \\
&\leq \text{diam}(\Theta) \max_{\theta \in \mathbb{R}^d: \|\theta\| \leq 1} |F_{n,\theta}(x) - F_\theta(x)|^p,
\end{aligned}$$

where $F_{n,\theta}$ and F_θ are respectively the cumulative distributions of $\theta^\top P_n$ and $\theta^\top P$. Furthermore, we have the following relation:

$$\max_{\theta \in \mathbb{R}^d: \|\theta\| \leq 1} |F_{n,\theta}(x) - F_\theta(x)| = \sup_{A \in \mathcal{A}} |P_n(A) - P(A)|,$$

where \mathcal{A} is the set of half-spaces $\{y \in \mathbb{R}^d : \theta^\top y \leq x\}$ for all $\theta \in \mathbb{R}^d$ such that $\|\theta\| \leq 1$. The Vapnik-Chervonenkis (VC) dimension of \mathcal{A} is upper bounded by $d + 1$ (see the reference [66]). Therefore, with probability at least $1 - \delta$, we obtain that

$$\sup_{A \in \mathcal{A}} |P_n(A) - P(A)| \leq \sqrt{\frac{32}{n} [(d + 1) \log_2(n + 1) + \log_2(8/\delta)]}.$$

Putting the above results together, we can conclude that

$$\mathbb{E}[\text{Max-SW}_p(P_n, P)] \leq C \sqrt{(d + 1) \log_2 n/n},$$

for some universal constant $C > 0$. As a consequence, we obtain the conclusion of the proposition.

C.6 Proof of Theorem 2

Symmetry: For any $p \geq 1$, it is clear that $\text{Max-HSW}_p(\mu, \nu) = \text{Max-HSW}_p(\nu, \mu)$ for any probability measures μ and ν .

Non-negativity: The non-negativity of $\text{Max-HSW}_{p,k}(\cdot, \cdot)$ comes directly from the non-negativity of the Wasserstein metric.

Existence of the max directions: The unit hyperspheres \mathbb{S}^{d-1} and \mathbb{S}^{k-1} are compact and we have $W_p((\mathcal{HR}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_\nu)(\cdot, \theta_{1:k}, \psi))$ is continuous in terms of $\theta_{1:k}$ and ψ . Therefore, there exists

$$\theta_{1:k}^*, \psi^* = \arg \max_{\theta_{1:k} \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{HR}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_\nu)(\cdot, \theta_{1:k}, \psi))$$

Identity: For any $p \geq 1$ and $k \geq 1$, it is clear that when $\mu = \nu$, then $\text{Max-HSW}_p(\mu, \nu) = 0$. When $\text{Max-HSW}_p(\mu, \nu) = 0$, we have $W_p((\mathcal{HR}f_\mu)(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_\nu)(\cdot, \theta_{1:k}, \psi)) = 0$ for almost all $\psi \in \mathbb{S}^{k-1}$, $\theta_{1:k} \in \mathbb{S}^{d-1}$. Applying the identity property of the Wasserstein distance, we have $(\mathcal{HR}f_\mu)(\cdot, \theta_{1:k}, \psi) = (\mathcal{HR}f_\nu)(\cdot, \theta_{1:k}, \psi)$ almost all $\psi \in \mathbb{S}^{k-1}$, $\theta_{1:k} \in \mathbb{S}^{d-1}$. Since the Hierarchical Radon Transform is injective, we obtain $\mu = \nu$.

Triangle Inequality: For any probability measures μ_1, μ_2, μ_3 , we find that

$$\begin{aligned}
\text{Max-HSW}_{p,k}(\mu_1, \mu_3) &= \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}, \psi)) \\
&= W_p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}^*, \psi^*), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}^*, \psi^*)) \\
&\leq W_p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}^*, \psi^*), (\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}^*, \psi^*)) \\
&\quad + W_p((\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}^*, \psi^*), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}^*, \psi^*)) \\
&\leq \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{HR}f_{\mu_1})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi)) \\
&\quad + \max_{\theta_1, \dots, \theta_k \in \mathbb{S}^{d-1}, \psi \in \mathbb{S}^{k-1}} W_p((\mathcal{HR}f_{\mu_2})(\cdot, \theta_{1:k}, \psi), (\mathcal{HR}f_{\mu_3})(\cdot, \theta_{1:k}, \psi)) \\
&= \text{Max-HSW}_{p,k}(\mu_1, \mu_2) + \text{Max-HSW}_{p,k}(\mu_2, \mu_3)
\end{aligned}$$

where the first equality is due to the existence of the max directions, the first inequality is due to the triangle inequality with Wasserstein metric, and the second inequality is an application of the definition of maxima.

D Deep Generative Models

In this section, we first review the standard setup of training generative models with sliced Wasserstein variants [14]. We then provide additional experimental results including convergence speed of the SW and the HSW in IS score, generated images from the SW and the HSW, and comparison between the Max-SW and the Max-HSW.

D.1 Training Detail

We follow the standard setup in [14]. We parameterize the model distribution $p_\phi(x) \in \mathcal{P}(\mathbb{R}^d)$ and $p_\phi(x) = G_\phi \# \epsilon$ where ϵ is the standard multivariate Gaussian of 128 dimension and G_ϕ is a neural network with Resnet architecture [19]. Since the ground truth metric between images is unknown, we need a discriminator as a type of ground metric learning. We denote the discriminator as a function $T_{\beta_2} \circ T_{\beta_1}$ where $T_{\beta_1} : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ and $T_{\beta_2} : \mathbb{R}^{d'} \rightarrow \mathbb{R}$. In greater detail, T_{β_1} maps the original images to their corresponding features maps and T_{β_2} maps their features maps to their corresponding discriminative scores. Let the data distribution is μ , our training objectives are:

$$\begin{aligned}
&\min_{\beta_1, \beta_2} (\mathbb{E}_{x \sim \mu} [\min(0, -1 + T_{\beta_2}(T_{\beta_1}(x)))] + \mathbb{E}_{z \sim \epsilon} [\min(0, -1 - T_{\beta_2}(T_{\beta_1}(G_\phi(z))))]), \\
&\min_{\phi} \mathbb{E}_{X \sim \mu^{\otimes m}, Y \sim \epsilon^{\otimes m}} \mathcal{D}(T_{\beta_1} \# P_X, T_{\beta_1} \# G_\phi \# P_Y),
\end{aligned}$$

where $m \geq 1$ is the mini-batch size and $\mathcal{D}(\cdot, \cdot)$ is the (Max-)SW distance or (Max-)HSW distance.

Relation to m -mini-batch energy distance: The objective for training the generator can be seen as the application of mini-batch optimal transport [15, 48, 47] with sliced Wasserstein variants kernel. It can be also seen as a type of m -mini-batch energy distance [24, 59].

D.2 Additional Results

Convergence of generative models from SW and HSW in IS scores: We plots the IS scores over training epochs of the SW and the HSW with the same setting as in Table 1 in Figure 5. We

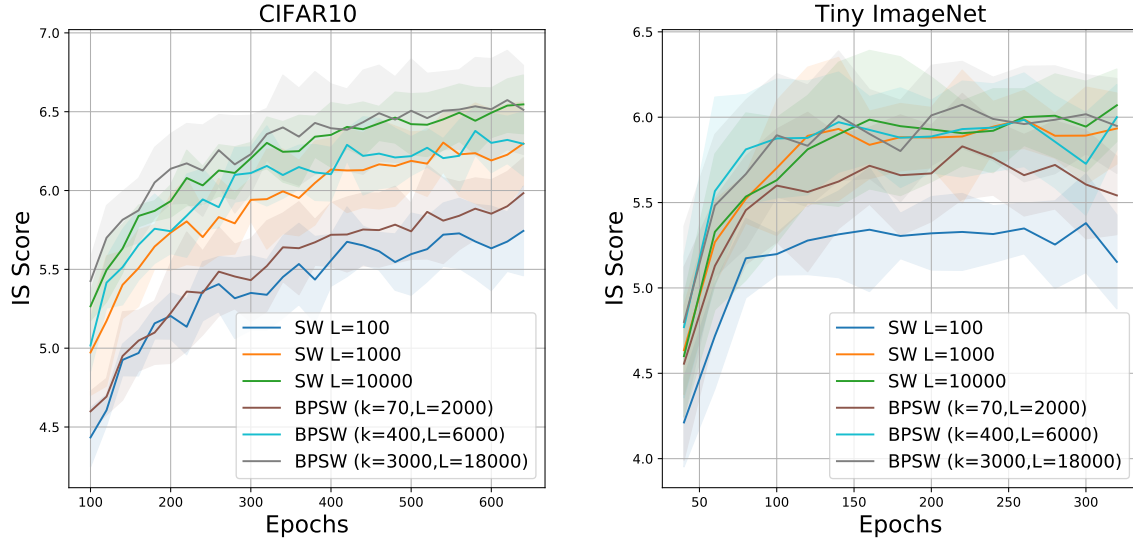


Figure 5: The IS scores over epochs of different training losses on datasets. We observe that HSW helps the generative models converge faster.

observe that IS scores from models trained by the HSW (with better computation) increases faster than ones trained from the SW. We would like to recall that reason is that the HSW has a higher number of final projections than the SW, hence, it is more discriminative signal than the SW.

Radom generated images: We show random generated images from CIFAR10 and Tiny Imagenet from the SW and the HSW generative models in Figure 6- 7. From these images, we observe that increasing L enhance the performance of the SW and increasing both L and k yield better images for the HSW. Also, the qualitative comparison from generated images between the SW and the HSW is consistent to the FID scores and IS scores in Table 1.

Comparison between the Max-HSW and the Max-SW : We would like to recall that Max-HSW is the generalization of the Max-SW since Max-HSW with $k = 1$ is equivalent to the Max-SW. Therefore, the performance of Max-HSW is a least the same as the Max-SW. To find out the benefit of $k > 1$ in Max-HSW, we run Max-HSW with $k \in \{1, 10, 100, 1000\}$, the slice learning rate $\eta \in \{0.001, 0.01, 0.1\}$, and the maximum number of iteration $T \in \{1, 10, 100\}$. We report the best FID scores and IS scores in Table 3. We observe that Max-HSW $k > 1$ gives a better FID score than the Max-SW on CelebA dataset. The reason might be due to the overparametrization of the Max-HSW that leads to a better final slice. However, on CIFAR10 and Tiny ImageNet, the Max-HSW $k > 1$ does not show any improvement compared to the Max-HSW $k = 1$ (Max-SW). This might be because the projected gradient ascent does not work well in a multiplayer structure like in HRT. We will leave the investigation about optimization of the Max-HSW in future works since the Max-HSW has the potential to explain partially adversarial training frameworks.

E Experimental Settings

Architectures of neural network: We illustrate the detail of neural network architectures including the generative networks and the discriminative networks on CIFAR10 in Table 4, CelebA

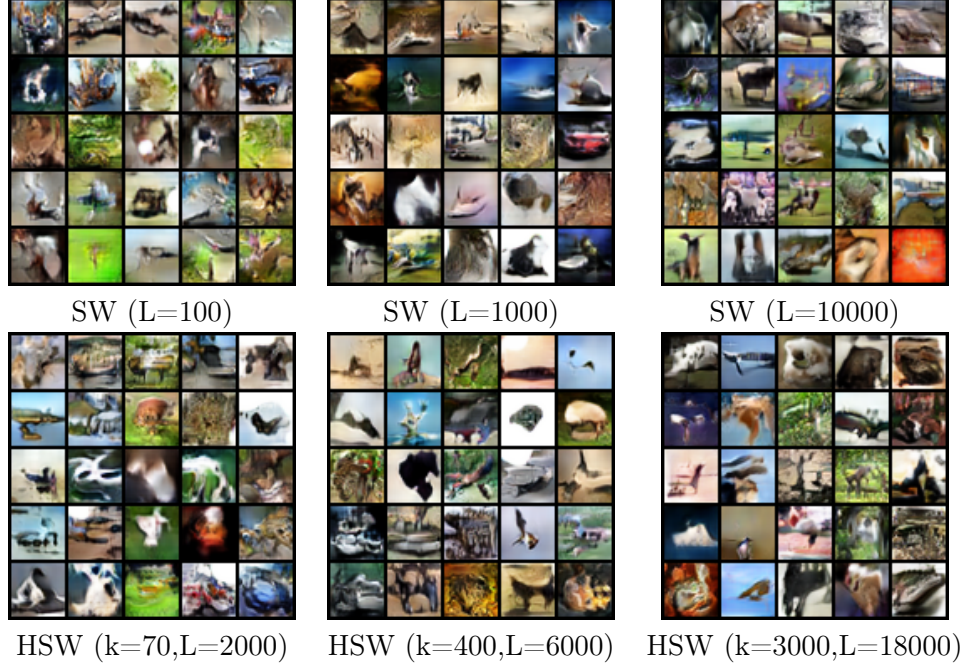


Figure 6: Random generated images of the SW and the HSW on CIFAR10.



Figure 7: Random generated images of the SW and the HSW on Tiny ImageNet.

in Table 5, and Tiny ImageNet in Table 6.

Other settings: We set the number of training iterations to 50000 for all datasets. For each 5 iterations, we update the generator G_ϕ by the corresponding SW and CSW variants. For the

Table 3: FID scores and IS scores the Max-SW and the Max-HSW of on CIFAR10 (32x32), CelebA (64x64), and Tiny ImageNet (64x64).

Method	CIFAR10		CelebA	Tiny ImageNet	
	FID (\downarrow)	IS (\uparrow)	FID (\downarrow)	FID (\downarrow)	IS (\uparrow)
Max-SW (Max-HSW k=1)	43.67 \pm 2.34	6.49 \pm 0.11	17.17 \pm 1.72	82.47 \pm 5.73	6.03 \pm 0.52
Max-HSW	43.67 \pm 2.34	6.49 \pm 0.11	15.92\pm0.87	82.47 \pm 5.73	6.03 \pm 0.52

Table 4: CIFAR10 architectures.

(a) G_ϕ	(b) T_{β_1}	(c) T_{β_2}
Input: $\epsilon \in \mathbb{R}^{128} \sim \mathcal{N}(0, 1)$	Input: $\mathbf{x} \in [-1, 1]^{32 \times 32 \times 3}$	Input: $\mathbf{x} \in \mathbb{R}^{128 \times 8 \times 8}$
128 \rightarrow 4 \times 4 \times 256, dense linear	ResBlock down 128	ReLU
ResBlock up 256	ResBlock down 128	Global sum pooling
ResBlock up 256	ResBlock down 128	128 \rightarrow 1
ResBlock up 256	ResBlock 128	Spectral normalization
BN, ReLU, 3 \times 3 conv, 3 Tanh	ResBlock 128	

discriminator, we update T_{β_1} and T_{β_2} every iterations. We set the mini-batch size m to 128 on CIFAR10, CelebA, and Tiny ImageNet. The learning rate of G_ϕ , T_{β_1} , and T_{β_2} is set to 0.0002. We use Adam [23] for optimization problems with $(\beta_1, \beta_2) = (0, 0.9)$. We use $p = 2$ for both the SW and the HSW.

Calculation of scores: For the FID scores and the Inception scores, we calculate them based on 50000 random samples from trained models³. For the FID scores, we calculate the statistics of datasets on all training samples. Our calculation co

³We use the scores calculation from <https://github.com/GongXinyuu/sngan.pytorch>

Table 5: CelebA architectures.

(a) G_ϕ	(b) T_{β_1}	(c) T_{β_2}
Input: $\epsilon \in \mathbb{R}^{128} \sim \mathcal{N}(0, 1)$	Input: $\mathbf{x} \in [-1, 1]^{64 \times 64 \times 3}$	Input: $\mathbf{x} \in \mathbb{R}^{128 \times 8 \times 8}$
$128 \rightarrow 4 \times 4 \times 256$, dense linear	ResBlock down 128	ReLU
ResBlock up 256	ResBlock down 128	Global sum pooling
ResBlock up 256	ResBlock down 128	$128 \rightarrow 1$
ResBlock up 256	ResBlock 128	Spectral normalization
ResBlock up 256	ResBlock 128	
ResBlock up 256	ResBlock 128	
BN, ReLU, 3×3 conv, 3 Tanh		

Table 6: Tiny ImageNet architectures.

(a) G_ϕ	(b) T_{β_1}	(c) T_{β_2}
Input: $\epsilon \in \mathbb{R}^{128} \sim \mathcal{N}(0, 1)$	Input: $\mathbf{x} \in [-1, 1]^{64 \times 64 \times 3}$	Input: $\mathbf{x} \in \mathbb{R}^{128 \times 8 \times 8}$
$128 \rightarrow 4 \times 4 \times 256$, dense linear	ResBlock down 128	ReLU
ResBlock up 256	ResBlock down 128	Global sum pooling
ResBlock up 256	ResBlock down 128	$128 \rightarrow 1$
ResBlock up 256	ResBlock 128	Spectral normalization
ResBlock up 256	ResBlock 128	
ResBlock up 256	ResBlock 128	
BN, ReLU, 3×3 conv, 3 Tanh		