Revisiting Sliced Wasserstein on Images: From Vectorization to Convolution

Anonymous Author(s)
Affiliation
Address
email

Abstract

The conventional sliced Wasserstein is defined between two probability measures that have realizations as *vectors*. When comparing two probability measures over images, practitioners first need to vectorize images and then project them to one-dimensional space by using matrix multiplication between the sample matrix and the projection matrix. After that, the sliced Wasserstein is evaluated by averaging the two corresponding one-dimensional projected probability measures. However, this approach has two limitations. The first limitation is that the spatial structure of images is not captured efficiently by the vectorization step; therefore, the later slicing process becomes harder to gather the discrepancy information. The second limitation is memory inefficiency since each slicing direction is a vector that has the same dimension as the images. To address these limitations, we propose novel slicing methods for sliced Wasserstein between probability measures over images that are based on the convolution operators. We derive *convolution sliced Wasserstein* (CSW) and its variants via incorporating stride, dilation, and non-linear activation function into the convolution operators. We investigate the metricity of CSW as well as its sample complexity, its computational complexity, and its connection to conventional sliced Wasserstein distances. Finally, we demonstrate the favorable performance of CSW over the conventional sliced Wasserstein in comparing probability measures over images and in training deep generative modeling on images.

1 Introduction

Optimal transport and Wasserstein distance [55, 47] have become popular tools in machine learning and data science. For example, optimal transport has been utilized in generative modeling tasks to generate realistic images [2, 54], in domain adaptation applications to transfer knowledge from source to target domains [9, 3], in clustering applications to capture the heterogeneity of data [21], and in other applications [29, 58, 59]. Despite having appealing performance, Wasserstein distance has been known to suffer from high computational complexity, namely, its computational complexity is at the order of $\mathcal{O}(m^3 \log m)$ [45] when the probability measures have at most m supports. In addition, Wasserstein distance also suffers from the curse of dimensionality, namely, its sample complexity is at the order of $\mathcal{O}(n^{-1/d})$ [15] where n is the sample size. A popular line of work to improve the speed of computation and the sample complexity of the Wasserstein distance is by adding an entropic regularization term to the Wasserstein distance [10]. This variant is known as entropic regularized optimal transport (or equivalently entropic regularized Wasserstein). By using the entropic version, we can approximate the value of Wasserstein distance with the computational complexities

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

being at the order of $\mathcal{O}(m^2)$ [1, 35, 36, 34] (up to some polynomial orders of 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 useful line of work to improve both the computational and sample complexities of the Wasserstein distance is based on the closed-form solution of optimal transport in one dimension. A notable distance along this direction is sliced Wasserstein (SW) distance [6]. Due to the fast computational complexity $\mathcal{O}(m \log_2 m)$ and no curse of dimensionality $\mathcal{O}(n^{-1/2})$, the sliced Wasserstein has been applied successfully in several applications, such as generative modeling [57, 13, 25], domain adaptation [31], and clustering [26]. The sliced Wasserstein is defined between two probability measures that have supports belonging to a vector space, e.g, $\mathbb{R}^d$. As defined in [6], the sliced Wasserstein 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 the sliced Wasserstein distance. The number of samples is often called the number of projections and it is denoted as L. On the computational side, the computation of sliced Wasserstein can be decomposed into two steps. In the first step, L projecting directions are first sampled and then stacked as a matrix (the projection matrix). After that, the projection matrix is multiplied by the two data matrices resulting in two matrices that represent L one-dimensional projected probability measures. In the second step, L one-dimensional Wasserstein distances are computed between the two corresponding projected measures with the same projecting direction. Finally, the average of those distances is yielded as the value of the sliced Wasserstein.

Despite being applied widely in tasks that deal with probability measures over images [57, 13], the conventional formulation of sliced Wasserstein is not well-defined to the nature of images. In particular, an image is not a vector but is a tensor. Therefore, a probability measure over images should be defined over the space of tensors instead of vectors. The conventional formulation leads to an extra step in using the sliced Wasserstein on the domain of images which is vectorization. Namely, all images (supports of two probability measures) are transformed into vectors by a deterministic one-one mapping which is the "reshape" operator. This extra step does not keep the spatial structures of the supports, which are crucial information of images. Furthermore, the vectorization step also poses certain challenges to design efficient ways of projecting (slicing) samples to one dimension based on prior knowledge about the domain of samples. Finally, prior empirical investigations indicate that there are several slices in the conventional Wasserstein collapsing the two probability measures to the Dirac Delta at zero [13, 12, 24]. Therefore, these slices do not contribute to the overall discrepancy. These works suggest that the space of projecting directions in the conventional sliced Wasserstein (the unit hyper-sphere) is potentially not optimal, at least for images.

Contribution. To address these issues of the sliced Wasserstein over images, we propose to replace the conventional formulation of the sliced Wasserstein with a new formulation that is defined on the space of probability measures over tensors. Moreover, we also propose a novel slicing process by changing the conventional matrix multiplication to the convolution operators [16, 18]. In summary, our main contributions are two-fold: **1.** We leverage the benefits of the convolution operators on images, including their efficient parameter sharing and memory saving as well as their superior performance in several tasks on images [28, 19], to introduce efficient slicing methods on sliced Wasserstein, named *convolution slicers*. With the convolution slicers, we derive a novel variant of sliced Wasserstein, named *convolution sliced Wasserstein* (CSW). We investigate the metricity of CSW, its sample and computational complexities, and its connection to other variants of sliced Wasserstein. **2.** We then illustrate the favorable performance of CSW in comparing probability measures over images. In particular, we show that CSW provides an almost identical discrepancy between MNIST's digits compared to that of the SW while having much less slicing memory. Furthermore, we compare SW and CSW in training deep generative models on standard benchmark image datasets, including CIFAR10, CelebA, STL10, and CeleA-HQ. By considering the quality of the trained generative models, training speed, and training memory of CSW and SW, we observe that the CSW has more favorable performance than the vanilla SW.

Organization. The remainder of the paper is organized as follows. We first provide background about Wasserstein distance, the conventional slicing process in the sliced Wasserstein distance, and the convolution operator in Section 2. In Section 3, we propose the convolution slicing and the convolution sliced Wasserstein, and analyze some of its theoretical properties. Section 4 contains the application of CSW to generative models, qualitative experimental results, and quantitative experimental results on standard benchmarks. We conclude the paper In Section 5. Finally, we defer the proofs of key results and extra materials in the Appendices.

Notation. For any $d \geq 2$, $\mathbb{S}^{d-1} := \{\theta \in \mathbb{R}^d \mid ||\theta||_2^2 = 1\}$ denotes the d dimensional unit hyper-sphere in $\mathcal{L}_2$ norm, and $\mathcal{U}(\mathbb{S}^{d-1})$ is the uniform measure over $\mathbb{S}^{d-1}$. Moreover, δ denotes the Dirac delta function. For $p \geq 1$, $\mathcal{P}_p(\mathbb{R}^d)$ is the set of all probability measures on $\mathbb{R}^d$ that have finite p-moments. For $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, $\Pi(\mu, \nu) := \{\pi \in \mathcal{P}_p(\mathbb{R}^d \times \mathbb{R}^d) \mid \int_{\mathbb{R}^d} \pi(x, y)dx = \nu, \int_{\mathbb{R}^d} \pi(x, y)dy = \mu\}$ is the set of transportation plans between μ and ν. For $m \geq 1$, we denotes $\mu^{\otimes m}$ as the product measure which has the supports are the joint vector of m random variables that follows μ. For a vector $X \in \mathbb{R}^{dm}$, $X := (x_1, \ldots, x_m)$, P_X denotes the empirical measures $\frac{1}{m}\sum_{i=1}^m \delta_{x_i}$. For any two sequences a_n and b_n, the notation $a_n = \mathcal{O}(b_n)$ means that $a_n \leq Cb_n$ for all $n \geq 1$ where C is some universal constant.

2 Background

In this section, we first review the definitions of the Wasserstein distance, the conventional slicing, and the sliced Wasserstein distance, and discuss its limitation. We then review the convolution and the padding operators on images.

Sliced Wasserstein: For any $p \geq 1$ and dimension $d' \geq 1$, we first define the Wasserstein-p distance [55, 46] between two probability measures $\mu \in \mathcal{P}_p(\mathbb{R}^{d'})$ and $\nu \in \mathcal{P}_p(\mathbb{R}^{d'})$, which is given by $\mathrm{W}_p(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathbb{R}^{d'} \times \mathbb{R}^{d'}} \|x - y\|_p^p d\pi(x, y)\right)^{\frac{1}{p}}$. When $d' = 1$, the Wasserstein distance has a closed form which is $W_p(\mu, \nu) = (\int_0^1 |F_\mu^{-1}(z) - F_\nu^{-1}(z)|^p dz)^{1/p}$ where F_μ and F_ν are the cumulative distribution function (CDF) of μ and ν respectively.

Given this closed-form property of Wasserstein distance in one dimension, the sliced Wasserstein distance [6] between μ and ν had been introduced and admitted the following formulation: $\mathrm{SW}_p^p(\mu, \nu) := \int_{\mathbb{S}^{d-1}} \mathrm{W}_p^p(\theta\sharp\mu, \theta\sharp\nu)d\theta$, where $\theta\sharp\mu$ is the push-forward probability measure of μ through the function $T_\theta : \mathbb{R}^{d'} \to \mathbb{R}$ with $T_\theta(x) = \theta^\top x$. For each $\theta \in \mathbb{S}^{d'-1}$, $\mathrm{W}_p^p(\theta\sharp\mu, \theta\sharp\nu)$ can be computed in linear time $\mathcal{O}(m \log_2 m)$ where m is the number of supports of μ and ν. However, the integration over the unit sphere in the sliced Wasserstein distance is intractable to compute. Therefore, Monte Carlo scheme is employed to approximate the integration, namely, $\theta_1, \ldots, \theta_L \sim \mathcal{U}(\mathbb{S}^{d'-1})$ are drawn uniformly from the unit sphere and the approximation of the sliced Wasserstein distance is given by: $\widehat{\mathrm{SW}}_p^p(\mu, \nu) \approx \frac{1}{L}\sum_{i=1}^L \mathrm{W}_p^p(\theta_i\sharp\mu, \theta_i\sharp\nu)$. In practice, L should be chosen to be sufficiently large compared to the dimension d', which can be undesirable.

Sliced Wasserstein on Images: Now, we focus on two probability measures over images: $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^{c\times d\times d})$ for number of channels $c \geq 1$ and dimension $d \geq 1$. In this case, the sliced Wasserstein between μ and ν is defined as:

$$\mathrm{SW}_p(\mu, \nu) = \mathrm{SW}_p(\mathcal{R}\sharp\mu, \mathcal{R}\sharp\nu), \tag{1}$$

where $\mathcal{R} : \mathbb{R}^{c\times d\times d} \to \mathbb{R}^{cd^2}$ is a deterministic one-to-one "reshape" mapping.

The slicing process: The slicing of sliced Wasserstein distance on probability measures over images consists of two steps: vectorization and projection. Suppose that the probability measure $\mu \in \mathcal{P}(\mathbb{R}^{c\times d\times d})$ has n supports. Then the supports of μ are transformed into vectors in $\mathbb{R}^{cd^2}$ and are stacked as a matrix of size $n \times cd^2$. A projection matrix of size $L \times cd^2$ is then sampled and has each column as a random vector following the uniform measure over the unit hyper-sphere. Finally, the multiplication of those two matrices returns L projected probability measures of n supports in one dimension. We illustrate this process in Figure 3 in Appendix C.

Limitation of the conventional slicing: First of all, images contain spatial relations across channels and local information. Therefore, transforming images into vectors makes it challenging to obtain

that information. Second, vectorization leads to the usage of projecting directions from the unit hyper-sphere, which can have several directions that do not have good discriminative power. Finally, sampling projecting directions in high-dimension is also time-consuming and memory-consuming. As a consequence, avoiding the vectorization step can improve the efficiency of the whole process.

Convolution operator: We now define the convolution operator on tensors [16], which will be used as an alternative way of projecting images to one dimension in the sliced Wasserstein. The definition of the convolution operator with stride and dilation is as follows.

Definition 1 *(Convolution) Given the number of channels $c \geq 1$, the dimension $d \geq 1$, the stride size $s \geq 1$, the dilation size $b \geq 1$, the size of kernel $k \geq 1$, the convolution of a tensor $X \in \mathbb{R}^{c\times d\times d}$ with a kernel size $K \in \mathbb{R}^{c\times k\times k}$ is $X \overset{s,b}{*} K = Y$, $Y \in \mathbb{R}^{1\times d'\times d'}$ where $d' = \frac{d-b(k-1)-1}{s} + 1$. For $i = 1,\ldots,d'$ and $j = 1,\ldots,d'$, $Y_{1,i,j}$ is defined as: $Y_{1,i,j} = \sum_{h=1}^{c}\sum_{i'=0}^{k-1}\sum_{j'=0}^{k-1} X_{h,s(i-1)+bi'+1,s(j-1)+bj'+1} \cdot K_{h,i'+1,j'+1}$.*

From its definition, we can check that the computational complexity of the convolution operator is $\mathcal{O}\left(c\left(\frac{d-b(k-1)-1}{s}+1\right)^2 k^2\right)$.

3 Convolution Sliced Wasserstein

In this section, we will define a convolution slicer that maps a tensor to a scalar by convolution operators. Moreover, we discuss the convolution slicer and some of its specific forms including the convolution-base slicer, the convolution-stride slicer, the convolution-dilation slicer, and their non-linear extensions. After that, we derive the convolution sliced Wasserstein (CSW), a family of variants of sliced Wasserstein, that utilizes a convolution slicer as the projecting method. Finally, we discuss some theoretical properties of CSW, namely, its metricity, its computational complexity, its sample complexity, and its connection to other variants of sliced Wasserstein.

3.1 Convolution Slicer

We first start with the definition of the convolution slicer, which plays an important role in defining convolution sliced Wasserstein.

Definition 2 *(Convolution Slicer) For $N \geq 1$, given a sequence of kernels $K^{(1)} \in \mathbb{R}^{c^{(1)}\times d^{(1)}\times d^{(1)}},\ldots,K^{(N)} \in \mathbb{R}^{c^{(N)}\times d^{(N)}\times d^{(N)}}$, a* convolution slicer *$\mathcal{S}(\cdot|K^{(1)},\ldots,K^{(N)})$ on $\mathbb{R}^{c\times d\times d}$ is a composition of N convolution functions with kernels $K^{(1)},\ldots,K^{(N)}$ (with stride or dilation if needed) such that $\mathcal{S}(X|K^{(1)},\ldots,K^{(N)}) \in \mathbb{R} \quad \forall X \in \mathbb{R}^{c\times d\times d}$.*

As indicated in Definition 2, the idea of the convolution slicer is to progressively map a given data X to a one-dimensional subspace through a sequence of convolution kernels, which capture spatial relations across channels as well as local information of the data. It is starkly different from the vectorization step in standard sliced Wasserstein on images (1). The illustration of the convolution slicer is given in Figure 4 in Appendix C.

We consider three particular types of convolution slicers based on using linear function on the convolution operator, named convolution-base, convolution-stride, and convolution-dilation slicers. We defer the definition of convolution-dilation slicers to Definition 6 in Appendix B. We first start with the definition of the convolution-base slicer.

Definition 3 *(Convolution-base Slicer) Given $X \in \mathbb{R}^{c\times d\times d}$ $(d \geq 2)$,*

1. When d is even, N is the biggest integer that satisfies $d = 2^{N-1}a$ with a is also an integer, sliced kernels are defined as $K^{(1)} \in \mathbb{R}^{c\times(2^{-1}d+1)\times(2^{-1}d+1)}$ and $K^{(h)} \in \mathbb{R}^{1\times(2^{-h}d+1)\times(2^{-h}d+1)}$ for $h = 2,\ldots,N-1$, and $K^{(N)} \in \mathbb{R}^{1\times a\times a}$ where $a = \frac{d}{2^{N-1}}$. Then, the convolution-base slicer *$\mathcal{CS}$-$b(X|K^{(1)},\ldots,K^{(N)})$ is defined as:*

$$\mathcal{CS}\text{-}b(X|K^{(1)},\ldots,K^{(N)}) = X^{(N)}, \quad X^{(h)} = \begin{cases} X & h = 0 \\ X^{(h-1)} \overset{1,1}{*} K^{(h)} & 1 \leq h \leq N, \end{cases} \tag{2}$$

2. When d is odd, the convolution-base slicer $\mathcal{CS}$-$b(X|K^{(1)},\ldots,K^{(N)})$ *takes the form:*

$$\mathcal{CS}\text{-}b(X|K^{(1)},\ldots,K^{(N)}) = \mathcal{CS}\text{-}b(X \overset{1,1}{*} K^{(1)}|K^{(2)},\ldots,K^{(N)}), \tag{3}$$

where $K^{(1)} \in \mathbb{R}^{c\times 2\times 2}$ and $K^{(2)},\ldots,K^{(N)}$ are the corresponding sliced kernels that are defined on the dimension $d-1$.

The idea of the convolution-base slicer in Definition 3 is to reduce the width and the height of the image by half after each convolution operator. If the width and the height of the image are odd, the first convolution operator is to reduce the size of the image by one via convolution with kernels of size 2×2, and then the same procedure as that of the even case is applied. We would like to remark that the conventional slicing of sliced Wasserstein in Section 2 is equivalent to a convolution-base slicer $\mathcal{S}(\cdot|K^{(1)})$ where $K^{(1)} \in \mathbb{R}^{c\times d\times d}$ that satisfies the constraint $\sum_{h=1}^{c}\sum_{i=1}^{d}\sum_{j=1}^{d} K_{h,i,j}^{(1)2} = 1$.

We now discuss the second variant of the convolution slicer, named convolution-stride slicer, where we further incorporate stride into the convolution operators. Its definition is as follows.

Definition 4 *(Convolution-stride Slicer) Given $X \in \mathbb{R}^{c\times d\times d}$ $(d \geq 2)$,*

1. When d is even, N is the biggest integer that satisfies $d = 2^{N-1}a$ with a is also an integer, sliced kernels are defined as $K^{(1)} \in \mathbb{R}^{c\times 2\times 2}$ and $K^{(h)} \in \mathbb{R}^{1\times 2\times 2}$ for $h = 2,\ldots,N-1$, and $K^{(N)} \in \mathbb{R}^{1\times a\times a}$ where $a = \frac{d}{2^{N-1}}$. Then, the convolution-stride slicer $\mathcal{CS}$-$s(X|K^{(1)},\ldots,K^{(N)})$ *is defined as:*

$$\mathcal{CS}\text{-}s(X|K^{(1)},\ldots,K^{(N)}) = X^{(N)}, \quad X^{(h)} = \begin{cases} X & h = 0 \\ X^{(h-1)} \overset{2,1}{*} K^{(h)} & 1 \leq h \leq N-1, \\ X^{(h-1)} \overset{1,1}{*} K^{(h)} & h = N, \end{cases} \tag{4}$$

2. When d is odd, the convolution-stride slicer $\mathcal{CS}$-$s(X|K^{(1)},\ldots,K^{(N)})$ *takes the form:*

$$\mathcal{CS}\text{-}s(X|K^{(1)},\ldots,K^{(N)}) = \mathcal{CS}\text{-}s(X \overset{1,1}{*} K^{(1)}|K^{(2)},\ldots,K^{(N)}), \tag{5}$$

where $K^{(1)} \in \mathbb{R}^{c\times 2\times 2}$ and $K^{(2)},\ldots,K^{(N)}$ are the corresponding sliced kernels that are defined on the dimension $d-1$.

Similar to the convolution-base slicer in Definition 3, the convolution-stride slicer reduces the width and the height of the image by half after each convolution operator. We use the same procedure of reducing the height and the width of the image by one when the height and the width of the image are odd. The benefit of the convolution-stride slicer is that the size of its kernels does not depend on the width and the height of images as that of the convolution-base slicer. This difference improves the computational complexity and time complexity of the convolution-stride slicer over those of the convolution-base slicer (cf. Proposition 3 in Appendix B).

Non-linear convolution-base slicer: The composition of convolution functions in the linear convolution slicer and its linear variants is still a linear function, which may not be effective when the data lie in a complex and highly non-linear low-dimensional subspace. A natural generalization of linear convolution slicers to enhance the ability of the slicers to capture the non-linearity of the data is to apply a non-linear activation function after convolution operators. This enables us to define a non-linear slicer in Definition 7 in Appendix E. The non-linear slicer can be seen as a defining function in generalized Radon Transform [48] which was used in generalized sliced Wasserstein [24].

3.2 Convolution Sliced Wasserstein

Given the definition of convolution slicers, we now state general definition of convolution sliced Wasserstein. An illustration of the convolution sliced Wasserstein is given in Figure 4 in Appendix C.

Definition 5 *For any $p \geq 1$, the* convolution sliced Wasserstein *(CSW) of order $p > 0$ between two given probability measures $\mu,\nu \in \mathcal{P}_p(\mathbb{R}^{c\times d\times d})$ is given by:*

$$CSW_p(\mu,\nu) := \left(\mathbb{E}\left[W_p^p\left(\mathcal{S}(\cdot|K^{(1)},\ldots,K^{(N)})\sharp\mu, \mathcal{S}(\cdot|K^{(1)},\ldots,K^{(N)})\sharp\nu\right)\right]\right)^{\frac{1}{p}},$$

where the expectation is taken with respect to $K^{(1)} \sim \mathcal{U}(\mathcal{K}^{(1)}), \ldots, K^{(N)} \sim \mathcal{U}(\mathcal{K}^{(N)})$. *Here,* $\mathcal{S}(\cdot|K^{(1)}, \ldots, K^{(N)})$ *is a convolution slicer with* $K^{(l)} \in \mathbb{R}^{c^{(l)} \times k^{(l)} \times k^{(l)}}$ *for any* $l \in [N]$ *and* $\mathcal{U}(\mathcal{K}^{(l)})$ *is the uniform distribution with the realizations being in the set* $\mathcal{K}^{(l)}$ *which is defined as* $\mathcal{K}^{(l)} := \left\{K^{(l)} \in \mathbb{R}^{c^{(l)} \times k^{(l)} \times k^{(l)}} | \sum_{h=1}^{c^{(l)}} \sum_{i'=1}^{k^{(l)}} \sum_{j'=1}^{k^{(l)}} K_{h,i',j'}^{(i)2} = 1\right\}$, *namely, the set* $\mathcal{K}^{(l)}$ *consists of tensors* $K^{(l)}$ *whose squared* ℓ_2 *norm is 1.*

The constraint that ℓ_2 norms of $K^{(l)}$ is 1 is for guaranteeing the distances between projected supports are bounded. When we specifically consider the convolution slicer as convolution-base slicer ($\mathcal{CS}$-b), convolution-stride slicer ($\mathcal{CS}$-s), and convolution-dilation slicer ($\mathcal{CS}$-d), we have the corresponding notions of convolution-base sliced Wasserstein (CSW-b), convolution-stride sliced Wasserstein (CSW-s), and convolution-dilation sliced Wasserstein (CSW-d).

Monte Carlo estimation and implementation: Similar to the conventional sliced Wasserstein, the expectation with respect to kernels $K^{(1)}, \ldots, K^{(N)}$ uniformly drawn from the sets $\mathcal{K}^{(1)}, \ldots, \mathcal{K}^{(N)}$ in the convolution sliced Wasserstein is intractable to compute. Therefore, we also make use of Monte Carlo method to approximate the expectation, which leads to the following approximation of the convolution sliced Wasserstein:

$$\text{CSW}_p^p(\mu, \nu) \approx \frac{1}{L} \sum_{i=1}^{L} W_p^p \left(\mathcal{S}(\cdot|K_i^{(1)}, \ldots, K_i^{(N)})\sharp\mu, \mathcal{S}(\cdot|K_i^{(1)}, \ldots, K_i^{(N)})\sharp\nu\right), \tag{6}$$

where $K_i^{(\ell)}$ are uniform samples from the sets $\mathcal{K}^{(\ell)}$ (which is equivalent to sample uniformly from $\mathbb{S}^{c^{(l)} \cdot k^{(l)2}}$ then applying the one-to-one reshape mapping) for any $\ell \in [N]$ and $i \in [L]$. Since each of the convolution slicer $\mathcal{S}(\cdot|K_i^{(1)}, \ldots, K_i^{(N)})$ is in one dimension, we can utilize the closed-form expression of Wasserstein metric in one dimension to compute $W_p \left(\mathcal{S}(\cdot|K_i^{(1)}, \ldots, K_i^{(N)})\sharp\mu, \mathcal{S}(\cdot|K_i^{(1)}, \ldots, K_i^{(N)})\sharp\nu\right)$ with a complexity of $\mathcal{O}(m \log_2 m)$ for each $i \in [L]$ where m is the maximum number of supports of μ and ν. Therefore, the total computational complexity of computing the Monte Carlo approximation (6) is $\mathcal{O}(Lm \log_2 m)$ when the probability measures μ and ν have at most m supports. It is comparable to the computational complexity of sliced Wasserstein on images (1) where we directly vectorize the images and apply the Radon transform to these flatten images. Finally, for the implementation, we would like to remark that L convolution slicers in equation (6) can be computed *independently* and *parallelly* using the group convolution implementation which is supported in almost all libraries.

Properties of convolution sliced Wasserstein: We first have the following result for the metricity of the convolution sliced Wasserstein.

Theorem 1 *For any* $p \geq 1$, *the convolution sliced Wasserstein* $CSW_p(.,.)$ *is a pseudo-metric on the space of probability measures on* $\mathbb{R}^{c \times d \times d}$, *namely, it is symmetric, satisfies the triangle inequality, and* $CSW_p(\mu, \nu) = 0 \iff \mu = \nu$.

Proof of Theorem 1 is in Appendix D.1. Our next result establishes the connection between the convolution sliced Wasserstein and max-sliced Wasserstein and Wasserstein distances.

Proposition 1 *For any* $p \geq 1$, *we find that* $CSW_p(\mu, \nu) \leq$ *Max-*$SW_p(\mu, \nu) \leq W_p(\mu, \nu)$, *where* *Max-*$SW_p(\mu, \nu) := \max_{\theta \in \mathbb{R}^{cd^2}: \|\theta\| \leq 1} W_p(\theta\sharp\mu, \theta\sharp\nu)$ *is max-sliced Wasserstein of order* p.

Proof of Proposition 1 is in Appendix D.2. Given the bounds in Proposition 1, we demonstrate that the convolution 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 2 *Assume that* P *is a probability measure supported on compact set of* $\mathbb{R}^{c \times d \times d}$. *Let* $X_1, X_2, \ldots, X_n$ *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*

$$\mathbb{E}[CSW_p(P_n, P)] \leq C\sqrt{(cd^2 + 1) \log n / n},$$

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

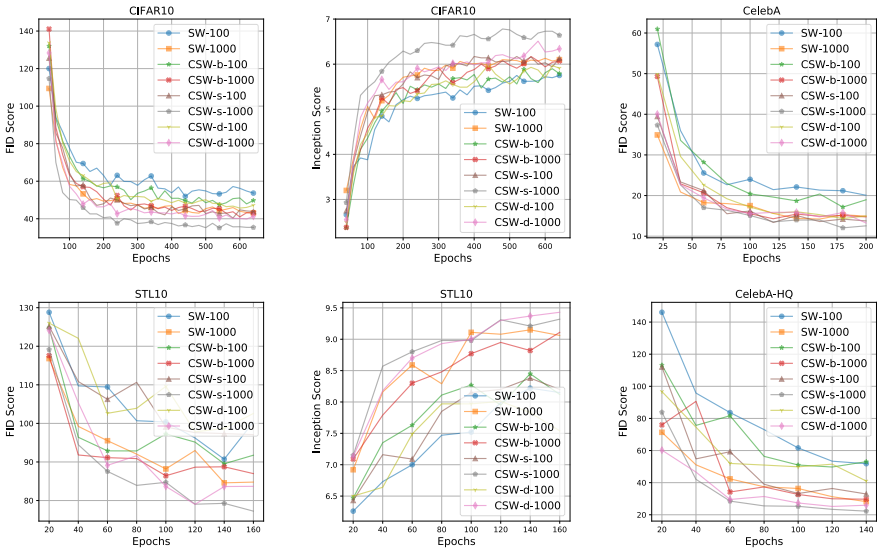

Figure 1: FID scores and IS scores over epochs of different training losses on datasets. We observe that CSW's variants usually help the generative models converge faster.

Proof of Proposition 2 is in Appendix D.3. The result of Proposition 2 indicates that the sample complexity of the convolution sliced Wasserstein is comparable to that of the sliced Wasserstein on images (1), which is at the order of $\mathcal{O}(n^{-1/2})$ [4], and better than that of the Wasserstein metric, which is at the order of $\mathcal{O}(n^{-1/(2cd^2)})$ [15].

Extension to non-linear convolution sliced Wasserstein: In Appendix E, we provide a non-linear version of the convolution sliced Wasserstein, named non-linear convolution sliced Wasserstein. The high-level idea of the non-linear version is to incorporate non-linear activation functions to the convolution-base, convolution-stride, and convolution-dilation slicers. The inclusion of non-linear activation functions is to enhance the ability of slicers to capture the non-linearity of the data. By plugging these non-linear convolution slicers into the general definition of the convolution sliced Wasserstein in Definition 5, we obtain the non-linear variants of convolution sliced Wasserstein.

4 Experiments

In this section, we focus on comparing the sliced Wasserstein (SW) (with the conventional slicing), the convolution-base sliced Wasserstein (CSW-b), the convolution sliced Wasserstein with stride (CSW-s), and the convolution sliced Wassersstein with dilation (CSW-d) (see Definition 6 in Appendix B)) in training generative models on standard benchmark image datasets such as CIFAR10 (32x32) [27], STL10 (96x96) [8], CelebA (64x64), and CelebA-HQ (128x128) [37]. We recall that the number of projections in SW and CSW's variants is denoted as L. Finally, we also show the values of the SW and the CSW variants between probability measures over digits of the MNIST dataset [30] in Appendix F.1. From experiments on MNIST, we observe that values of CSW variants are similar to values of SW while having better projection complexities.

In generative modeling, we follow the framework of the sliced Wasserstein generator in [13] with some modifications of neural network architectures. The details of the training are given in Appendix F.2. We train the above model on standard benchmarks such as CIFAR10 (32x32) [27], STL10 (96x96) [8], CelebA (64x64), and CelebAHQ (128x128) [37]. To compare models, we use the FID score [20] and the Inception score (IS) [51]. The detailed settings about architectures, hyperparameters, and evaluation of FID and IS are given in Appendix G. We first show the FID scores and IS scores of generative models trained by SW and CSW's variants with the number of projections $L \in \{1, 100, 1000\}$ in Table 1. In the table, we report the performance of models at the last training epoch. We do not report the IS scores on CelebA and CelebA-HQ since the IS scores are not suitable for face images. We then demonstrate the FID scores and IS scores across training epochs in Figure 1 for investigating the convergence of generative models trained by SW and CSW's variants. After

Table 1: Summary of FID and IS scores of methods on CIFAR10 (32x32), CelebA (64x64), STL10 (96x96), and CelebA-HQ (128x128). Some results on CIFAR10 are reported from 5 different runs.

Method	CIFAR10 (32x32)		CelebA (64x64)	STL10 (96x96)		CelebA-HQ (128x128)
	FID (↓)	IS (↑)	FID (↓)	FID (↓)	IS (↑)	FID (↓)
SW (L=1)	87.97	3.59	128.81	170.96	3.68	**275.44**
CSW-b (L=1)	84.38	4.28	85.83	173.33	**3.89**	315.91
CSW-s (L=1)	80.10	4.31	**66.52**	**168.9**3	3.75	303.57
CSW-d (L=1)	**63.94**	**4.89**	89.37	212.61	2.48	321.06
SW (L=100)	52.36±0.76	5.79±0.16	20.08	100.35	8.14	51.80
CSW-b (L=100)	49.67±2.00	5.87±0.15	18.96	**91.75**	8.11	53.05
CSW-s (L=100)	**43.73±2.09**	**6.17±0.06**	**13.76**	97.08	**8.20**	**32.94**
CSW-d (L=100)	47.23±1.12	5.97±0.11	14.96	102.58	7.53	41.01
SW (L=1000)	44.25±1.21	6.02±0.03	14.92	84.78	9.06	28.19
CSW-b (L=1000)	42.88±0.98	6.11±0.10	14.75	86.98	9.11	29.69
CSW-s (L=1000)	**36.80±1.44**	**6.55±0.12**	**12.55**	**77.24**	9.31	**22.25**
CSW-d (L=1000)	40.44±1.02	6.38±0.14	13.24	83.36	**9.42**	25.93

Table 2: Computational time and memory of methods (reported in the number of iterations per a second and megabytes (MB).

Method	CIFAR10 (32x32)		CelebA (64x64)		STL10 (96x96)		CelebA-HQ (128x128)	
	Iters/s (↑)	Mem (↓)	Iters/s (↑)	Mem (↓)	Iters/s (↑)	Mem (↓)	Iters/s (↑)	Mem (↓)
SW (L=1)	18.98	2071	6.21	8003	9.59	4596	10.35	4109
SW (L=100)	18.53	2080	6.16	8015	9.47	4601	10.22	4117
SW (L=1000)	18.15	2169	6.10	8102	9.13	4647	10.17	4202
CSW-b (L=1)	18.43	2070	6.21	8003	9.56	4596	10.33	4109
CSW-b (L=100)	18.35	2077	6.15	8009	9.40	4598	10.19	4110
CSW-b (L=1000)	18.06	2117	6.10	8049	9.07	4613	10.12	4134
CSW-s (d) (L=1)	18.69	2070	6.21	8003	9.56	4596	10.33	4109
CSW-s (d) (L=100)	18.50	2073	6.16	8005	9.41	4597	10.20	4109
CSW-s (d) (L=1000)	18.10	2098	6.10	8029	9.10	4603	10.12	4114

that, we report the training time and training memory of SW and CSW variants in Table 2. Finally, we show randomly generated images from SW's models and CSW-s' models on CelebA dataset in Figure 2. Generated images of all models on all datasets are given in Figures 5-8 in Appendix F.2.

Summary of FID scores and IS scores: According to Table 1, on CIFAR10, CSW-d gives the lowest values of FID scores and IS scores when $L = 1$ while CSW-s gives the lowest FID scores when $L = 100$ and $L = 1000$. Compared to CSW-s, CSW-d and CSW-b yield higher FID scores and lower IS scores. However, CSW-d and CSW-b are still better than SW. On CelebA, CSW-s performs the best in all settings. On STL10, CSW's variants are also better than the vanilla SW; however, it is unclear which is the best variant. On CelebA-HQ, SW gives the lowest FID score when $L = 1$. In contrast, when $L = 100$ and $L = 1000$, CSW-s is the best choice for training the generative model. Since the FID scores of $L = 1$ are very high on CelebA-HQ and STL10, the scores are not very meaningful for comparing SW and CSW's variants. For all models, increasing L leads to better quality. Overall, we observe that CSW's variants enhance the performance of generative models.

FID scores and IS scores across epochs: From Figure 1, we observe that CSW's variants help the generative models converge faster than SW when $L = 100$ and $L = 1000$. Increasing the number of projections from 100 to 1000, the generative models from both SW and CSW's variants become better. Overall, CSW-s is the best option for training generative models among CSW's variants since its FID curves are the lowest and its IS curves are the highest.

Training time and training memory: We report in Table 2 the training speed in the number of iterations per second and the training memory in megabytes (MBs). We would like to recall that the time complexity and the projection memory complexity of CSW-s and CSW-d are the same. Therefore, we measure the training time and the training memory of CSW-s as the result for both CSW-s and CSW-d. We can see that increasing the number of projections L costs more memory and also slows down the training speed. However, the rate of increasing memory of CSW is smaller than SW. For CSW-s and CSW-d, the extent of saving memory is even better. As an example, $L = 1000$ in CSW-s and CSW-d costs less memory than SW with $L = 100$ while the performance is better (see Table 1). In terms of training time, CSW-s and CSW-d are comparable to SW and they can

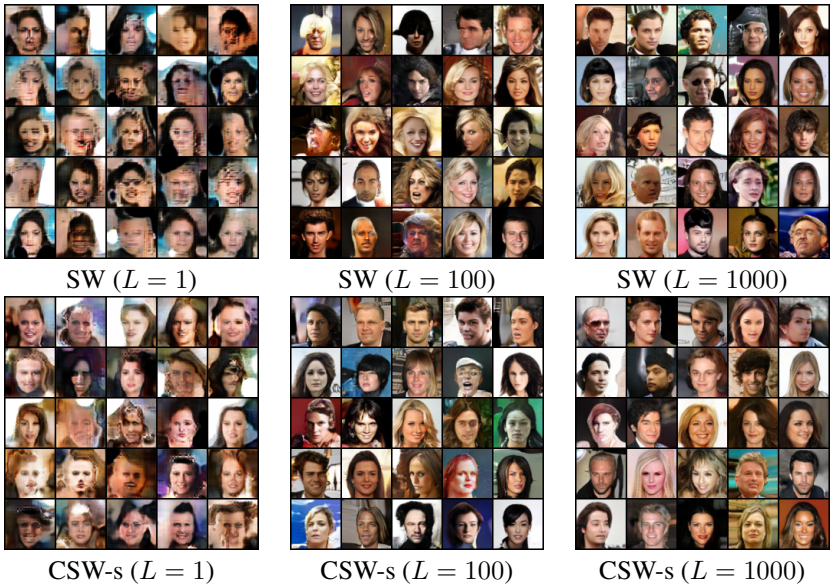

Figure 2: Random generated images of SW and CSW-s on CelebA.

be computed faster than CSW. We refer the readers to Section 3 for a detailed discussion about the computational time and projection memory complexity of CSW's variants.

Generated images: We show randomly generated images on CelebA dataset in Figure 2 and Figure 6 (Appendix F), and generated images on CIFAR10, CelebA, STL10, and CelebA-HQ in Figures 5-8 as qualitative comparison between SW and CSW variants. From the figures, we can see that generated images of CSW-s is more realistic than ones of SW. The difference is visually clear when the number of projections L is small e.g., $L = 1$ and $L = 100$. When $L = 1000$, we can still figure out that CSW-s is better than SW by looking at the sharpness of the generated images. Also, we can visually observe the improvement of SW and CSW-s when increasing the number of projections. In summary, the qualitative results are consistent with the quantitative results (FID scores and IS scores) in Table 1. For the generated images of CSW-b and CSW-d, we also observe the improvement compared to the SW which is consistent with the improvement of FID scores and IS scores.

Non-linear convolution sliced Wasserstein: We also compare non-linear extensions of SW and CSW variants in training generative models on CIFAR10 in Appendix F. For details of non-linear extensions, we refer to Appendix E. From experiments, we observe that convolution can also improve the performance of sliced Wasserstein in non-linear projecting cases. Compared to linear versions, non-linear versions can enhance the quality of the generative model or yield comparable results.

5 Conclusion

We have addressed the issue of the conventional slicing process of sliced Wasserstein when working with probability measures over images. In particular, sliced Wasserstein is defined on probability measures over vectors which leads to the step of vectorization for images. As a result, the conventional slicing process cannot exploit the spatial structure of data for designing the space of projecting directions and projecting operators. To address the issue, we propose a new slicing process by using the convolution operator which has been shown to be efficient on images. Moreover, we investigate the computational complexity and projection memory complexity of the new slicing technique. We show that convolution slicing is comparable to conventional slicing in terms of computational complexity while being better in terms of projection memory complexity. By utilizing the new slicing technique, we derive a novel family of sliced Wassersein variants, named convolution sliced Wasserstein. We investigate the properties of the convolution sliced Wasserstein including its metricity, its computational and sample complexities, and its connection to other variants of sliced Wasserstein in literature. Finally, we carry out extensive experiments in comparing digits images and training generative models on standard benchmark datasets to demonstrate the favorable performance of the convolution sliced Wasserstein.

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.

[2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.

[3] 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.

[4] S. Bobkov and M. Ledoux. 'One-dimensional empirical measures, order statistics, and Kantorovich transport distances. *Memoirs of the American Mathematical Society*, 261, 2019.

[5] C. Bonet, N. Courty, F. Septier, and L. Drumetz. Sliced-Wasserstein gradient flows. *arXiv preprint arXiv:2110.10972*, 2021.

[6] 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.

[7] X. Chen, Y. Yang, and Y. Li. Augmented sliced Wasserstein distances. *International Conference on Learning Representations*, 2022.

[8] A. Coates, A. Ng, and H. Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international Conference on Artificial Intelligence and Statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011.

[9] 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.

[10] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.

[11] B. Dai and U. Seljak. Sliced iterative normalizing flows. In *International Conference on Machine Learning*, pages 2352–2364. PMLR, 2021.

[12] 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.

[13] 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.

[14] K. Fatras, Y. Zine, R. Flamary, R. Gribonval, and N. Courty. Learning with minibatch Wasserstein: asymptotic and gradient properties. In *AISTATS 2020-23nd International Conference on Artificial Intelligence and Statistics*, volume 108, pages 1–20, 2020.

[15] 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.

[16] K. Fukushima and S. Miyake. Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. In *Competition and cooperation in neural nets*, pages 267–285. Springer, 1982.

[17] Z. Goldfeld and K. Greenewald. Sliced mutual information: A scalable measure of statistical dependence. *Advances in Neural Information Processing Systems*, 34, 2021.

[18] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.

[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.

[20] 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.

[21] 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.

[22] M. Huang, S. Ma, and L. Lai. A Riemannian block coordinate descent method for computing the projection robust Wasserstein distance. In *International Conference on Machine Learning*, pages 4446–4455. PMLR, 2021.

[23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[24] 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.

[25] S. Kolouri, P. E. Pope, C. E. Martin, and G. K. Rohde. Sliced Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018.

[26] 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.

[27] 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.

[28] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 2012.

[29] 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.

[30] 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.

[31] 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.

[32] 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.

[33] T. Lin, C. Fan, N. Ho, M. Cuturi, and M. Jordan. Projection robust Wasserstein distance and Riemannian optimization. *Advances in Neural Information Processing Systems*, 33:9383–9397, 2020.

[34] 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.

[35] 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.

[36] 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.

[37] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

[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.

[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.

[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.

[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.

[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.

[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.

[44] F.-P. Paty and M. Cuturi. Subspace robust Wasserstein distances. In *International Conference on Machine Learning*, pages 5072–5081, 2019.

[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] 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.

[47] G. Peyré and M. Cuturi. Computational optimal transport, 2020.

[48] J. Radon. 1.1 über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten. *Classic papers in modern diagnostic radiology*, 5:21, 2005.

[49] A. Rakotomamonjy and R. Liva. Differentially private sliced Wasserstein distance. In *International Conference on Machine Learning*, pages 8810–8820. PMLR, 2021.

[50] 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.

[51] 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.

[52] J. Solomon, F. De Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (ToG)*, 34(4):1–11, 2015.

[53] V. Titouan, R. Flamary, N. Courty, R. Tavenard, and L. Chapel. Sliced Gromov-Wasserstein. *Advances in Neural Information Processing Systems*, 32, 2019.

[54] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018.

[55] C. Villani. *Optimal transport: Old and New*. Springer, 2008.

[56] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019.

[57] 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.

[58] 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.

[59] 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.

[60] M. Yi and S. Liu. Sliced Wasserstein variational inference. In *Fourth Symposium on Advances in Approximate Bayesian Inference*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]