



HAL
open science

Gaussian-Smoothed Sliced Probability Divergences

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

► **To cite this version:**

Mokhtar Z. Alaya, Alain Rakotomamonjy, Maxime Berar, Gilles Gasso. Gaussian-Smoothed Sliced Probability Divergences. 2023. hal-04282819v1

HAL Id: hal-04282819

<https://hal.science/hal-04282819v1>

Preprint submitted on 13 Nov 2023 (v1), last revised 22 Apr 2024 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Gaussian-Smoothed Sliced Probability Divergences

Mokhtar Z. Alaya¹ Alain Rakotomamonjy² Maxime Bérar³ Gilles Gasso⁴

November 13, 2023

Abstract

Gaussian smoothed sliced Wasserstein distance has been recently introduced for comparing probability distributions, while preserving privacy on the data. It has been shown that it provides performances similar to its non-smoothed (non-private) counterpart. However, the computational and statistical properties of such a metric have not yet been well-established. This work investigates the theoretical properties of this distance as well as those of generalized versions denoted as Gaussian-smoothed sliced divergences $G_\sigma SD^p$. We first show that smoothing and slicing preserve the metric property and the weak topology. To study the sample complexity of such divergences, we then introduce $\hat{\mu}_n$ the double empirical distribution for the smoothed-projected μ . The distribution $\hat{\mu}_n$ is a result of a double sampling process: one from sampling according to the origin distribution μ and the second according to the convolution of the projection of μ on the unit sphere and the Gaussian smoothing. We particularly focus on the Gaussian smoothed sliced Wasserstein distance and prove that it suffers from an unavoidable bias of approximation of order $\log(n/2)$. We also derive other properties, including continuity, of different divergences with respect to the smoothing parameter. We support our theoretical findings with empirical studies in the context of privacy-preserving domain adaptation.

1 Introduction

Divergences for comparing two distributions have been shown to be important for achieving good performance in the contexts of generative modeling (Arjovsky et al., 2017; Salimans et al., 2018), domain adaptation (Long et al., 2015; Courty et al., 2016; Lee et al., 2019), and in computer vision (Bonneel et al., 2011; Solomon et al., 2015) among many more applications (Kolouri et al., 2017; Peyré and Cuturi, 2019). Examples of divergences that have proved useful for these tasks are the Maximum Mean Discrepancy (Gretton et al., 2012; Long et al., 2015; Sutherland et al., 2017), the Wasserstein distance (Monge, 1781; Kantorovich, 1942; Villani, 2009) or its variant the sliced Wasserstein distance (SW) (Bonneel and Coeurjolly, 2019; Kolouri et al., 2019a; Nguyen et al., 2020; Kolouri et al., 2016).

The SW distance has the advantage of being computationally efficient, since it uses a closed-form solution for distributions with support on \mathbb{R} , by computing the expectation of one-dimensional (1D) random projections of distributions in \mathbb{R}^d . Owing to this efficiency and the resulting scalability, this distance has been successfully applied in several applications ranging from generative models to

¹LMAC EA 2222, Université de Technologie de Compiègne, *e-mail*: alayaelm@utc.fr

²Criteo AI Lab & LITIS EA4108, Université de Rouen Normandie, *e-mail*: alain.rakoto@insa-rouen.fr

³LITIS EA4108, Université de Rouen Normandie, *e-mail*: maxime.berar@univ-rouen.fr

⁴LITIS EA4108, INSA Rouen, *e-mail*: gilles.gasso@insa-rouen.fr

domain adaptation (Kolouri et al., 2019b; Deshpande et al., 2019; Wu et al., 2019; Lee et al., 2019) and its statistical properties have been well-studied in Nadjahi et al. (2020).

Recently, Gaussian smoothed variants of the Wasserstein distance and the sliced Wasserstein distance have been introduced respectively in Nietert et al. (2021) and in Rakotomamonjy and Ralaivola (2021). One main motivation behind these variants is to provide a privacy guarantee for the distribution comparison task as Gaussian smoothing is known to be a mechanism for achieving differential privacy (Dwork et al., 2014). While the properties of the Gaussian smoothed Wasserstein distance have been extensively studied by Nietert et al. (2021), the properties of the Gaussian smoothed sliced Wasserstein distance have not been fully investigated yet although they are known to be more computationally efficient. In this work, we fill this gap by providing a theoretical analysis of the Gaussian smoothed sliced Wasserstein distance.

We investigate the theoretical properties of the Gaussian-smoothed sliced Wasserstein distance and the ones of more general Gaussian-smoothed sliced divergences induced by some base distances or divergences for distributions defined in \mathbb{R}^d . Specifically, as for main contributions, we establish the topological properties of these divergences.

Then, we focus on the sample complexity of such divergences by introducing the double empirical distribution for the smoothed-projected origin distribution μ . The new empirical distribution is a result of double sampling process: one from sampling according to the origin distribution and the second according to the convolution of the projection of μ on the unit sphere and the Gaussian smoothing. We particularly focus on the Gaussian smoothed sliced Wasserstein distance since the Wasserstein base distance enjoys an order statistic representation Bobkov and Ledoux (2019). Under some mild assumptions we also prove that the Gaussian-smoothed sliced divergences satisfy an order relation with respect to the noise level and are also continuous with respect to this parameter. Given the importance of the noise level in the privacy/utility trade-off achieved by the divergence, this latter property is of high impact as it supports a computationally cheap warm-start/fine-tuning procedure when looking for a privacy/utility compromise of the divergence. Our theoretical study is backed by some numerical experiments on toy problems and on domain adaptation illustrating how owing to the topology induced by our metric and its continuity, differential privacy comes almost for free (without loss of performance) and multiple models with different level of privacy can be cheaply computed.

The paper is organized as follows: after introducing the notation and some background in Section 2, we detail the topological properties of Gaussian-smoothed sliced divergence in Section 3.1 while the statistical properties are established in Section 3.2. The noise analyses are provided in Section 3.3. Experimental analyses for supporting the theory and showcasing the relevance of our divergences in domain adaptation are depicted in Section 4. Discussions on the perspectives and limitations are in Section 5. All the proofs of the theoretical results and some additional experiments are postponed to the appendices in the supplementary.

2 Preliminaries

For the reader’s convenience, we provide a brief summary of standard notations and definitions used throughout the paper.

Notation. For $d \in \mathbb{N}^*$, let $\mathcal{P}(\mathbb{R}^d)$ be the set of Borel probability measures on \mathbb{R}^d and $\mathcal{P}_p(\mathbb{R}^d) \subset \mathcal{P}(\mathbb{R}^d)$, those with finite moment of order p , i.e., $\mathcal{P}_p(\mathbb{R}^d) \triangleq \{\mu \in \mathcal{P} : \int \|x\|^p d\mu(x) < \infty\}$, where

$\|\cdot\|$ is the Euclidean norm. We denote $M_p(\mu) = \int_x \|x\|^p d\mu(x)$. For two probability distributions μ and ν , we denote their convolution as $\mu * \nu \in \mathcal{P}(\mathbb{R}^d)$, namely $(\mu * \nu)(A) = \int_x \int_y \mathbf{1}_A(x+y) d\mu(x) d\nu(y)$, where $\mathbf{1}_A(\cdot)$ is the indicator function over A . Given two independent random variables $X \sim \mu$ and $Y \sim \nu$, we remind that $X + Y \sim \mu * \nu$. The d -dimensional unit-sphere is noted as $\mathbb{S}^{d-1} \triangleq \{\theta \in \mathbb{R}^d : \|\theta\| = 1\}$. We denote by u_d the uniform distribution on \mathbb{S}^{d-1} and we use $\delta(\cdot)$ to denote the Kronecker delta function. We note as $\mathbf{E}_\mu f$ the expectation of the function f with respect to μ .

Sliced Wasserstein distance. We remind in this paragraph several measures of similarity between two distributions. The Wasserstein distance of order $p \in [1, \infty)$ between two measures in $\mathcal{P}_p(\mathbb{R}^d)$ is given by the relaxation of the optimal transport problem, and it is defined as

$$W_p^p(\mu, \nu) = \inf_{\gamma \in \Pi(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - x'\|^p \gamma(x, x') dx dx'$$

where $\Pi(\mu, \nu) \triangleq \{\gamma \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d) | \pi_{1\#}\gamma = \mu, \pi_{2\#}\gamma = \nu\}$ and π_1, π_2 are the marginal projectors of γ on each of its coordinates. When $d = 1$, the Wasserstein distance can be calculated in closed-form owing to the cumulative distributions of μ and ν (Rachev and Rüschendorf, 1998). In practice for empirical distributions, the closed-form solution requires only the sorting of the samples, which makes it very efficient. Because of this efficiency, efforts have been devoted to derive a metric for high-dimensional distributions based on 1D Wasserstein distance. The main idea is to project high-dimensional probability distributions onto a random one-dimensional space and then to compute the Wasserstein distance. This operation can be theoretically formalized through the use of the Radon transform, leading to the so-called sliced Wasserstein distance (Bonneel and Coeurjolly, 2019; Kolouri et al., 2019a; Nguyen et al., 2020; Kolouri et al., 2016).

Definition 1. For any $p \in [1, \infty)$ and two measures $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, the sliced Wasserstein distance (SW) reads as

$$SW^p(\mu, \nu) \triangleq \int_{\mathbb{S}^{d-1}} W_p^p(\mathcal{R}_{\mathbf{u}}\mu, \mathcal{R}_{\mathbf{u}}\nu) u_d(\mathbf{u}) d\mathbf{u}.$$

where $\mathcal{R}_{\mathbf{u}}$ is the Radon transform of a probability distribution, namely $\mathcal{R}_{\mathbf{u}}\mu(\cdot) = \int_{\mathbb{R}^d} \mu(\mathbf{s}) \delta(\cdot - \mathbf{s}^\top \mathbf{u}) d\mathbf{s}$. In practice, the integral is approximated through a Monte-Carlo simulation leading to a sum of 1D Wasserstein distances over a fixed number of random directions \mathbf{u} .

Gaussian-smoothed sliced Wasserstein distance. Based on this definition of SW, replacing the Radon projected measures with their Gaussian-smoothed counterpart leads to the following definition:

Definition 2. The σ -Gaussian-smoothed p -Sliced Wasserstein distance between probability distributions μ and ν in $\mathcal{P}_p(\mathbb{R}^d)$ writes as

$$G_\sigma SW^p(\mu, \nu) \triangleq \int_{\mathbb{S}^{d-1}} W_p^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u},$$

where $\mathcal{N}_\sigma = \mathcal{N}(0, \sigma^2)$ is the zero-mean σ^2 -variance Gaussian measure.

It is important to note here that the smoothing (convolution) operation occurs after projection onto the one-dimensional space. Hence, assuming $X \sim \mu, Y \sim \nu$, for a given direction \mathbf{u} , we compute in

the integral the one-dimensional Wasserstein distance between the probability laws of $\mathbf{u}^\top X + Z$ and $\mathbf{u}^\top Y + Z'$ where $Z, Z' \sim \mathcal{N}_\sigma$ are independent random variables. The metric properties of $G_\sigma SW_p$ for $p \geq 1$ have been discussed in a recent work (Rakotomamonjy and Ralaivola, 2021). This latter work has also shown, in the context of differential privacy, the importance of convolving the Radon projected distribution with a Gaussian instead of computing the SW distance of the original distribution smoothed with a d -dimensional Gaussian $\mu * \mathcal{N}_{\sigma \mathbf{I}_d}$, where \mathbf{I}_d denotes the $d \times d$ identity matrix.

Gaussian-smoothed sliced divergence. The idea of slicing high-dimensional distributions before feeding them to a divergence between probability distributions can be extended to distances other than the Wasserstein distance. These sliced divergences have been studied by Nadjahi et al. (Nadjahi et al., 2020). In a similar way, we can define a Gaussian-smoothed sliced divergence, given a divergence $D_{\mathbb{R}^d} : \mathcal{P}_p(\mathbb{R}^d) \times \mathcal{P}_p(\mathbb{R}^d) \rightarrow \mathbb{R}^+$ for $d \geq 1$ as:

Definition 3. *The σ -Gaussian-smoothed p -Sliced Divergence between probability distributions μ and ν in $\mathcal{P}_p(\mathbb{R}^d)$ associated to the base divergence $D \triangleq D_{\mathbb{R}^d}$, $p \geq 1$ is*

$$G_\sigma SD^p(\mu, \nu) \triangleq \int_{\mathbb{S}^{d-1}} D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u}.$$

where the superscript p refers to a power.

Typical relevant divergences are the maximum mean discrepancy (MMD) (Gretton et al., 2012) or the Sinkhorn divergence (Genevay et al., 2018; Peyré and Cuturi, 2019). In Section 4, we report empirical findings based on these divergences as well as on the Wasserstein distance.

3 Theoretical properties

In this section, we analyze the properties of the Gaussian-smoothed sliced divergence, in terms of topological and statistical properties and the influence of the Gaussian smoothing parameter σ on the distance.

3.1 Topology

It has already been shown in Rakotomamonjy and Ralaivola (2021) that the Gaussian-smoothed sliced Wasserstein is a metric on $\mathcal{P}(\mathbb{R}^d)$. In the next, we extend these results to any divergence $D(\cdot, \cdot)$ under certain assumptions.

Theorem 1. *For any $\sigma > 0, p \geq 1$, the following properties hold:*

1. *if $D(\cdot, \cdot)$ is non-negative (or symmetric), then $G_\sigma SD(\cdot, \cdot)$ is non-negative (or symmetric);*
2. *if $D(\cdot, \cdot)$ satisfies the identity of indiscernibles, i.e. for $\mu', \nu' \in \mathcal{P}(\mathbb{R})$, $D(\mu', \nu') = 0$ if and only if $\mu' = \nu'$, then this identity also holds for $G_\sigma SD(\cdot, \cdot)$ for any $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$;*
3. *if $D(\cdot, \cdot)$ satisfies the triangle inequality then $G_\sigma SD(\cdot, \cdot)$ satisfies the triangle inequality.*

The above theorem shows that under mild hypotheses over the base divergence D , as being a metric for instance, the metric property of its Gaussian-smoothed sliced version naturally derives. As exposed in the appendix, the more involved property to prove is the identity of indiscernibles. We further postponed to the appendix two other topological properties: (i) $G_\sigma\text{SD}$ metrizes the weak topology on $\mathcal{P}_p(\mathbb{R}^d)$ and (ii) $G_\sigma\text{SD}$ is lower semi-continuous with respect to the weak topology in $\mathcal{P}_p(\mathbb{R}^d)$.

3.2 Statistical properties

The next theoretical question we are interested in is about the incurred error when the true distribution μ is approximated by its empirical distribution $\hat{\mu}_n$. Such a case is common in practical applications where only (high-dimensional) empirical samples are at disposal. Specifically, we are interested in quantifying two key properties of empirical Gaussian-smoothed divergence: (i) the convergence of $G_\sigma\text{SD}(\hat{\mu}_n, \hat{\nu}_n)$ to $G_\sigma\text{SD}(\mu, \nu)$ (ii) the convergence of $\widehat{G_\sigma\text{SD}}(\mu, \nu)$ to $G_\sigma\text{SD}(\mu, \nu)$, i.e., when approximating the expectation over the random projection with sample mean.

Let $\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and $\hat{\nu}_n = \frac{1}{n} \sum_{i=1}^n \delta_{Y_i}$ be the empirical probability measures of independent observations. The smoothed Gaussian sliced divergence between $\hat{\mu}_n$ and $\hat{\nu}_n$ is given by

$$G_\sigma\text{SD}^p(\hat{\mu}_n, \hat{\nu}_n) = \int_{\mathbb{S}^{d-1}} D^p(\mathcal{R}_\mathbf{u}\hat{\mu}_n * \mathcal{N}_\sigma, \mathcal{R}_\mathbf{u}\hat{\nu}_n * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u}.$$

Remark 1. Remark that for a fixed $\mathbf{u} \in \mathbb{S}^{d-1}$, the distributions $\mathcal{R}_\mathbf{u}\hat{\mu}_n * \mathcal{N}_\sigma$ and $\mathcal{R}_\mathbf{u}\hat{\nu}_n * \mathcal{N}_\sigma$ are continuous, in particular they are a mixture of Gaussian distributions centered on the projected samples with variance σ^2 .

Lemma 1. Conditionally on the samples $\{X_i\}_{i=1, \dots, n}$ and $\{Y_j\}_{j=1, \dots, n}$, one has: $\mathcal{R}_\mathbf{u}\hat{\mu}_n * \mathcal{N}_\sigma = \frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)$ and $\mathcal{R}_\mathbf{u}\hat{\nu}_n * \mathcal{N}_\sigma = \frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)$.

Note that we further need to sample with respect to the continuous mixture Gaussian measures in Lemma 1 in order to get a *fully* empirical measure version of $G_\sigma\text{SD}(\mu, \nu)$. To this end, we define in the next section the *double* empirical divergence of $G_\sigma\text{SD}$.

3.2.1 Double empirical divergence of $G_\sigma\text{SD}$

Let T_1^x, \dots, T_n^x and T_1^y, \dots, T_n^y be i.i.d. observations of $\mathcal{R}_\mathbf{u}\hat{\mu}_n * \mathcal{N}_\sigma$ and $\mathcal{R}_\mathbf{u}\hat{\nu}_n * \mathcal{N}_\sigma$, respectively. Sampling i.i.d. $\{T_i^x\}_{i=1, \dots, n}$ is given by the following scheme: for $i = 1, \dots, n$, we first choose the component $\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)$ from the mixture $\frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)$ then we generate $T_i^x = \mathbf{u}^\top X_i + Z_i^x$, where $Z_i^x \sim \mathcal{N}_\sigma$. Hence, we define, for a given \mathbf{u}

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n \delta_{T_i^x} = \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}$$

and

$$\hat{\nu}_n = \frac{1}{n} \sum_{i=1}^n \delta_{T_i^y} = \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top Y_i + Z_i^y}.$$

The measure $\hat{\mu}_n \in \mathcal{P}(\mathbb{R})$ defines an empirical version of the continuous $\mathcal{R}_\mathbf{u}\hat{\mu}_n * \mathcal{N}_\sigma$ denoted as $\widehat{\mathcal{R}_\mathbf{u}\hat{\mu}_n * \mathcal{N}_\sigma}$ (similarly $\hat{\nu}_n = \widehat{\mathcal{R}_\mathbf{u}\hat{\nu}_n * \mathcal{N}_\sigma}$). Using the aforementioned notation, we define:

Definition 4. *The double empirical smoothed Gaussian sliced divergence reads as*

$$\hat{G}_\sigma \text{SD}^p(\hat{\mu}_n, \hat{\nu}_n) \triangleq \int_{\mathbb{S}^{d-1}} \text{D}^p(\hat{\mu}_n, \hat{\nu}_n) u_d(\mathbf{u}) d\mathbf{u}.$$

Remark 2. (i) *It is worth to comment the double randomnesses showing in the definition of $\hat{G}_\sigma \text{SD}^p(\hat{\mu}_n, \hat{\nu}_n)$: the first comes from sampling according to the original probability measure (μ or ν) whereas the second takes place from sampling according to the mixture $\frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)$.*
(ii) *The empirical measure of the convolution $\widehat{\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma}$ could be written as $\frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}$ allowing to sample in a one shot n i.i.d. samples $U_i^x + Q_i^x$ such that $U_i^x \sim \mathcal{R}_{\mathbf{u}}\mu$ and $Q_i^x \sim \mathcal{N}_\sigma$. From an empirical view, sampling according to $\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma$ is intractable. For that reason, our theoretical results and numerical experiments are based on $\hat{\mu}_n, \hat{\nu}_n$, and hence with respect to $\hat{G}_\sigma \text{SD}(\hat{\mu}_n, \hat{\nu}_n)$.*

3.2.2 Sample complexity of $G_\sigma \text{SW}^p$

Let $T_{(l)}^x$ the l -th order variable of T^x , namely $T_{(1)}^x = \min_{i=1, \dots, n} T_i^x$, $T_{(n)}^x = \max_{i=1, \dots, n} T_i^x$ and $T_{(l)}^x$ defines the l -th smallest of $\{T_i^x : i = 1, \dots, n\}$. The (ordered) collection of observed values $\{T_{(1)}^x \leq T_{(2)}^x \leq \dots \leq T_{(n)}^x\}$ is called the *order statistic* of T^x .

Herein, our goal is to quantify the error made when approximating $G_\sigma \text{SW}(\mu, \nu)$ with $\hat{G}_\sigma \text{SW}(\hat{\mu}_n, \hat{\nu}_n)$. More precisely, we are interested in establishing an order of the convergence rate of $\hat{G}_\sigma \text{SD}(\hat{\mu}_n, \hat{\nu}_n)$ towards $G_\sigma \text{SD}(\mu, \nu)$, according to the sample size n . This rate stands for the so-called *sample complexity*.

We denote hereafter by $\mu^{\otimes n}$ and $\mathcal{N}_\sigma^{\otimes n}$ the n -fold product extensions of μ and \mathcal{N}_σ , respectively. We denote by $\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}}[\cdot] = \mathbf{E}_{\mu^{\otimes n}}[\mathbf{E}_{\mathcal{N}_\sigma^{\otimes n}}[\cdot | X_1, \dots, X_n]]$. Next, we focus on the sample complexity for the special case of Gaussian-smoothed sliced Wasserstein distance.

Proposition 1. *Fix $\sigma > 0, p \geq 1$. Then,*

$$\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}}[\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \mu)] \leq \Xi_{p, \sigma; \mu} \left(4 \log\left(\frac{n}{2}\right) + \frac{\log n}{n} \right),$$

where $\Xi_{p, \sigma; \mu} = 2^{3p-1} (C_p + e^{5/3}) (M_{2p}(\mu) + 2^p \sigma^{4p})^{1/2}$ with C_p is a positive constant depending only on p .

Let us sketch the proof of Proposition 1 to show why we get a bias-like term of order $\log(\frac{n}{2})$. By application of triangle inequality for Wasserstein distance, we get $W_p^p(\hat{\mu}_n, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) \leq 2^{p-1} (\mathbf{I} + \mathbf{II})$ where

$$\mathbf{I} = W_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}, \frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x} \right)$$

and

$$\mathbf{II} = W_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma \right).$$

Remark 3. Due to the fact that in the **I** term we have to upper bound a p -Wasserstein distance between two empirical measures, our proof shows that the expectation of **I** is of order $O(\log(\frac{n}{2}))$. It is indeed an expectation of the absolute order statistics of the samples. To best of our knowledge, the (only) tight upper bound of absolute order statistics can be found in [Gribkova \(2020\)](#) (see [Theorem 13](#) in [Appendix C](#)). We further highlight that the true empirical version of $\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma$ is untractable and unfeasible to be implemented. We then approximate this latter using the double empirical version $\hat{\mu}_n$, which is prone to this unavoidable bias of approximation.

Remark 4. The term **II** is p -Wasserstein between $\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma$ and its true empirical measure, for that reason we get a standard order of $O(\frac{\log n}{n})$, which comes from a by-product of [Fournier and Guillin \(2015\)](#) (see [Theorem 11](#) in [Appendix C](#)).

Remark 5. An algorithmic consequence of this bias of approximation would tend not to use the Gaussian smoothed sliced in a generative modelling framework, i.e. $\hat{\theta}_n^\sigma = \arg \min_{\theta \in \mathcal{P}_\Theta(\mathbb{R}^d)} \hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \mu_\theta)$ with μ_θ a generative model.

Remark 6. The obtained complexity sample showing in [Proposition 1](#) is given only for the Gaussian smoothed Wasserstein distance, since the based Wasserstein distance satisfies a result for order statistic representation (see [Appendix 12](#)). For the other types of based divergence, one could establish a line of theoretical research to investigate such representations, specifically for MMD and Sinkhorn divergences.

Remark 7. Note that the sample complexity depends on the amount of smoothing through the moment of the Gaussian noise : the larger the amount of smoothing (and thus the privacy), the worse is the constant of the complexity. Hence, a trade-off on privacy and statistical estimation appears here as a reasonable guarantee on the differential privacy usually requires a large Gaussian variance.

Proposition 2. Fix $\sigma > 0, p \geq 1$. Then,

$$\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \mathbf{E}_{\nu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n)] \leq 3^{p-1} G_\sigma \text{SW}^p(\mu, \nu) + \Xi_{p,\sigma;\mu,\nu} \left(4 \log\left(\frac{n}{2}\right) + \frac{\log n}{n} \right),$$

and

$$G_\sigma \text{SW}^p(\mu, \nu) \leq 3^{p-1} \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \mathbf{E}_{\nu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n)] + \Xi_{p,\sigma;\mu,\nu} \left(4 \log\left(\frac{n}{2}\right) + \frac{\log n}{n} \right),$$

where $\Xi_{p,\sigma;\mu,\nu} = 3^{p-1} 2^{3p+1} (C_p + e^{5/3}) \{ (M_{2p}(\mu) + 2^p \sigma^{4p})^{1/2} + (M_{2p}(\nu) + 2^p \sigma^{4p})^{1/2} \}$.

Despite that our theoretical results are addressed only for Gaussian-smoothed sliced Wasserstein distance, our empirical results show that given other base divergences D , shows that the sample complexity of $G_\sigma \text{SD}^p$ is proportional to the one dimensional sample complexity of D^p ($p = 2$). [Figure 1](#) provides an empirical illustration of this statement.

3.2.3 Projection complexity

To compute the Gaussian-smoothed sliced divergence, one may resort to a Monte Carlo scheme to numerically approximate the integral in $G_\sigma \text{SD}^p(\mu, \nu)$. Towards this, let define the following sum:

$$\widehat{G_\sigma \text{SD}^p}(\mu, \nu) = \frac{1}{L} \sum_{l=1}^L D^p(\mathcal{R}_{\mathbf{u}_l} \mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}_l} \nu * \mathcal{N}_\sigma),$$




Figure 1: Measuring the divergence between two sets of samples in \mathbb{R}^{50} , of increasing size, randomly drawn from $\mathcal{N}(0, \mathbf{I})$. We compare three sliced divergences and their Gaussian-smoothed sliced versions with a $\sigma = 3$: (top) dimension has been set to $d = 50$; (bottom) sample complexity with different dimensions. This plot confirms that the complexity is dimension-independent.

where \mathbf{u}_l is a random vector uniformly drawn from \mathbb{S}^{d-1} , for $l = 1, \dots, L$. Theorem 2 shows that for a fixed dimension d , the root mean square error of Monte Carlo (MC) approximation is of order $O(\frac{1}{\sqrt{L}})$, which corresponds to the projection complexity. We denote by $u_d^{\otimes L}$ and the L -fold product extensions of the uniform measure u_d on the unit sphere.

Theorem 2. *Let $\sigma > 0, p \geq 1$. Then the error related to the MC-estimation of $G_\sigma \text{SD}^p$ is bounded as follows*

$$\mathbf{E}_{u_d^{\otimes L}} [|\widehat{G_\sigma \text{SD}^p}(\mu, \nu) - G_\sigma \text{SD}^p(\mu, \nu)|] \leq \frac{A(p, \sigma)}{\sqrt{L}},$$

where $A^2(p, \sigma) = \int_{\mathbb{S}^{d-1}} (D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) - \bar{\vartheta}_p)^2 u_d(\mathbf{u}) d\mathbf{u}$, with $\bar{\vartheta}_p = \int_{\mathbb{S}^{d-1}} D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u}$.

The term $A^2(p, \sigma)$ corresponds to the variance of $D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma)$ with respect to $\mathbf{u} \sim u_d$. It is worth to note that the precision of the Monte Carlo scheme approximation depends on the number of projections L and the variance of the evaluations of the divergence D^p . The estimation error decreases at the rate $L^{-1/2}$ according to the number of projections used to compute the smoothed sliced divergence.

Remark 8. *Given the above results, we provide a finer analysis of $G_\sigma \text{SW}^p(\mu, \nu)$'s sample complexity, that is bounded by the sample and projection complexities,*

$$\text{complexity}(G_\sigma \text{SW}^p) = O\left(\log\left(\frac{n}{2}\right) + \frac{\log n}{n} + \frac{A(p, \sigma)}{\sqrt{L}}\right).$$

If we consider the number of projections as $L = \lfloor n^\beta \rfloor$ for some $\beta \in (0, 1)$ then the overall complexity $\text{complexity}(G_\sigma \text{SW}^p(\mu, \nu)) = O(\log(\frac{n}{2}) + n^{-\beta/2})$.

3.3 Noise-level dependencies

The parameter σ of the Gaussian smoothing function \mathcal{N}_σ may significantly influence the attained privacy level. Hence, we provide theoretical results analyzing the effect of the noise level σ on the induced Gaussian-smoothed sliced divergence.




Figure 2: Absolute difference between the approximated Monte Carlo approximation of all divergences compared to the true one (evaluated with 10,000 number of projections). The two sets of 500 samples in \mathbb{R}^{50} are randomly drawn from $\mathcal{N}(0, \mathbf{I})$. The Gaussian-smoothed sliced divergences are parameterized with $\sigma = 3$.

Order relation. We first show that the noise level tends to reduce the difference between two distributions as measured using $G_\sigma \text{SD}^p(\mu, \nu)$ provided the base divergence D satisfies some mild assumptions.

Proposition 3. Let $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ and consider the noise levels σ_1, σ_2 such that $0 \leq \sigma_1 \leq \sigma_2 < \infty$. Assume that the base divergence D satisfies $D(\mu' * \mathcal{N}_{\sigma_2}, \nu' * \mathcal{N}_{\sigma_2}) \leq D(\mu' * \mathcal{N}_{\sigma_1}, \nu' * \mathcal{N}_{\sigma_1})$, for any $\mu', \nu' \in \mathcal{P}(\mathbb{R})$. Then, $G_{\sigma_2} \text{SD}^p(\mu, \nu) \leq G_{\sigma_1} \text{SD}^p(\mu, \nu)$.

Note that the assumption for the base divergence inequality holds for the Gaussian-smoothed Wasserstein distance [Nietert et al. \(2021\)](#). While we conjecture that it holds also for smoothed Sinkhorn and MMD, we leave the proofs for future works. Based on the property in Proposition 3, we show some specific properties of the metric with respect to the noise level σ .

Proposition 4. $G_\sigma \text{SD}^p(\mu, \nu)$ is decreasing with respect to σ and we have $\lim_{\sigma \rightarrow 0} G_\sigma \text{SD}^p(\mu, \nu) = D^p(\mu, \nu)$.

The proof of Proposition 4 comes straightforwardly from Proposition 3 by taking $\sigma_2 = \sigma$ and letting $\sigma_1 \rightarrow 0$. This property interestingly states that the $G_\sigma \text{SD}^p$ recovers the sliced divergence when the noise level vanishes.

Continuity. Now we analyze the continuity properties of some $G_\sigma \text{SD}^p(\mu, \nu)$ w.r.t. the noise level.

Proposition 5. For any two distributions μ and ν for which the sliced Wasserstein is well-defined, the Gaussian-smoothed sliced Wasserstein distance is continuous w.r.t. to σ .

Proposition 6. Assume that the kernel defining the maximum mean discrepancy (MMD) divergence is bounded. Then the Gaussian-smoothed sliced $G_\sigma \text{MMD}$ is continuous w.r.t. to σ .

The above propositions show that most distribution divergences are continuous with respect to σ under mild conditions.

4 Numerical Experiments

In this section, we report on a series of experiments that support the established theoretical results. We also highlight the usefulness of the findings in a context of privacy-preserving domain adaptation problem.

4.1 Supporting the theoretical results

Sample complexity. The first experiment (see Figure 1) analyzes the sample complexity of different base divergences. It shows that the sample complexity stays similar to the one of their original and sliced counterparts up to a constant (see Proposition 1). For this purpose, we have considered samples in \mathbb{R}^d randomly drawn from a Normal distribution $\mathcal{N}(0, \mathbf{I})$. For the Sinkhorn divergence, the entropy regularization has been set to 0.1 and for MMD, we used a Gaussian kernel for which the bandwidth has been set to the mean of all pairwise distances between samples. The number of projections has been fixed to $L = 50$ and we perform 20 runs per experiment. For the first study, the convergence rate has been evaluated by increasing the samples number up to 25,000 with fixed dimension $d = 50$. For the second one, we vary both the dimension and the number of samples.

Figure 1 shows the sample complexity of some sliced divergences, respectively noted as SWD, SKD and MMD for Sliced Wasserstein distance, Sinkhorn divergence and Maximum Mean discrepancy and their Gaussian-smoothed sliced versions, named as GS SWD, GS SKD and GS MMD. On the top plot, we can see that all Gaussian-smoothed sliced divergences preserve the complexity rate with just a slight to moderate overhead. The worst difference is for Sinkhorn divergence, while MMD almost comes for free in term of complexity. From the bottom plot where sample complexities for different dimensions d are given, we confirm the finding that Gaussian smoothing keeps the independence of the convergence rate to the dimension of sliced divergences.

Two other experiments on the sample complexity and identity of indiscernibles are also reported in the supplementary material.

Projection complexity. We have also investigated the impact of the number of projections when estimating the distance between two sets of 500 samples drawn from the same distribution, $\mathcal{N}(0, \mathbf{I})$. Figure 2 plots the approximation error between the true expectation of the sliced divergences (computed for a number of $L = 10,000$ projections) and its approximated versions. We remark that, for all methods, the error ranges within 10-fold when approximating with 50 projections and decreases with the number of projections.

Performance path on the impact of the noise parameter. Since the Gaussian smoothing parameter σ is key in a privacy preserving context, as it impacts on the level of privacy of the Gaussian mechanism, we have analyzed its impact on the smoothed sliced divergence. We have reproduced the experiment for the sample complexity but with different values of σ . The number of projections has been set to 50. Figure 3 shows these sample complexities. The first very interesting point to note is that the smoothing parameter has almost no effect on the GS MMD sample complexity. For the GS SWD and GS SKD divergences, instead, the smoothing tends to increase the divergence at fixed number of samples. Another interpretation is that to achieve a given value of divergence, one needs more far samples when the smoothing is larger (*i.e.* for getting a given divergence value at $\sigma = 5$, one needs almost 10-fold more samples for $\sigma = 15$). This overhead of samples needed when smoothing




Figure 3: Measuring the divergence between two sets of samples in \mathbb{R}^{50} drawn from $\mathcal{N}(0, \mathbf{I})$. We plot the sample complexity for different Gaussian-smoothed sliced divergence at different level of noises.




Figure 4: Domain adaptation performances using different divergences on distributions with respect to the Gaussian smoothing. (Left) USPS to MNIST. (Middle) Office-31 Webcam to DSLR. (Right) Office-31 Amazon to Webcam.

increases is properly described, for the Gaussian-smoothed sliced SWD in our Proposition 1, as the sample complexity depends on the moments of the Gaussian.

As for conclusion from these analyses, we highlight that the Gaussian-smoothed sliced MMD seems to present several strong benefits: its sample complexity does not depend on the dimension and seems to be the best one among the divergence we considered. More interestingly, it is not impacted by the amount of Gaussian smoothing and thus not impacted by a desired privacy level.

4.2 Domain adaptation with G_σ SW

As an application, we have considered the problem of unsupervised domain adaptation for a classification task. In this context, given source examples \mathbf{X}_s and their label \mathbf{y}_s and unlabeled target examples \mathbf{X}_t , our goal is to design a classifier $h(\cdot)$ learned from the source examples that generalizes well on the target ones. A classical approach consists in learning a representation mapping $g(\cdot)$ that leads to invariant latent representations, invariance being measured as a distance between empirical distributions of mapped source and target samples. Formally, this leads to the following problem

$$\min_{g,h} \{ \mathcal{L}_c(h(g(\mathbf{X}_s)), \mathbf{y}_s) + \mathcal{D}(g(\mathbf{X}_s), g(\mathbf{X}_t)) \}$$

where \mathcal{L}_c can be the cross-entropy loss or a quadratic loss and \mathcal{D} a divergence between empirical distributions, in our case, \mathcal{D} will be any Gaussian-smoothed sliced divergence. We solve this problem




Figure 5: Domain adaptation performances using different divergences on distributions with respect to the Gaussian smoothing using **one-epoch-fine-tuned** models. (Left) USPS to MNIST. (Middle) Office-31 Webcam to DSLR. (Right) Office-31 Amazon to Webcam.

through stochastic gradient descent, similarly to many approaches that use sliced Wasserstein distance as a distribution distance [Lee et al. \(2019\)](#). Note that, in practice, using a smoothed divergence preserves the privacy of the target samples as shown by ([Rakotomamonjy and Ralaivola, 2021](#)).

When performing such model adaptation, a privacy/utility trade-off that has to be handled. In practice, one would prefer the most private model while not hurting its performance. Hence, one would seek the largest noise level $\sigma > 0$ to use while preserving accuracy on target domain. Hence, it is useful to evaluate how the model performs on a range of noise level (hence, privacy level). This can be computationally expensive as it requires to fully train several models on hundreds of epochs. Instead, we leverage on the continuity of our G_σ SD to employ a fine-tuning strategy: we train a domain adaptation model for the largest desired value of σ (over the full number of epochs) and when σ is decreased, we just fine-tune the lasted model by training on only one epoch.

Our experiments evaluate the studied Gaussian-smoothed sliced divergences in classical unsupervised domain adaptation. We have considered two datasets: a handwritten digit recognition (USPS/MNIST) and Office 31 datasets.

In our first analysis, we have compared our G_σ SD performances with non-smoothed divergences. The first one is the sliced Wasserstein distance (SWD) [Lee et al. \(2019\)](#) and the second one is the Jenson-Shannon approximation based on adversarial approach, known as DANN [Ganin and Lempitsky \(2015\)](#). For all methods and for each dataset, we used the same neural network architecture for representation mapping and for classification. Approaches differ only on how distance between distributions have been computed. Here for each noise value σ , we have trained the model from scratch for 100 epochs. Results are depicted in Figure 4. For the two problems, we can see that performances obtained with the Gaussian-smoothed sliced Wasserstein or MMD divergences are similar to those obtained with DANN or SWD across all ranges of noise. The smoothed version of Sinkhorn is less stable and induces a slight loss of performance. Owing to the metric property and the induced weak topology, the privacy preservation comes almost without loss of performance in this domain adaptation context.

In the second analysis, we have studied the privacy/utility trade-off when fine-tuning models, using only one epoch, for decreasing values of σ . Results are shown in Figure 5. They highlight that depending on the data and the used smoothed divergence, performance varies between one percent for Office 31 to four percent for USPS to MNIST. Note that except for the largest value of σ , we are

training a model using only one epoch instead of a hundred. A very large gain in complexity is thus achieved for swiping the full range of noise level. Hence depending on the importance this slight drop in performance will have, it is worth using a large value of σ and preserving strong privacy or go through a validation procedure of several (cheaply obtained) models.

5 Conclusion

This work provided the properties of Gaussian-smoothed sliced divergences for comparing distributions. We derived several theoretical results related to their topological and statistical properties and showed, under mild conditions on their base divergences, the smoothing and slicing operations preserves the metric property. From a statistical point of view, we introduced the double empirical distribution and focused on the sample complexity of the smoothed sliced Wasserstein distance. We proved that it suffers from an unavoidable bias of approximation. We analyzed the behavior of these divergences on domain adaptation problems and confirm the fact that using those divergences yields only to slight loss of performances while preserving privacy. An important direction for future research is establishing order statistic representations for Sinkhorn and MMD divergences. Furthermore, in the obtained bound we use upper bound of higher moments of the smoothing distribution, that would think to consider non Gaussian smoothing distribution enjoying this property.

Acknowledgments

Mokhtar Z. Alaya's work has been supported by the Research Direction of University of Technology of Compiègne under the Junior Researchers Program "UTC – Appel soutien Jeunes MC".

A Proofs and additional theoretical results

In the following sections, we give the proofs of the theoretical guarantees given in the main of the paper. For a sake of completeness, we add other results that we consider interesting.

A.1 Proof of Theorem 1: $G_\sigma\text{SD}$ is a proper metric on $\mathcal{P}_p(\mathbb{R}^d) \times \mathcal{P}_p(\mathbb{R}^d)$

Before starting the proof, we add this notation: the characteristic function of a probability distribution $\mu \in \mathcal{P}(\mathbb{R}^d)$ is $\varphi_\mu(t) = \mathbf{E}_\mu[e^{iX^\top t}]$. Given this definition, similarly to the Fourier transform, the characteristic function of the convolution of two probability distributions reads as $\varphi_{\nu * \mu}(t) = \varphi_\nu(t) \cdot \varphi_\mu(t)$.

- *Non-negativity (or symmetry)*. The non-negativity (or symmetry) follows directly from the non-negativity (or symmetry) of D^p , see Definition 3.

- *Identity property*. If the base divergence D^p satisfies the identity property in one dimensional measures, then for any $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ and $\mathbf{u} \in \mathbb{S}^{d-1}$, one has that $D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) = 0$, hence, by Definition 3, $G_\sigma\text{SD}^p(\mu, \mu) = 0$. Let us now prove the fact that for any $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, $G_\sigma\text{SD}^p(\mu, \nu) = 0$ entails $\mu = \nu$ a.s. On one hand, $G_\sigma\text{SD}^p(\mu, \nu) = 0$ gives the fact that $D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) = 0$ for u_d -almost every $\mathbf{u} \in \mathbb{S}^{d-1}$, hence $\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma = \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma$ for u_d -almost every $\mathbf{u} \in \mathbb{S}^{d-1}$. Following the techniques in proof of Proposition 5.1.2 in Bonnotte (2013), for any measure $\eta \in \mathcal{P}(\mathbb{R}^m)$ (with $m \geq 1$), $\mathcal{F}[\eta](\cdot)$ stands for the Fourier transform of η and is given as $\mathcal{F}[\eta](\mathbf{v}) = \int_{\mathbb{R}^m} e^{-i\mathbf{s}^\top \mathbf{v}} d\eta(\mathbf{s})$ for any $\mathbf{v} \in \mathbb{R}^m$. Then

$$\begin{aligned} \mathcal{F}[\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma](v) &= \int_{\mathbb{R}} e^{-ivt} d(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma)(t) \\ &= \int_{\mathbb{R}} \int_{\mathbb{R}} e^{-i(r+t)v} d\mathcal{R}_{\mathbf{u}}\mu(r) d\mathcal{N}_\sigma(t) \quad (\text{by the definition of the convolution operator}) \\ &= \int_{\mathbb{R}^d} \int_{\mathbb{R}} e^{-i(\langle \mathbf{u}, \mathbf{s} \rangle + t)v} d\mu(\mathbf{s}) d\mathcal{N}_\sigma(t) \quad (\text{by the definition of Radon Transform}) \\ &= \int_{\mathbb{R}} e^{-itv} d\mathcal{N}_\sigma(t) \int_{\mathbb{R}^d} e^{-i(\langle \mathbf{u}, \mathbf{s} \rangle)v} d\mu(\mathbf{s}) \\ &= \mathcal{F}[\mathcal{N}_\sigma](v) \mathcal{F}[\mu](v\mathbf{u}). \end{aligned}$$

Since for u_d -almost every $\mathbf{u} \in \mathbb{S}^{d-1}$, $\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma = \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma$, and hence $\mathcal{F}[\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma] = \mathcal{F}[\mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma] \Leftrightarrow \mathcal{F}[\mathcal{N}_\sigma]\mathcal{F}[\mu] = \mathcal{F}[\mathcal{N}_\sigma]\mathcal{F}[\nu]$ (by the Fourier transform of the convolution) $\Leftrightarrow \mathcal{F}[\mu] = \mathcal{F}[\nu]$. Since the Fourier transform is injective, we conclude that $\mu = \nu$.

- *Triangle inequality*. Assume that D is a metric and let $\mu, \nu, \eta \in \mathcal{P}_p(\mathbb{R}^d)$. We then have

$$\begin{aligned} G_\sigma\text{SD}(\mu, \nu) &= \left\{ \int_{\mathbb{S}^{d-1}} D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u} \right\}^{1/p} \\ &\leq \left\{ \int_{\mathbb{S}^{d-1}} \left(D(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\eta * \mathcal{N}_\sigma) + D(\mathcal{R}_{\mathbf{u}}\eta * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) \right)^p u_d(\mathbf{u}) d\mathbf{u} \right\}^{1/p} \\ &\stackrel{(*)}{\leq} \left\{ \int_{\mathbb{S}^{d-1}} \left(D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\eta * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u} \right)^{1/p} \right. \\ &\quad \left. + \left\{ \int_{\mathbb{S}^{d-1}} D^p(\mathcal{R}_{\mathbf{u}}\eta * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u} \right\}^{1/p} \right. \end{aligned}$$

$$= G_\sigma \text{SD}(\mu, \eta) + G_\sigma \text{SD}(\eta, \nu),$$

where inequality in (\star) follows from the application of Minkowski inequality.

A.2 $G_\sigma \text{SD}$ metrizes the weak topology and is lower semi-continuous

Now, we establish under which conditions on the divergence D , the convergence of a sequence in $G_\sigma \text{SD}$ implies weak convergence in $\mathcal{P}_p(\mathbb{R}^d)$. We say that $\{\mu_k\}_{k \in \mathbb{N}}$ converges weakly to μ and write, $\mu_k \Rightarrow \mu$, if $\int f(x) d\mu_k(x) \rightarrow \int f(x) d\mu(x)$, as $k \rightarrow \infty$, for every f in the space of all bounded continuous real functions.

Theorem 3. *Let $\sigma > 0, p \geq 1$, $\mu \in \mathcal{P}_p(\mathbb{R}^d)$, and $\{\mu_k \in \mathcal{P}_p(\mathbb{R}^d)\}_{k \in \mathbb{N}}$ a sequence of distributions. Assume that the divergence D is bounded and metrizes the weak topology on $\mathcal{P}(\mathbb{R})$. Then, $\lim_{k \rightarrow \infty} G_\sigma \text{SD}(\mu_k, \mu) = 0$ if and only if $\mu_k \Rightarrow \mu$.*

Note that Theorem 3 extends the results of [Nadjahi et al. \(2020\)](#) to Gaussian-smoothed distributions, as we retrieve them as a special case for $\sigma = 0$. In addition, based on Theorem 3.2 by [Lin et al. \(2021\)](#) and the above, we can also claim that the Gaussian-smoothed SWD metrizes the weak convergence.

Proposition 7. *Let $\sigma > 0, p \geq 1$ and assume that the base divergence D is lower semi-continuous w.r.t. the weak topology in $\mathcal{P}(\mathbb{R})$. Then, $G_\sigma \text{SD}$ is lower semi-continuous with respect to the weak topology in $\mathcal{P}_p(\mathbb{R}^d)$.*

When the base divergence D is equal to the Wasserstein distance W_p , that is lower semi-continuous ([Villani, 2009](#)), then Proposition 7 shows that the smoothed sliced Wasserstein distance is semi-lower continuous too.

A.2.1 Proof of Theorem 3

The proof is done by double implications:

“ \Rightarrow ” Assume that $\mu_k \Rightarrow \mu$. Fix $\mathbf{u} \in \mathbb{S}^{d-1}$, the mapping $\mathbf{u} \mapsto \mathcal{R}_{\mathbf{u}}$ is continuous from \mathbb{R}^d to \mathbb{R} , then an application of continuous mapping Theorem 6 entails that $\mathcal{R}_{\mathbf{u}}\mu_k \Rightarrow \mathcal{R}_{\mathbf{u}}\mu$. By Lévy’s continuity Theorem 7 $\mathcal{R}_{\mathbf{u}}\mu_k * \mathcal{N}_\sigma \Rightarrow \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma$. Therefore, $\lim_{k \rightarrow \infty} D(\mathcal{R}_{\mathbf{u}}\mu_k, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) = 0$. Since we suppose that the divergence D is bounded, then there exists $K \geq 0$ such that for any k , $D^p(\mathcal{R}_{\mathbf{u}}\mu_k, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) \leq K$. An application of bounded convergence Theorem 5 yields

$$\lim_{k \rightarrow \infty} G_\sigma \text{SD}^p(\mu_k, \mu) = \lim_{k \rightarrow \infty} \int_{\mathbb{S}^{d-1}} D^p(\mathcal{R}_{\mathbf{u}}\mu_k * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u} = 0.$$

“ \Leftarrow ” (By contrapositive). Suppose that μ_k doesn’t converge weakly to μ and assume that $\lim_{k \rightarrow \infty} G_\sigma \text{SD}^p(\mu_k, \mu) = 0$. On one hand, since \mathbb{R}^d is a complete separable space then using Theorem 10 that tells that weak convergence is equivalent to the convergence corresponding to Lévy-Prokhorov distance Λ (see Definition 5), there exists $\varepsilon > 0$ and a subsequence $\{\mu_{s(k)}\}_{k \in \mathbb{N}}$ such that $\Lambda(\mu_{s(k)}, \mu) > \varepsilon$. On the other hand, we have $\lim_{k \rightarrow \infty} G_\sigma \text{SD}^p(\mu_{s(k)}, \mu) = 0$, that is equivalent to $\{D(\mathcal{R}_{\mathbf{u}}\mu_{s(k)} * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma)\}_k$ converges to 0 in $L^p(\mathbb{S}^{d-1}) = \{f : \mathbb{S}^{d-1} \rightarrow \mathbb{R} \mid \int_{\mathbb{S}^{d-1}} f(\mathbf{u}) u_d(\mathbf{u}) d\mathbf{u} < \infty\}$. By application of Proposition 10, there exists a subsequence $\{\mu_{s(t(k))}\}_k$ such that $\lim_{k \rightarrow \infty} D(\mathcal{R}_{\mathbf{u}}\mu_{s(t(k))} * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) = 0$ almost everywhere

for all $\mathbf{u} \in \mathbb{S}^{d-1}$. Recall that the divergence D metrizes the weak convergence in $\mathcal{P}(\mathbb{R})$ then $\mathcal{R}_{\mathbf{u}}\mu_{s(t(k))} * \mathcal{N}_\sigma \Rightarrow \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma$ almost everywhere for all $\mathbf{u} \in \mathbb{S}^{d-1}$. Therefore, $\mathcal{R}_{\mathbf{u}}\mu_{s(t(k))} \Rightarrow \mathcal{R}_{\mathbf{u}}\mu$ almost everywhere for all $\mathbf{u} \in \mathbb{S}^{d-1}$. Using Cramér-Wold device, see Proposition 11, we get $\mu_{s(t(k))} \Rightarrow \mu$. Since the Lévy-Prokhorov distance metrizes the weak convergence, it entails that $\lim_{k \rightarrow \infty} \Lambda(\mu_{s(t(k))}, \mu_k) = 0$, that contradicts the fact that $\Lambda(\mu_{s(k)}, \mu) > \varepsilon$. We then conclude by contrapositive that $\mu_k \Rightarrow \mu$.

A.2.2 Proof of Proposition 7

Recall that the base divergence D is lower semi-continuous w.r.t. the weak topology in $\mathcal{P}(\mathbb{R})$, namely for every sequence of measures $\{\mu'_k\}_{k \in \mathbb{N}}$ and $\{\nu'_k\}_{k \in \mathbb{N}}$ in $\mathcal{P}(\mathbb{R})$ such that $\mu'_k \Rightarrow \mu'$ and $\nu'_k \Rightarrow \nu'$, one has $D(\mu', \nu') \leq \liminf_{k \rightarrow \infty} D(\mu'_k, \nu'_k)$.

Now, let $\{\mu_k\}_{k \in \mathbb{N}}$ and $\{\nu_k\}_{k \in \mathbb{N}}$ are two sequences of measure in $\mathcal{P}_p(\mathbb{R}^d)$ such that $\mu_k \Rightarrow \mu$ and $\nu_k \Rightarrow \nu$. By continuous mapping theorem 6 and Levy's continuity theorem, we obtain $\mathcal{R}_{\mathbf{u}}\mu_k * \mathcal{N}_\sigma \Rightarrow \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma$ and $\mathcal{R}_{\mathbf{u}}\nu_k * \mathcal{N}_\sigma \Rightarrow \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma$ for all $\mathbf{u} \in \mathbb{S}^{d-1}$. Since the base divergence D is a lower semi-continuous with respect to weak topology in $\mathcal{P}(\mathbb{R})$, then

$$D^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) \leq \left(\liminf_{k \rightarrow \infty} D(\mathcal{R}_{\mathbf{u}}\mu_k * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu_k * \mathcal{N}_\sigma) \right)^p \leq \liminf_{k \rightarrow \infty} D^p(\mathcal{R}_{\mathbf{u}}\mu_k * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu_k * \mathcal{N}_\sigma).$$

It gives

$$G_\sigma \text{SD}^p(\mu, \nu) \leq \int_{\mathbb{S}^{d-1}} \liminf_{k \rightarrow \infty} D^p(\mathcal{R}_{\mathbf{u}}\mu_k * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu_k * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u}.$$

Furthermore, by application of Fatou's lemma 8, we get

$$G_\sigma \text{SD}^p(\mu, \nu) \leq \liminf_{k \rightarrow \infty} \int_{\mathbb{S}^{d-1}} D^p(\mathcal{R}_{\mathbf{u}}\mu_k * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu_k * \mathcal{N}_\sigma) u_d(\mathbf{u}) d\mathbf{u} = \liminf_{k \rightarrow \infty} G_\sigma \text{SD}^p(\mu_k, \nu_k),$$

which is the desired result.

A.3 Proofs of statistical properties

A.3.1 Proof of Lemma 1

Straightforwardly, for every Borelian $I \in \mathcal{B}(\mathbb{R})$, we have

$$\begin{aligned} \mathcal{R}_{\mathbf{u}}\hat{\mu}_n * \mathcal{N}_\sigma(I) &= \int_r \int_s \mathbf{1}_I(r+s) d\mathcal{R}_{\mathbf{u}}\hat{\mu}_n d\mathcal{N}_\sigma(s) \\ &= \int_r \int_s \mathbf{1}_I(r+s) d\left\{ \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i} \right\}(r) d\mathcal{N}_\sigma(s) \\ &= \frac{1}{n} \sum_{i=1}^n \int_s \mathbf{1}_I(\mathbf{u}^\top X_i + s) d\mathcal{N}_\sigma(s) \\ &= \frac{1}{n} \sum_{i=1}^n \int_s \mathbf{1}_I(\mathbf{u}^\top X_i + s) f_{\mathcal{N}_\sigma}(s) ds \\ &= \frac{1}{n} \sum_{i=1}^n \int_{s'} \mathbf{1}_I(s') f_{\mathcal{N}_\sigma}(s' - \mathbf{u}^\top X_i) ds' \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{n} \sum_{i=1}^n \int_{s'} \mathbf{1}_I(s') f_{\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)}(s') ds' \quad (\text{since } f_{\mathcal{N}_\sigma}(s' - \mathbf{u}^\top X_i) = f_{\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)}(s')) \\
&= \frac{1}{n} \sum_{i=1}^n \int_{s'} \mathbf{1}_I(s') d\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)(s') \\
&= \frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)(I).
\end{aligned}$$

Thanks to Theorem of Cramér and Wold ([Cramér and Wold, 1936](#)), we conclude the equality between the measures $\mathcal{R}_{\mathbf{u}} \hat{\mu}_n * \mathcal{N}_\sigma = \frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)$.

A.3.2 Additional Proposition

Proposition 8. *Assume that the divergence D equal to the Wasserstein distance W_p . Then,*

$$G_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n) \leq \frac{1}{n} \sum_{i=1}^n \int_{\mathbb{S}^{d-1}} W_p^p(\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)) u_d(\mathbf{u}) d\mathbf{u}.$$

In particular, $G_\sigma \text{SW}^2(\hat{\mu}_n, \hat{\nu}_n) \leq \frac{1}{n} \sum_{i=1}^n \|X_i - Y_i\|^2$.

Proposition 8 ensures that $G_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n)$ is upper bounded by the mean of the SWD between Gaussian distributions centered on the projected samples and with variance σ^2 . In the case of $p = 2$ and using the closed form of Wasserstein distance between Gaussian distributions (Bures distance), we have the mean of the squared L_2 -norms between samples as an upper bound for $G_\sigma \text{SW}^2(\hat{\mu}_n, \hat{\nu}_n)$. *Proof of Proposition 8.* Setting the divergence D equals to the Wasserstein distance W_p . By Lemma 1, we have

$$G_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n) = \int_{\mathbb{S}^{d-1}} W_p^p\left(\frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \frac{1}{n} \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)\right) u_d(\mathbf{u}) d\mathbf{u}.$$

Then using scaling property of W_p (i.e. $W_p(a\alpha, a\beta) = |a| W_p(\alpha, \beta)$ for all probability measures α, β and $a \in \mathbb{R}$), hence

$$G_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n) = \frac{1}{n^p} \int_{\mathbb{S}^{d-1}} W_p^p\left(\sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)\right) u_d(\mathbf{u}) d\mathbf{u}.$$

Remark that $\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2)$ and $\mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)$ are independent then from an additive property, see [Panaretos and Zemel \(2019\)](#), we get

$$\begin{aligned}
W_p^p\left(\sum_{i=1}^n \mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \sum_{j=1}^n \mathcal{N}(\mathbf{u}^\top Y_j, \sigma^2)\right) &\leq \left(\sum_{i=1}^n W_p(\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2))\right)^p \\
&\leq n^{p-1} \sum_{i=1}^n W_p^p(\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)),
\end{aligned}$$

where the last inequality is due to $\left(\sum_{i=1}^n a_i\right)^p \leq n^{p-1} \sum_{i=1}^n a_i^p$ for $a_i \geq 0$ and $p \geq 1$. Then,

$$G_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n) \leq \frac{1}{n} \sum_{i=1}^n \int_{\mathbb{S}^{d-1}} W_p^p(\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)) u_d(\mathbf{u}) d\mathbf{u}.$$

For $p = 2$, using the closed form of Wasserstein distance between Gaussian, we have

$$W_2^2(\mathcal{N}(\mathbf{u}^\top X_i, \sigma^2), \mathcal{N}(\mathbf{u}^\top Y_i, \sigma^2)) = |\mathbf{u}^\top (X_i - Y_i)|^2,$$

hence

$$G_\sigma \text{SW}^2(\hat{\mu}_n, \hat{\nu}_n) \leq \frac{1}{n} \sum_{i=1}^n \int_{\mathbb{S}^{d-1}} |\mathbf{u}^\top (X_i - Y_i)|^2 u_d(\mathbf{u}) d\mathbf{u} \leq \frac{1}{n} \sum_{i=1}^n \|X_i - Y_i\|^2.$$

A.3.3 Order statistics representation of $\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n)$

Lemma 2. For fixed $\mathbf{u} \in \mathbb{S}^{d-1}$, one has

$$W_p^p(\hat{\mu}_n, \hat{\nu}_n) = \frac{1}{n} \sum_{l=1}^n |T_{(l)}^x - T_{(l)}^y|^p = \frac{1}{n} \sum_{l=1}^n \left| (\mathbf{u}^\top X_i + Z_i^x)_{(l)} - (\mathbf{u}^\top Y_i + Z_i^y)_{(l)} \right|^p.$$

and the double-empirical smoothed Gaussian sliced Wasserstein satisfies

$$\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n) = \frac{1}{n} \sum_{l=1}^n \int_{\mathbb{S}^{d-1}} \left| (\mathbf{u}^\top X_i + Z_i^x)_{(l)} - (\mathbf{u}^\top Y_i + Z_i^y)_{(l)} \right|^p u_d(\mathbf{u}) d\mathbf{u}.$$

Proof of Lemma 2. By Lemma 3, we obtain

$$W_p^p(\hat{\mu}_n, \hat{\nu}_n) = W_p^p\left(\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}, \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top Y_i + Z_i^y}\right) = \frac{1}{n} \sum_{l=1}^n \left| (\mathbf{u}^\top X_i + Z_i^x)_{(l)} - (\mathbf{u}^\top Y_i + Z_i^y)_{(l)} \right|^p.$$

Then,

$$\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n) = \frac{1}{n} \sum_{l=1}^n \int_{\mathbb{S}^{d-1}} \left| (\mathbf{u}^\top X_i + Z_i^x)_{(l)} - (\mathbf{u}^\top Y_i + Z_i^y)_{(l)} \right|^p u_d(\mathbf{u}) d\mathbf{u}.$$

A.3.4 Proof of Proposition 1

For this proof and the following ones, we use frequently the triangle inequality for Wasserstein distance and the quantities $\hat{\mu}_n$ and $\frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}$, the empirical measure probabilities of $\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}$ and $\mathcal{R}_{\mathbf{u}} \mu * \mathcal{N}_\sigma$, respectively.

On one hand, using triangle inequality of Wasserstein distance, we have

$$\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \mu)] = \int_{\mathbb{S}^{d-1}} \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [W_p^p(\hat{\mu}_n, \mathcal{R}_{\mathbf{u}} \mu * \mathcal{N}_\sigma)] u_d(\mathbf{u}) d\mathbf{u}$$

$$\begin{aligned}
&= \int_{\mathbb{S}^{d-1}} \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\mathbf{W}_p^p \left(\frac{1}{n} \sum_{j=1}^n \delta_{\mathbf{u}^\top X_i + Z_j^x}, \mathcal{R}_{\mathbf{u}} \mu * \mathcal{N}_\sigma \right) \right] u_d(\mathbf{u}) \mathbf{d}\mathbf{u} \\
&\leq \int_{\mathbb{S}^{d-1}} 2^{p-1} \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\mathbf{W}_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}, \frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x} \right) \right. \\
&\quad \left. + \mathbf{W}_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}, \mathcal{R}_{\mathbf{u}} \mu * \mathcal{N}_\sigma \right) \right] u_d(\mathbf{u}) \mathbf{d}\mathbf{u}.
\end{aligned}$$

• Control of $\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\mathbf{W}_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}, \frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x} \right) \right]$.

By application of Theorem 12 for $\hat{\alpha}_n = \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}$ and $\hat{\beta}_n = \frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}$, we have

$$\begin{aligned}
&\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\mathbf{W}_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}, \frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x} \right) \right] \\
&= \frac{1}{n} \sum_{l=1}^n \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| (\mathbf{u}^\top X_i + Z_i^x)_{(l)} - (U_i^x + Q_i^x)_{(l)} \right|^p \right] \\
&\leq \frac{2^{p-1}}{n} \sum_{l=1}^n \left\{ \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| (\mathbf{u}^\top X_i + Z_i^x)_{(l)} \right|^p \right] + \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| (U_i^x + Q_i^x)_{(l)} \right|^p \right] \right\}.
\end{aligned}$$

On other hand, using Theorem 13 with $\rho = \frac{1}{2}$ we get $C(\rho) = C(\frac{1}{2}) = \sqrt{2}e^{5/3}$, then

$$\begin{aligned}
\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| (\mathbf{u}^\top X_i + Z_i^x)_{(l)} \right|^p \right] &\leq \sqrt{2}e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} \left(\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| \mathbf{u}^\top X_1 + Z_1^x \right|^{2p} \right] \right)^{1/2} \\
&\leq \sqrt{2}e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} \left(2^{2p-1} \left\{ \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| \mathbf{u}^\top X_1 \right|^{2p} \right] \right. \right. \\
&\quad \left. \left. + \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| Z_1^x \right|^{2p} \right] \right\} \right)^{1/2} \\
&\leq \sqrt{2}e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} \left(2^{2p-1} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|)) \right)^{1/2} \\
&\leq 2^p e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2},
\end{aligned}$$

where we recall that for any integer k , $M_k(\mu) = \int_{\mathbb{R}^d} \|x\|^k \mathbf{d}\mu(x)$ and $M_k(|\mathcal{N}_\sigma|) = \int_{\mathbb{R}} |t|^k \mathbf{d}\mathcal{N}_\sigma(t)$. Similarly,

$$\begin{aligned}
\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| (U_i^x + Q_i^x)_{(l)} \right|^p \mathbf{b} \right] &\leq \sqrt{2}e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} \left(\mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| U_1^x + Q_1^x \right|^{2p} \right] \right)^{1/2} \\
&\leq \sqrt{2}e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} \left(2^{2p-1} \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| U_1^x \right|^{2p} \right] + \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \left[\left| Q_1^x \right|^{2p} \right] \right)^{1/2} \\
&\leq \sqrt{2}e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} \left(2^{2p-1} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|)) \right)^{1/2} \\
&\leq 2^p e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2}.
\end{aligned}$$

Therefore,

$$\begin{aligned} \mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}} [|(\mathbf{u}^\top X_i + Z_i^x)_{(l)}|^p] + \mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}} [|(U_i^x + Q_i^x)_{(l)}|^p] \\ \leq 2^{p+1} e^{5/3} \frac{n+1}{\sqrt{l(n-l+1)}} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2}. \end{aligned}$$

Hence,

$$\begin{aligned} \mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}} \left[W_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{u}^\top X_i + Z_i^x}, \frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x} \right) \right] \\ \leq \frac{2^{2p}}{n} e^{5/3} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}}. \end{aligned}$$

- *Control of* $\mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}} \left[W_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}, \mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma \right) \right]$.

By application of Theorem 11, we get

$$\mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}} \left[W_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}, \mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma \right) \right] \leq C_{p,q} M_q^{p/q} (\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma) \Delta_n(p, q).$$

Moreover, Let us first upper bound the q -th moment of $M_q(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma)$, for all $k \geq 1$. For all $\mathbf{u} \in \mathbb{S}^{d-1}$, one has

$$\begin{aligned} M_q(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma) &= \int_{\mathbb{R}} |t|^q d(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma)(t) \\ &= \int_{\mathbb{R}} \int_{\mathbb{R}} |r+t|^q d\mathcal{R}_{\mathbf{u}\mu}(r) d\mathcal{N}_\sigma(t) \\ &= \int_{\mathbb{R}^d} \int_{\mathbb{R}} |\langle \mathbf{u}, s \rangle + t|^q d\mu(s) d\mathcal{N}_\sigma(t). \end{aligned}$$

Using the elementary inequality $(a+b)^q \leq 2^{q-1}(a^q + b^q)$ for $k \geq 1$, $a \geq 0$, and $b \geq 0$, we obtain

$$\begin{aligned} M_q(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma) &\leq 2^{q-1} \int_{\mathbb{R}^d} \int_{\mathbb{R}} (|\langle \mathbf{u}, s \rangle|^q + |t|^q) d\mu(s) d\mathcal{N}_\sigma(t) \\ &\leq 2^{q-1} \left(\|\mathbf{u}\| \int_{\mathbb{R}^d} \|s\|^q d\mu(s) + \int_{\mathbb{R}} |t|^q d\mathcal{N}_\sigma(t) \right) \\ &\leq 2^{q-1} \left(\int_{\mathbb{R}^d} \|s\|^q d\mu(s) + \int_{\mathbb{R}} |t|^q d\mathcal{N}_\sigma(t) \right) \\ &= 2^{q-1} (M_q(\mu) + M_q(|\mathcal{N}_\sigma|)). \end{aligned}$$

Then we arrive at

$$\mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}} \left[W_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{U_i^x + Q_i^x}, \mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma \right) \right] \leq 2^{q-1} C_{p,q} (M_q(\mu) + M_q(|\mathcal{N}_\sigma|))^{p/q} \Delta_n(p, q).$$

This implies that

$$\begin{aligned} \mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}}[\hat{\mathbf{G}}_\sigma \text{SW}^p(\hat{\mu}_n, \mu)] &\leq \frac{2^{3p-1}}{n} e^{5/3} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}} + \\ &\quad + 2^{(p+q)-2} C_{p,q} (M_q(\mu) + M_q(|\mathcal{N}_\sigma|))^{p/q} \Delta_n(p, q). \end{aligned}$$

For $q = 2p$ we have $\Delta_n(p, q) = \frac{\log n}{n}$, it yields

$$\begin{aligned} \mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}}[\hat{\mathbf{G}}_\sigma \text{SW}^p(\hat{\mu}_n, \mu)] &\leq \frac{2^{3p-1}}{n} e^{5/3} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}} + \\ &\quad + 2^{3p-2} C_p (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \frac{\log(n)}{n} \\ &\leq \frac{2^{3p-1}}{n} e^{5/3} (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}} + \\ &\quad + 2^{3p-2} C_p (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \frac{\log(n)}{n} \\ &\leq 2^{3p-1} (C_p + e^{5/3}) (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \\ &\quad \times \left(\frac{1}{n} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}} + \frac{\log(n)}{n} \right). \end{aligned}$$

Finally, for any $q \geq 1$ the q -th moment of $|\mathcal{N}_\sigma|$ satisfies

$$\mathbf{E}[|\mathcal{N}_\sigma|^q] = \frac{2^q \Gamma((q+1)/2)}{\Gamma(1/2)} \sigma^{2q} \leq 2^{q/2} \sigma^{2q},$$

where $\Gamma : \mathbb{R} \rightarrow \mathbb{R}$ is the Gamma function expressed as $\Gamma(v) = \int_0^\infty t^{v-1} e^{-t} dt$. Now let us examine. Let $\lceil n/2 \rceil$ be the upper integer part of $n/2$. Then, for any $1 \leq l \leq \lceil n/2 \rceil$, one has $l \leq n-l+1$. This implies

$$\sum_{l=1}^n \frac{1}{\sqrt{l(n-l+1)}} \leq 2 \sum_{l=1}^{\lceil n/2 \rceil} \frac{1}{\sqrt{l(n-l+1)}} \leq 2 \sum_{l=1}^{\lceil n/2 \rceil} \frac{1}{l} \leq 4 \log\left(\frac{n}{2}\right).$$

Finally,

$$\mathbf{E}_{\mu^{\otimes n}|\mathcal{N}_\sigma^{\otimes n}}[\hat{\mathbf{G}}_\sigma \text{SW}^p(\hat{\mu}_n, \mu)] \leq 2^{3p-1} (C_p + e^{5/3}) (M_{2p}(\mu) + 2^p \sigma^{4p})^{1/2} \left(4 \log\left(\frac{n}{2}\right) + \frac{\log n}{n} \right).$$

which is the desired result.

A.3.5 Proof of Proposition 2

Using triangle inequality, we have

$$W_p(\hat{\mu}_n, \hat{\nu}_n) \leq W_p(\hat{\mu}_n, \mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma) + W_p(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_\sigma) + W_p(\mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_\sigma, \hat{\nu}_n).$$

and then

$$W_p^p(\hat{\mu}_n, \hat{\nu}_n) \leq 3^{p-1} \{ W_p^p(\hat{\mu}_n, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) + W_p^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) + W_p^p(\mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma, \hat{\nu}_n) \}.$$

This implies that

$$\begin{aligned} & \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \mathbf{E}_{\nu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n)] \\ & \leq 3^{p-1} G_\sigma \text{SW}^p(\mu, \nu) \\ & \quad + 3^{p-1} \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \mu)] + 3^{p-1} \mathbf{E}_{\nu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\nu}_n, \nu)]. \end{aligned}$$

From the proof of Proposition 1, we have

$$\begin{aligned} & \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \mu)] \\ & \leq 2^{3p-1} (C_p + e^{5/3}) (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \left(\frac{1}{n} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}} + \frac{\log(n)}{n} \right). \end{aligned}$$

Similarly,

$$\begin{aligned} & \mathbf{E}_{\nu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\nu}_n, \nu)] \\ & \leq 2^{3p-1} (C_p + e^{5/3}) (M_{2p}(\nu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \left(\frac{1}{n} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}} + \frac{\log(n)}{n} \right). \end{aligned}$$

This gives that

$$\begin{aligned} & \mathbf{E}_{\mu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} \mathbf{E}_{\nu^{\otimes n} | \mathcal{N}_\sigma^{\otimes n}} [\hat{G}_\sigma \text{SW}^p(\hat{\mu}_n, \hat{\nu}_n)] \\ & \leq 3^{p-1} G_\sigma \text{SW}^p(\mu, \nu) \\ & \quad + 3^{p-1} 2^{3p} (C_p + e^{5/3}) \{ (M_{2p}(\mu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} + (M_{2p}(\nu) + M_{2p}(|\mathcal{N}_\sigma|))^{1/2} \} \\ & \quad \times \left(\frac{1}{n} \sum_{l=1}^n \frac{n+1}{\sqrt{l(n-l+1)}} + \frac{\log(n)}{n} \right). \end{aligned}$$

This ends the proof of the first statement in Proposition 2. For the second one, we also use a triangle inequality

$$W_p^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) \leq 3^{p-1} \{ W_p^p(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, \hat{\mu}_n) + W_p^p(\hat{\mu}_n, \hat{\nu}_n) + W_p^p(\hat{\nu}_n, \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma) \}.$$

Then we control each term as we did before.

A.4 Proof of Theorem 2: projection complexity

Using Holder's inequality, we have

$$\begin{aligned} \mathbf{E}_{\mathbf{u} \sim u_d} [|\widehat{G}_\sigma \text{SD}^p(\mu, \nu) - G_\sigma \text{SD}^p(\mu, \nu)|] & \leq \left(\mathbf{E}_{\mathbf{u} \sim u_d} [|\widehat{G}_\sigma \text{SD}^p(\mu, \nu) - G_\sigma \text{SD}^p(\mu, \nu)|^2] \right)^{1/2} \\ & = \left(\mathbf{V}_{\mathbf{u} \sim u_d} [|\widehat{G}_\sigma \text{SD}^p(\mu, \nu)|] \right)^{1/2} \\ & = \left(\mathbf{V}_{\mathbf{u} \sim u_d} [G_\sigma \text{SD}^p(\mu, \nu)] \right)^{1/2} \\ & = \frac{A(p, \sigma)}{\sqrt{L}}. \end{aligned}$$

A.5 Proof of Proposition 3

For all $\mathbf{u} \in \mathbb{S}^{d-1}$ we have $\mathcal{R}_{\mathbf{u}\mu}, \mathcal{R}_{\mathbf{u}\nu} \in \mathcal{P}(\mathbb{R})$. By application of the inequality of noise level satisfied by D in one dimension we get

$$D^p(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_{\sigma_2}, \mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_{\sigma_2}) \leq D^p(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_{\sigma_1}, \mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_{\sigma_1}).$$

Then, computing the expectation over the projections \mathbf{u} since the divergence is non-negative concludes the proof.

A.6 Relation between $G_\sigma \text{SW}^p(\mu, \nu)$ under two noise levels

We provide a relation between Gaussian-smoothed sliced Wasserstein distances under two noise levels.

Proposition 9. *Let $0 \leq \sigma_1 \leq \sigma_2$ be two noise levels. Then, one has*

$$G_{\sigma_1} \text{SW}^p(\mu, \nu) \leq 2^{p-1} G_{\sigma_2} \text{SW}^p(\mu, \nu) + 2^{\frac{5p}{2}} (\sigma_2^2 - \sigma_1^2)^p.$$

Proof of Proposition 9

The proof follows the same lines in proof of Lemma 1 in [Nietert et al. \(2021\)](#). First, we have that $\mathcal{N}_{\sigma_2} = \mathcal{N}_{\sigma_1} * \mathcal{N}_{\sqrt{\sigma_2^2 - \sigma_1^2}}$. Setting the following random variables: $X_{\mathbf{u}} \sim \mathcal{R}_{\mathbf{u}\mu}, Y_{\mathbf{u}} \sim \mathcal{R}_{\mathbf{u}\nu}, Z_X \sim \mathcal{N}_{\sigma_1}, Z_Y \sim \mathcal{N}_{\sigma_1}, Z'_X \sim \mathcal{N}_{\sqrt{\sigma_2^2 - \sigma_1^2}}, Z'_Y \sim \mathcal{N}_{\sqrt{\sigma_2^2 - \sigma_1^2}}$. The sliced Wasserstein distance $W_p^p(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_{\sigma_2}, \mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_{\sigma_2})$ is given as a minimization over couplings $(X_{\mathbf{u}}, Z_X, Z'_X)$ and $(Y_{\mathbf{u}}, Z_Y, Z'_Y)$. Using the inequality $\mathbf{E}[|X|^p] - 2^{p-1} \mathbf{E}[|Y|^p] \leq 2^{p-1} \mathbf{E}[|X + Y|^p]$ for any random variables $X, Y \in \mathbb{L}_p$ integrable, we obtain,

$$\begin{aligned} 2^{p-1} \mathbf{E}[|(X_{\mathbf{u}} + Z_X) - (Y_{\mathbf{u}} + Z_Y) + (Z'_X + Z'_Y)|^p] \\ \geq \mathbf{E}[|(X_{\mathbf{u}} + Z_X) - (Y_{\mathbf{u}} + Z_Y)|^p] - 2^{p-1} \mathbf{E}[|(Z'_X + Z'_Y)|^p]. \end{aligned}$$

Hence,

$$\begin{aligned} 2^{p-1} W_p^p(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_{\sigma_2}, \mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_{\sigma_2}) &\geq \inf \left(\mathbf{E}[|(X_{\mathbf{u}} + Z_X) - (Y_{\mathbf{u}} + Z_Y)|^p] - 2^{p-1} \mathbf{E}[|(Z'_X + Z'_Y)|^p] \right) \\ &\geq W_p^p(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_{\sigma_1}, \mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_{\sigma_1}) - 2^{p-1} \sup \mathbf{E}[|(Z'_X + Z'_Y)|^p] \\ &\geq W_p^p(\mathcal{R}_{\mathbf{u}\mu} * \mathcal{N}_{\sigma_1}, \mathcal{R}_{\mathbf{u}\nu} * \mathcal{N}_{\sigma_1}) - 2^{2p-1} \sup \mathbf{E}[|(Z'_X)|^p]. \end{aligned}$$

Therefore,

$$2^{p-1} G_{\sigma_2} \text{SW}^p(\mu, \nu) \geq G_{\sigma_1} \text{SW}^p(\mu, \nu) - 2^{2p-1} \sup \mathbf{E}[|(Z'_X)|^p].$$

Hence,

$$G_{\sigma_1} \text{SW}^p(\mu, \nu) \leq 2^{p-1} G_{\sigma_2} \text{SW}^p(\mu, \nu) + 2^{2p-1} \sup \mathbf{E}[|(Z'_X)|^p]$$

then

$$G_{\sigma_1} \text{SW}^p(\mu, \nu) \leq 2^{p-1} G_{\sigma_2} \text{SW}^p(\mu, \nu) + 2^{\frac{5p}{2}} (\sigma_2^2 - \sigma_1^2)^p,$$

and concludes the proof.

A.7 Proof of Proposition 5: continuity of the smoothed Gaussian sliced Wasserstein w.r.t. σ

From Lemma 1 in (Nietert et al., 2021), we know that the Gaussian-smoothed Wasserstein is continuous with respect to σ , for any distribution $\mathcal{R}_{\mathbf{u}}\nu$ and $\mathcal{R}_{\mathbf{u}}\mu$. In addition, for any \mathbf{u} , we have $W_p(\mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma, \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) \leq W_p(\mathcal{R}_{\mathbf{u}}\nu, \mathcal{R}_{\mathbf{u}}\mu)$. Then by applying Lebesgue's dominated convergence theorem (see Appendix 4) to the above inequality with $W_p(\mathcal{R}_{\mathbf{u}}\nu, \mathcal{R}_{\mathbf{u}}\mu)$ as a dominating function, that is u_d -almost everywhere integrable because both measures are in $\mathcal{P}_p(\mathbb{R}^d)$, we then conclude that the Gaussian-smoothed SWD is continuous w.r.t. σ .

A.8 Proof of Proposition 6: continuity of the smoothed sliced squared-MMD w.r.t. σ

Let us first recall the definition of the MMD divergence. Let $k : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a measurable bounded kernel on \mathbb{R} and consider the reproducing kernel Hilbert space (RKHS) \mathcal{H}_k associated with k and equipped with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$ and norm $\| \cdot \|_{\mathcal{H}_k}$. Let $\mathcal{P}_{\mathcal{H}_k}(\mathbb{R})$ be the set of probability measures η such that $\int_{\mathbb{R}} \sqrt{k(t, t)} d\eta(x) < \infty$. The kernel mean embedding is defined as $\Phi_k(\eta) = \int_{\mathbb{R}} k(\cdot, t) d\eta(t)$. The squared-maximum mean discrepancy between $\eta, \zeta \in \mathcal{P}(\mathbb{R})$ denoted as MMD : $\mathcal{P}_{\mathcal{H}_k}(\mathbb{R}) \times \mathcal{P}_{\mathcal{H}_k}(\mathbb{R}) \rightarrow \mathbb{R}_+$ is expressed as the distance between two such kernel mean embeddings. It is defined as Gretton et al. (2012)

$$\text{MMD}^2(\eta, \zeta) = \|\Phi_k(\eta) - \Phi_k(\zeta)\|_{\mathcal{H}_k}^2 = \mathbf{E}_{T, T' \sim \eta}[k(T, T')] - 2\mathbf{E}_{T \sim \eta, R \sim \zeta}[k(T, R)] + \mathbf{E}_{R, R' \sim \zeta}[k(R, R')]$$

where T and T' are independent random variables drawn according to η , R and R' are independent random variables drawn according to ζ , and T is independent of R . We define the Gaussian Smoothed Sliced squared-MMD as follows:

$$\begin{aligned} G_\sigma \text{MMD}^2(\mu, \nu) &= \int_{\mathbb{S}^{d-1}} \|\Phi_k(\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma) - \Phi_k(\mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma)\|_{\mathcal{H}_k}^2 u_d(\mathbf{u}) d\mathbf{u} \\ &= \int_{\mathbb{S}^{d-1}} (\mathbf{E}_{T, T' \sim \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma}[k(T, T')] - 2\mathbf{E}_{T \sim \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma, R \sim \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma}[k(T, R)] \\ &\quad + \mathbf{E}_{R, R' \sim \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma}[k(R, R')]) u_d(\mathbf{u}) d\mathbf{u}. \end{aligned}$$

From the definition of the smoothed sliced squared-MMD, we have

$$\begin{aligned} \mathbf{E}_{T, T' \sim \mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma}[k(T, T')] &= \iint_{\mathbb{R} \times \mathbb{R}} k(t, t') d\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma(t) d\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma(t') \\ &= \iint_{\mathbb{R} \times \mathbb{R}} \left(\int_{\mathbb{R}} k(t + z, t') d\mathcal{R}_{\mathbf{u}}\mu(z) \mathcal{N}_\sigma(t) \right) d\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma(t') \\ &= \iint_{\mathbb{R} \times \mathbb{R}} \left(\int_{\mathbb{R}^d} k(t + \mathbf{u}^\top x, t') d\mu(x) \mathcal{N}_\sigma(t) \right) d\mathcal{R}_{\mathbf{u}}\mu * \mathcal{N}_\sigma(t') \\ &= \iint_{\mathbb{R} \times \mathbb{R}} \iint_{\mathbb{R}^d \times \mathbb{R}^d} k(t + \mathbf{u}^\top x, t' + \mathbf{u}^\top x') d\mu(x) d\mu(x') d\mathcal{N}_\sigma(t) d\mathcal{N}_\sigma(t'). \end{aligned}$$

Similarly,

$$\mathbf{E}_{R, R' \sim \mathcal{R}_{\mathbf{u}}\nu * \mathcal{N}_\sigma}[k(R, R')] = \iint_{\mathbb{R} \times \mathbb{R}} \iint_{\mathbb{R}^d \times \mathbb{R}^d} k(r + \mathbf{u}^\top y, r' + \mathbf{u}^\top y') d\nu(y) d\nu(y') d\mathcal{N}_\sigma(r) d\mathcal{N}_\sigma(r')$$

and

$$\mathbf{E}_{T \sim \mathcal{R}_{\mathbf{u}} \mu * \mathcal{N}_\sigma, R \sim \mathcal{R}_{\mathbf{u}} \nu * \mathcal{N}_\sigma} [k(T, R)] = \iint_{\mathbb{R} \times \mathbb{R}} \iint_{\mathbb{R}^d \times \mathbb{R}^d} k(t + \mathbf{u}^\top x, r + \mathbf{u}^\top y) d\mu(x) d\nu(y) d\mathcal{N}_\sigma(t) d\mathcal{N}_\sigma(r).$$

Together the assumption of boundness of the kernel function k and the continuity of integrals, the three latter terms are continuous functions w.r.t. $\sigma \in (0, \infty)$. Again by the boundness of the kernel function k , there exists a positive finite constant C_k such that

$$|\mathbf{E}_{T, T' \sim \mathcal{R}_{\mathbf{u}} \mu * \mathcal{N}_\sigma} [k(T, T')] - 2\mathbf{E}_{T \sim \mathcal{R}_{\mathbf{u}} \mu * \mathcal{N}_\sigma, R \sim \mathcal{R}_{\mathbf{u}} \nu * \mathcal{N}_\sigma} [k(T, R)] + \mathbf{E}_{R, R' \sim \mathcal{R}_{\mathbf{u}} \nu * \mathcal{N}_\sigma} [k(R, R')]| \leq 4C_k.$$

We conclude the continuity of $\sigma \mapsto G_\sigma \text{MMD}^2(\mu, \nu)$ by an application of the continuity of integrals.

B Additional experiments

B.1 Sample complexity on CIFAR

We have also evaluated the sample complexity for the CIFAR dataset by sampling sets of increasing size. Results reported in Figure 6 confirms the findings obtained from the toy dataset.




Figure 6: Measuring the divergence between two sets of samples drawn iid from the CIFAR10 dataset. We compare three sliced divergences and their Gaussian smoothed versions with a $\sigma = 3$.

B.2 Identity of indiscernibles

The second experiment aims at checking whether our divergences converge towards a small value when the distributions to be compared are the same. For this, we consider samples from distributions μ and ν chosen as normal distributions with respectively mean $2 \times \mathbf{1}_d$ and $s\mathbf{1}_d$ with varying s (noted as the displacement). Results are depicted in Figure 7. We can see that all methods are able to attain their minimum when $s = 2$. Interestingly, the gap between the Gaussian smoothed and non-smoothed divergences for Wasserstein and Sinkhorn is almost indiscernible as the distance between distribution increases.




Figure 7: Measuring the divergence between two sets of samples in \mathbb{R}^{50} , one with mean 21_d and the other with mean $s1_d$ with increasing s . We compare three sliced divergences and their Gaussian smoothed version with a $\sigma = 3$.

C Useful results for the proofs

For the sake of completeness, we provide in this section useful results from the literature of functional analysis and probability measures that we frequently referred in the proofs of the main theoretical findings in this work.

Theorem 4 (Lebesgue’s Dominated Convergence Theorem, see [Bowers and Kalton \(2014\)](#)). *Let (E, \mathcal{T}, η) be a positive finite measure space and suppose $\{\psi_k\}_{k \in \mathbb{N}}$ be a sequence of scalar-valued measurable functions that converges η -almost everywhere (η -a.e.) to ψ . If there exists a function $g \in L_1(\eta)$ such that $|\psi_k| \leq g$ η -a.e. for all k , then ψ is integrable ($\psi \in L_1(\eta)$ and $\lim_{k \rightarrow \infty} \int_E \psi_k d\eta = \int_E \psi d\eta$).*

Theorem 5 (Bounded Convergence Theorem, see [Bowers and Kalton \(2014\)](#)). *Let (E, \mathcal{T}, η) be a positive finite measure space and suppose $\{\psi_k\}_{k \in \mathbb{N}}$ be a sequence of uniformly bounded scalar-valued measurable functions, i.e. there exists $K \geq 0$, $|\psi_k| \leq K$ for all k . If $\psi_k \rightarrow \psi$ η -a.e., then $f \in L_1(\eta)$ and $\lim_{k \rightarrow \infty} \int_E \psi_k d\eta = \int_E \psi d\eta$.*

Theorem 6 (Continuous Mapping Theorem, see [Bowers and Kalton \(2014\)](#)). *Let (E_1, d_1) and (E_2, d_2) be metric spaces and let $\psi : E_1 \rightarrow E_2$ measurable. Denote by \mathcal{D}_ψ the set of points of discontinuity of ψ . If $\eta, \eta_1, \eta_2, \dots$, are a finite measures on E_1 with $\eta(\mathcal{D}_\psi) = 0$ and $\eta_k \Rightarrow \eta$, then $\eta_k \circ \psi^{-1} \Rightarrow \eta \circ \psi^{-1}$.*

Theorem 7 (Levy’s Continuity Theorem, see [Athreya and Lahiri \(2006\)](#)). *Let $\eta, \{\eta_k\}_{k \in \mathbb{N}}$, be probability measures on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ with corresponding characteristic function φ and φ_k , $k \in \mathbb{N}$, respectively. Then the sequence $\{\eta_k\}_{k \in \mathbb{N}}$ converges weakly to η iff $\lim_{k \rightarrow \infty} \varphi_k(v) = \varphi(v)$ for all $v \in \mathbb{R}^d$.*

Theorem 8 (Fatou’s Lemma, see [Bowers and Kalton \(2014\)](#)). *Let (E, \mathcal{T}, η) be a positive finite measure space and suppose $\{\psi_k\}_{k \geq 1}$ be a sequence of scalar-valued measurable functions, then $\int_E (\liminf_{k \rightarrow \infty} |\psi_k|) d\eta \leq \liminf_{k \rightarrow \infty} \int_E |\psi_k| d\eta$.*

Theorem 9 (See [Billingsley \(1971\)](#)). *A necessary and sufficient condition for $\eta_k \Rightarrow \eta$ is that each that each subsequence $\{\eta_{s(k)}\}_k$ contains a further subsequence $\{\eta_{s(t(k))}\}_k$ converging weakly to η , where $s, t : \mathbb{N} \rightarrow \mathbb{N}$ are increasing applications.*

Definition 5 (Lévy-Prokhorov distance). *Let \mathcal{P} be the space of probability measures on a measurable metric space $((E, \rho), \mathcal{T})$. The Lévy-Prokhorov distance $\Lambda(\eta, \zeta)$ between $\eta, \zeta \in \mathcal{P}$ is defined as:*

$$\Lambda(\eta, \zeta) = \inf_{\varepsilon > 0} \{ \eta(A) < \zeta(A^\varepsilon) + \varepsilon, \quad \zeta(A) < \eta(A^\varepsilon) + \varepsilon, \quad \text{for all } A \in \mathcal{T} \},$$

where $A^\varepsilon = \{x \in E : \rho(x, A) < \varepsilon\}$.

Theorem 10 (see [Huber \(2011\)](#)). *Suppose E is a complete separable metric space. Then weak convergence is equivalent to Λ -convergence.*

Proposition 10 (See [Khoshnevisan \(2007\)](#)). *Suppose that (E, \mathcal{T}, η) is a measure space and $p \in [1, \infty)$. If $\{\psi_k\}_k$ is a sequence in $L^p(E)$ that converges in L^p to ψ , then there exists a subsequence $\{\psi_{s(k)}\}_k$ that converges pointwise η -a.e. to ψ .*

Proposition 11 (Cramér-Wold device, see [Huber \(2011\)](#)). *Let Z_0, Z_1, \dots be random variables with values in $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. Then $Z_n \Rightarrow Z_0$ iff $v^\top Z_n \Rightarrow v^\top Z_0$, for every $v \in \mathbb{R}^d$.*

Theorem 11 (See proof of Theorem 1 in [Fournier and Guillin \(2015\)](#)). *Let $\eta \in \mathcal{P}(\mathbb{R})$ and let $p \geq 1$. Assume that $M_q(\eta) < \infty$ for some $q > p$. There exists a constant $C_{p,q}$ depending only on p, q such that, for all $n \geq 1$,*

$$\mathbf{E}[W_p^p(\hat{\eta}_n, \eta)] \leq C_{p,q} M_q(\eta)^{p/q} \Delta_n(p, q),$$

where

$$\Delta_n(p, q) = \begin{cases} n^{-1/2} \mathbf{1}_{q > 2p}, \\ n^{-1/2} \log(n) \mathbf{1}_{q=2p} \\ n^{-(q-p)/q} \mathbf{1}_{p < q < 2p}. \end{cases}.$$

Lemma 3 (See Lemma 4.2 in [Bobkov and Ledoux \(2019\)](#)). *Given two collections of real numbers X_1, \dots, X_n and Y_1, \dots, Y_n , let $\hat{\alpha}_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and $\hat{\beta}_n = \frac{1}{n} \sum_{i=1}^n \delta_{Y_i}$ be the corresponding empirical measures. Then, for any $p \geq 1$, $W_p^p(\hat{\alpha}_n, \hat{\beta}_n) = \frac{1}{n} \sum_{l=1}^n |X_{(l)} - Y_{(l)}|^p$.*

Theorem 12 (Order statistics representation of $W_p(\hat{\alpha}_n, \alpha)$, see Theorem 4.3 in [Bobkov and Ledoux \(2019\)](#)). *Given random vectors (X_1, \dots, X_n) and (Y_1, \dots, Y_n) in \mathbb{R}^n , let $\hat{\alpha}_n$ and $\hat{\beta}_n$ be the corresponding empirical measures. Then, for any $p \geq 1$, $\mathbf{E}[W_p^p(\hat{\alpha}_n, \hat{\beta}_n)] = \frac{1}{n} \sum_{l=1}^n \mathbf{E}[|X_{(l)} - Y_{(l)}|^p]$. Additionally, if (Y_1, \dots, Y_n) is an independent copy of (X_1, \dots, X_n) and $\alpha = \frac{1}{n} \sum_{i=1}^n \text{Law}(X_i)$ is the mean marginal distribution, then $\mathbf{E}[W_p^p(\hat{\alpha}_n, \alpha)] \leq \frac{2^p}{n} \sum_{l=1}^n \mathbf{E}[|X_{(l)} - \mathbf{E}[X_{(l)}]|^p]$.*

Theorem 13 (See Theorem 1 in [Gribkova \(2020\)](#)). *Let R_1, \dots, R_n be i.i.d. real valued random variables and let $R_{(1)} \leq R_{(2)} \leq \dots \leq R_{(n)}$ denote the corresponding order statistics. Let $r > 0, q > 0$ be arbitrary positive numbers. Put $\rho = \frac{r}{q}$. For all $n \geq 2\rho + 1$ and for all $\rho \leq l \leq n - \rho + 1$ the following inequality holds*

$$\mathbf{E}[|R_{(l)}|^r] \leq C(\rho) \left(\mathbf{E}[|R_1|^q] \frac{(n+1)^2}{l(n-l+1)} \right)^\rho,$$

where $C(\rho) = 2\sqrt{\rho} e^{\rho+7/6}$.

References

- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223, International Convention Centre, Sydney, Australia. PMLR.
- Athreya, K. and Lahiri, S. (2006). *Measure Theory and Probability Theory*. Springer Texts in Statistics. Springer.
- Billingsley, P. (1971). *Weak Convergence of Measures: Applications in Probability*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104).
- Bobkov, S. and Ledoux, M. (2019). *One-Dimensional Empirical Measures, Order Statistics, and Kantorovich Transport Distances*. Memoirs of the American Mathematical Society. American Mathematical Society.
- Bonneel, N. and Coeurjolly, D. (2019). Spot: Sliced partial optimal transport. 38(4).
- Bonneel, N., van de Panne, M., Paris, S., and Heidrich, W. (2011). Displacement interpolation using lagrangian mass transport. *ACM Trans. Graph.*, 30(6):158:1–158:12.
- Bonnotte, N. (2013). *Unidimensional and Evolution Methods for Optimal Transportation*. Theses, Université Paris Sud - Paris XI ; Scuola normale superiore (Pise, Italie).
- Bowers, A. and Kalton, N. (2014). *An Introductory Course in Functional Analysis*. Universitext. Springer New York.
- Courty, N., Flamary, R., Tuia, D., and Rakotomamonjy, A. (2016). Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 39(9):1853–1865.
- Cramér, H. and Wold, H. (1936). Some theorems on distribution functions. *Journal of the London Mathematical Society*, s1-11(4):290–294.
- Deshpande, I., Hu, Y.-T., Sun, R., Pyrros, A., Siddiqui, N., Koyejo, S., Zhao, Z., Forsyth, D., and Schwing, A. G. (2019). Max-sliced Wasserstein distance and its use for gans. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10640–10648.
- Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.
- Fournier, N. and Guillin, A. (2015). On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3):707–738.
- Ganin, Y. and Lempitsky, V. (2015). Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR.
- Genevay, A., Peyré, G., and Cuturi, M. (2018). Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR.

- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. (2012). A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773.
- Gribkova, N. (2020). Bounds for absolute moments of order statistics. In *Exploring Stochastic Laws*, pages 129–134. De Gruyter.
- Huber, P. J. (2011). *Robust Statistics*, pages 1248–1251. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Kantorovich, L. (1942). On the transfer of masses (in russian). *Doklady Akademii Nauk*, 2:227–229.
- Khoshnevisan, D. (2007). *Probability*. Graduate studies in mathematics. American Mathematical Soc.
- Kolouri, S., Nadjahi, K., Simsekli, U., Badeau, R., and Rohde, G. (2019a). Generalized sliced Wasserstein distances. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 32, pages 261–272. Curran Associates, Inc.
- Kolouri, S., Park, S. R., Thorpe, M., Slepcev, D., and Rohde, G. K. (2017). Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Processing Magazine*, 34(4):43–59.
- Kolouri, S., Pope, P. E., Martin, C. E., and Rohde, G. K. (2019b). Sliced Wasserstein auto-encoders. In *International Conference on Learning Representations*.
- Kolouri, S., Zou, Y., and Rohde, G. K. (2016). Sliced Wasserstein kernels for probability distributions. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5258–5267.
- Lee, C.-Y., Batra, T., Baig, M. H., and Ulbricht, D. (2019). Sliced wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10285–10295.
- Lin, T., Zheng, Z., Chen, E., Cuturi, M., and Jordan, M. (2021). On projection robust optimal transport: Sample complexity and model misspecification. In *International Conference on Artificial Intelligence and Statistics*, pages 262–270. PMLR.
- Long, M., Cao, Y., Wang, J., and Jordan, M. (2015). Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pages 97–105. PMLR.
- Monge, G. (1781). Mémoire sur la théotie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences*, pages 666–704.
- Nadjahi, K., Durmus, A., Chizat, L., Kolouri, S., Shahrampour, S., and Simsekli, U. (2020). Statistical and topological properties of sliced probability divergences. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Nguyen, K., Ho, N., Pham, T., and Bui, H. (2020). Distributional sliced-Wasserstein and applications to generative modeling. *arxiv preprint 2002.07367*.

- Nietert, S., Goldfeld, Z., and Kato, K. (2021). Smooth p -wasserstein distance: Structure, empirical approximation, and statistical applications. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8172–8183. PMLR.
- Panaretos, V. M. and Zemel, Y. (2019). Statistical aspects of wasserstein distances. *Annual Review of Statistics and Its Application*, 6(1):405–431.
- Peyré, G. and Cuturi, M. (2019). Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607.
- Rachev, S. and Rüschendorf, L. (1998). *Mass Transportation Problems: Volume I: Theory*. Mass Transportation Problems. Springer.
- Rakotomamonjy, A. and Ralaivola, L. (2021). Differentially private sliced wasserstein distance. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8810–8820. PMLR.
- Salimans, T., Zhang, H., A.Radford, and Metaxas, D. (2018). Improving GANs using optimal transport. In *International Conference on Learning Representations*.
- Solomon, J., de Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., and Guibas, L. (2015). Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. Graph.*, 34(4):66:1–66:11.
- Sutherland, D. J., Tung, H.-Y., Strathmann, H., De, S., Ramdas, A., Smola, A. J., and Gretton, A. (2017). Generative models and model criticism via optimized maximum mean discrepancy. In *ICLR (Poster)*.
- Villani, C. (2009). *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer Berlin Heidelberg.
- Wu, J., Huang, Z., Acharya, D., Li, W., Thoma, J., Paudel, D. P., and Gool, L. V. (2019). Sliced wasserstein generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3713–3722.