



HAL
open science

Sliced-Wasserstein Estimation with Spherical Harmonics as Control Variates

Rémi Leluc, Aymeric Dieuleveut, François Portier, Johan Segers, Aigerim
Zhuman

► **To cite this version:**

Rémi Leluc, Aymeric Dieuleveut, François Portier, Johan Segers, Aigerim Zhuman. Sliced-Wasserstein Estimation with Spherical Harmonics as Control Variates. 2024. hal-04438124

HAL Id: hal-04438124

<https://hal.science/hal-04438124v1>

Preprint submitted on 5 Feb 2024

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.

Sliced-Wasserstein Estimation with Spherical Harmonics as Control Variates

Rémi Leluc¹ Aymeric Dieuleveut^{1*} François Portier^{2*} Johan Segers^{3*} Aigerim Zhuman^{3*}

Abstract

The Sliced-Wasserstein (SW) distance between probability measures is defined as the average of the Wasserstein distances resulting for the associated one-dimensional projections. As a consequence, the SW distance can be written as an integral with respect to the uniform measure on the sphere and the Monte Carlo framework can be employed for calculating the SW distance. Spherical harmonics are polynomials on the sphere that form an orthonormal basis of the set of square-integrable functions on the sphere. Putting these two facts together, a new Monte Carlo method, hereby referred to as Spherical Harmonics Control Variates (SHCV), is proposed for approximating the SW distance using spherical harmonics as control variates. The resulting approach is shown to have good theoretical properties, e.g., a no-error property for Gaussian measures under a certain form of linear dependency between the variables. Moreover, an improved rate of convergence, compared to Monte Carlo, is established for general measures. The convergence analysis relies on the Lipschitz property associated to the SW integrand. Several numerical experiments demonstrate the superior performance of SHCV against state-of-the-art methods for SW distance computation.

1. Introduction

The Sliced-Wasserstein (SW) distance between two probability measures, introduced in [Rabin et al. \(2012\)](#), shares similar topological properties with the standard Wasserstein distance ([Bonnotte, 2013](#); [Bayraktar & Guo, 2021](#); [Nadjahi et al., 2020b](#)) while having better properties in terms of computational complexity, especially for measures defined on spaces with relatively large dimension ([Nadjahi et al., 2020b](#)). As such, it has emerged as a powerful framework

for analyzing and solving a wide range of machine learning problems including generative modeling ([Deshpande et al., 2018; 2019](#); [Liutkus et al., 2019](#)), autoencoders ([Kolouri et al., 2018](#)), Bayesian computation ([Nadjahi et al., 2020a](#)), and image processing ([Bonnel et al., 2015](#)).

Informally, for $p \geq 1$ the SW_p distance between probability distributions μ and ν on \mathbb{R}^d is defined as the integral value, over θ uniformly distributed on the unit sphere \mathbb{S}^{d-1} , of the Wasserstein distance W_p between the projections of distributions μ and ν on the direction of θ . It takes advantage of the fact that (a) the Wasserstein distance between one-dimensional measures has a simple analytical form, being the $L^p([0, 1])$ -distance between their quantile functions (approximated using empirical quantiles), and (b) the integral is approximated by Monte Carlo sampling (uniformly on the unit sphere). As a consequence, two types of error affect the accuracy of most algorithms estimating the SW distance: first, the statistical error, due to the approximation of the distributions μ and ν by their empirical counterparts, which scales as $m^{-1/2}$, with m the size of the samples observed from μ and ν ([Nietert et al., 2022](#)); second, the Monte Carlo integration error, which scales as $n^{-1/2}$, with n the number of particles, as demonstrated in [Nadjahi et al. \(2020b\)](#) and [Nietert et al. \(2022\)](#). Importantly, these results establish that the sample complexity to approximate the SW distance does not depend on the dimension. As pointed out in [Nadjahi et al. \(2020b\)](#), even though the SW sample complexity is reduced compared to Wasserstein, it requires an additional Monte Carlo step which might offset the previous benefits if the associated variance is large.

In this paper, we propose and study a new approach to improve SW distance computation by improving the Monte Carlo estimation (step (b) above). The approach follows from a well-known variance reduction principle introducing *control variates* (functions having known integrals) in the Monte Carlo estimate. Following recent papers ([Oates et al., 2017](#); [Portier & Segers, 2019](#); [Leluc et al., 2021](#); [South et al., 2023](#)), the use of many control variates can significantly reduce the variance of the Monte Carlo method when the integrand f is well-approximated in some L^2 -space by the control variates. In [Portier & Segers \(2019\)](#) for instance, the obtained convergence rate is $n^{-1/2}\sigma_{cv}$, where σ_{cv}^2 is the variance of the residual function resulting from approximating f in the control variate space. By allowing the

*Equal contribution ¹CMAP, École Polytechnique, Institut Polytechnique de Paris, France ²CREST, ENSAI, France ³ISBA/LIDAM, UCLouvain, Belgium. Correspondence to: Rémi Leluc <remi.leluc@gmail.com>.

number of control variates to grow with n , the residual standard deviation σ_{cv} can vanish asymptotically, leading to an improvement of the $n^{-1/2}$ rate of the Monte Carlo estimate.

As the integral of interest in the SW distance is over the sphere, we propose to use as control variates the well-known *spherical harmonics*, which are homogeneous orthogonal polynomials defined on the sphere. Our approach, referred to as Sliced Wasserstein estimates with Spherical Harmonics as Control Variates (SHCV), leverages the property of spherical harmonics to be an orthonormal basis of $L_2(\mathbb{S}^{d-1})$ to build a precise and flexible approximation of the integrand leading to a more accurate Monte Carlo estimate.

Related methods. Several methods have been introduced for computing the SW distance and improve upon the standard Monte Carlo estimate. [Nguyen & Ho \(2024\)](#) use a *single* quadratic control variate to reduce the variance. This approach significantly differs from SHCV, which relies on a potentially large number of control variates. Further, for functions defined on \mathbb{R}^d , [Leluc et al. \(2023\)](#) use the *leave-one-out* nearest neighbor regression estimate as a control variate. As soon as the integrand is Lipschitz, the rate of convergence is $n^{-1/2-1/d}$, which improves upon the $n^{-1/2}$ rate of the Monte Carlo estimate. Even though this technique has not been studied for other spaces than \mathbb{R}^d (such as the sphere), good practical performance has been reported in [Leluc et al. \(2023\)](#) for approximating the SW distance.

In an attempt to improve the Monte Carlo estimate, Quasi-Monte Carlo (QMC) and Randomized Quasi Monte Carlo (RQMC) methods have been proposed to approximate the SW distance ([Nguyen et al., 2024](#)) based on deterministic sequences that are well spread over the integration domain. Finally, we mention the orthogonal approach of [Nadjahi et al. \(2021\)](#) which is based on Gaussian approximation as the dimension d grows, thus targets large dimension, and only applies for $p = 2$.

Contributions. The main contributions of this paper are summarized as follows:

- We propose a novel variance-reduced SW estimate based on control variate with spherical harmonics.
- We derive some theoretical properties as well as the convergence rate in probability on the SHCV integration error, providing rigorous underpinnings for the proposed methodology. Our analysis reveals that SHCV can achieve faster convergence rates than traditional unbiased estimators of the SW distance.
- In numerical experiments, our approach outperforms not only existing Monte Carlo approaches based on different control variates but also some recently introduced quasi Monte Carlo (QMC) methods (see related methods described above), without prohibitive additional cost.

Outline. Section 2 introduces the mathematical background of (Sliced-)Wasserstein distances while Section 3 presents MC integration with control variates. In Section 4, we introduce our estimator, based on spherical harmonics. It is empirically evaluated in Section 5. Section 6 provides a theoretical analysis and Section 7 concludes the paper. The appendices in the supplement contain proofs, algorithms and additional numerical results.

Notation. For a Borel set $X \subseteq \mathbb{R}^d$, $\mathcal{P}(X)$ is the set of probability measures with support contained in X . Let $p \in [1, \infty)$ and $\mathcal{P}_p(X) = \{\mu \in \mathcal{P}(X) : \int_X \|x\|^p d\mu(x) < +\infty\}$ be the set of probability measures on X with finite moment of order p . The unit sphere on \mathbb{R}^d is $\mathbb{S}^{d-1} = \{\theta \in \mathbb{R}^d : \|\theta\| = 1\}$ with surface area $|\mathbb{S}^{d-1}| = (2\pi^{d/2})/\Gamma(d/2)$, where Γ is the Euler gamma function. $\langle \cdot, \cdot \rangle$ is the Euclidean inner product and $\mathcal{N}(m, \Sigma)$ is the d -Gaussian distribution with mean $m \in \mathbb{R}^d$ and positive definite $d \times d$ covariance matrix Σ . The Laplace operator is $\Delta = \partial_1^2 + \dots + \partial_d^2$ where ∂_i is the i -th partial derivative.

2. Optimal Transport Distances: Wasserstein and Sliced-Wasserstein

2.1. Wasserstein Distance

The Wasserstein distance $W_p(\mu, \nu)$ of order $p \in [1, \infty)$ between probability measures $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ is defined by

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

where $\Pi(\mu, \nu) \subset \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ denotes the set of couplings for (μ, ν) , i.e., probability measures whose marginals with respect to the first and second variables are μ and ν , respectively. While the Wasserstein distance enjoys attractive theoretical properties ([Villani et al., 2009](#), Chapter 6), it suffers from a high computational cost. When computing $W_p(\mu_m, \nu_m)$ for discrete distributions μ_m and ν_m supported on m points, the worst-case computational complexity scales as $\mathcal{O}(m^3 \log m)$ ([Peyré et al., 2019](#)). However, for *univariate* distributions $\mu, \nu \in \mathcal{P}_p(\mathbb{R})$, the Wasserstein distance can be expressed through the quantile functions F^{-1} (inverse c.d.f.) as (see [Peyré et al., 2019](#), Section 2)

$$W_p^p(\mu, \nu) = \int_0^1 |F_\mu^{-1}(t) - F_\nu^{-1}(t)|^p dt.$$

In particular, for discrete measures $\mu_m = m^{-1} \sum_{i=1}^m \delta_{x_i}$ and $\nu_m = m^{-1} \sum_{i=1}^m \delta_{y_i}$, with $x_i, y_i \in \mathbb{R}$, the Wasserstein distance can be computed by sorting the atoms, yielding

$$W_p^p(\mu_m, \nu_m) = \frac{1}{m} \sum_{i=1}^m |x_{(i)} - y_{(i)}|^p, \quad (1)$$

where $(x_{(i)})_{i=1}^m$ and $(y_{(i)})_{i=1}^m$ are the order statistics of $(x_i)_{i=1}^m$ and $(y_i)_{i=1}^m$ respectively. The number of operations induced by the sorting step is $\mathcal{O}(m \log m)$.

2.2. Sliced-Wasserstein Distances

The Sliced-Wasserstein distance (Rabin et al., 2012; Bonneel et al., 2015; Kolouri et al., 2019) alleviates the high computational of Wasserstein distance, by *slicing* distributions and by taking advantage of the fast computation of the Wasserstein distance between *univariate* distributions (see Equation (1)).

Definition 1 (Sliced-Wasserstein distances). For $\theta \in \mathbb{S}^{d-1}$, let $\theta^* : \mathbb{R}^d \rightarrow \mathbb{R}$ denote the linear functional $\theta^*(x) = \langle \theta, x \rangle$ for $x \in \mathbb{R}^d$. Let $P \in \mathcal{P}(\mathbb{S}^{d-1})$ be a probability distribution on the unit sphere. The SW distance of order $p \in [1, \infty)$ based on P is defined for $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ as

$$\text{SW}_p^p(\mu, \nu, P) = \int_{\mathbb{S}^{d-1}} W_p^p(\theta_\#^* \mu, \theta_\#^* \nu) dP(\theta), \quad (2)$$

where for any measure $\xi \in \mathcal{P}(\mathbb{R}^d)$, $\theta_\#^* \xi$ is the push-forward measure by function θ^* , i.e. the distribution of the projection $\theta^*(X)$ of a random vector $X \in \mathbb{R}^d$ with distribution ξ .

We introduce the integrand $f_{\mu, \nu}^{(p)} : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ defined by

$$\forall \theta \in \mathbb{S}^{d-1}, \quad f_{\mu, \nu}^{(p)}(\theta) = W_p^p(\theta_\#^* \mu, \theta_\#^* \nu). \quad (3)$$

In practice, the SW distance of Equation (2) is approximated using a standard Monte Carlo estimate by (i) sampling a number $n \geq 1$ of independent random directions $\theta_1, \dots, \theta_n$ from P on \mathbb{S}^{d-1} and (ii) by averaging the values $(f_{\mu, \nu}^{(p)}(\theta_i))_i$. This is described in Algorithm 1.

Computing the SW distance between discrete measures $\mu_m = m^{-1} \sum_{i=1}^m \delta_{x_i}$ and $\nu_m = m^{-1} \sum_{i=1}^m \delta_{y_i}$ with atoms $x_i, y_i \in \mathbb{R}^d$ amounts to projecting $(x_i)_{i=1}^m$ and $(y_i)_{i=1}^m$ along the n random directions $\theta_1, \dots, \theta_n$ followed by computing the one-dimensional Wasserstein distances using Equation (1). This scheme requires $\mathcal{O}(n(dm + m \log m))$ operations which is, in general, faster than computing $W_p^p(\mu_m, \nu_m)$, especially for large m .

Remark 1 (Gaussian case when $p = 2$). The Wasserstein distance between two Gaussians $\mu = \mathcal{N}(a, \mathbf{A})$ and $\nu = \mathcal{N}(b, \mathbf{B})$ is known in closed form (Dowson & Landau, 1982) as $W_2^2(\mu, \nu) = \|a - b\|_2^2 + \mathfrak{B}^2(\mathbf{A}, \mathbf{B})$ where \mathfrak{B} is the Bures distance (Bhatia et al., 2019) on positive definite matrices, $\mathfrak{B}^2(\mathbf{A}, \mathbf{B}) = \text{Tr}(\mathbf{A}) + \text{Tr}(\mathbf{B}) - 2 \text{Tr}[(\mathbf{A}^{1/2} \mathbf{B} \mathbf{A}^{1/2})^{1/2}]$. The Sliced-Wasserstein distance between μ and ν relies on the projected distributions $\theta_\#^* \mu \sim \mathcal{N}(\theta^\top a, \theta^\top \mathbf{A} \theta)$ and $\theta_\#^* \nu \sim \mathcal{N}(\theta^\top b, \theta^\top \mathbf{B} \theta)$ via the integration over $\theta \in \mathbb{S}^{d-1}$ of $W_2^2(\theta_\#^* \mu, \theta_\#^* \nu) = |\theta^\top(a - b)|^2 + (\sqrt{\theta^\top \mathbf{A} \theta} - \sqrt{\theta^\top \mathbf{B} \theta})^2$.

3. Monte Carlo with Control Variates

This section presents the mathematical framework of Monte Carlo integration with control variates.

Algorithm 1 Sliced Wasserstein Monte Carlo

Require: $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, number of random projections n , probability $P \in \mathcal{P}(\mathbb{S}^{d-1})$

- 1: Sample random projections $\theta_1, \dots, \theta_n \sim P$
- 2: Compute $f_n = (f_{\mu, \nu}^{(p)}(\theta_i))_{i=1}^n$
- 3: Return average $\text{MC}_n = (\mathbb{1}_n/n)^\top f_n$

Monte Carlo integration. Consider a square-integrable, real-valued function $f \in L_2(P)$ on a probability space (Θ, \mathcal{F}, P) of which we would like to compute the integral

$$I(f) = \int_{\Theta} f(\theta) dP(\theta). \quad (4)$$

When $I(f)$ does not admit a closed form or when only calls from f are available, one may rely on Monte Carlo estimates. Let $\theta_1, \dots, \theta_n \sim P$ be an independent and identically distributed (i.i.d.) random sample from P . The naive Monte Carlo estimate $I_n(f)$ of $I(f)$ is given by the empirical mean

$$I_n(f) = \frac{1}{n} \sum_{i=1}^n f(\theta_i).$$

The Monte Carlo estimate is unbiased and has variance equal to $\sigma^2(f)/n$ where $\sigma^2(f) = \mathbb{I}\{f - I(f)\}^2$. Increasing the sample size n reduces the variance but at the cost of an increased computation time. Instead, it is possible to find another estimate with smaller variance.

Control Variates and OLSMC. A classical way to reduce the variance of the Monte Carlo estimate consists in incorporating knowledge about some specially chosen functions. Control variates are functions $\varphi_1, \dots, \varphi_s : \Theta \rightarrow \mathbb{R}$ with known expectations; without loss of generality, we assume that $\mathbb{I}(\varphi_j) = 0$ for $j = 1, \dots, s$. Let $\varphi = (\varphi_1, \dots, \varphi_s)^\top$ denote the \mathbb{R}^s -valued function with s control variates as elements. Since all control variates are centered, adding any linear combination of the form $\beta^\top \varphi$ with $\beta \in \mathbb{R}^s$ gives $\mathbb{I}(f - \beta^\top \varphi) = \mathbb{I}(f)$. Thus, given an i.i.d. sample $\theta_1, \dots, \theta_n$ from P , any coefficient vector $\beta \in \mathbb{R}^s$ yields an unbiased estimate of $I(f)$, given by

$$I_n(f, \beta) = \frac{1}{n} \sum_{i=1}^n (f(\theta_i) - \beta^\top \varphi(\theta_i)), \quad (5)$$

with variance equal to $\sigma^2(f - \beta^\top \varphi)/n$. The optimal coefficient β_* is defined as the one minimizing the variance,

$$\beta_* \in \arg \min_{\beta \in \mathbb{R}^s} \mathbb{I}\{f - I(f) - \beta^\top \varphi\}^2.$$

By the Hilbert projection theorem, $\beta_*^\top \varphi$ is the $L_2(P)$ -projection of $f - I(f)$ on the linear space of control variates $\text{Span}\{\varphi_1, \dots, \varphi_s\}$, as illustrated in Figure 1.




Figure 1. Visualization of the L_2 -projection of the integrand f onto the linear space $\text{Span}\{\varphi_1, \dots, \varphi_s\}$ of control variates.

Different approaches have been proposed to estimate β_* , each of which, when used in Equation (5), produces an enhanced Monte Carlo estimate with reduced variability (Glynn & Szechtman, 2002). Following Owen (2013, Section 8.9), or more recently Portier & Segers (2019) and Leluc et al. (2021), we frame the problem as an ordinary least-squares minimization problem. The integral $I(f)$ appears as the intercept of the linear regression model with explanatory variables $\varphi_1, \dots, \varphi_s$ and target response f , i.e.,

$$(I(f), \beta_*(f)) \in \arg \min_{(\alpha, \beta) \in \mathbb{R} \times \mathbb{R}^s} I[(f - \alpha - \beta^\top \varphi)^2].$$

In practice, these quantities are estimated using the particles $\theta_1, \dots, \theta_n$ sampled from P . The Ordinary Least Squares Monte Carlo (OLSMC) estimate $I_n^{\text{ols}}(f)$ is defined¹ by

$$(I_n^{\text{ols}}(f), \beta_n(f)) \in \arg \min_{(\alpha, \beta) \in \mathbb{R} \times \mathbb{R}^s} \|f_n - \alpha \mathbb{1}_n - \Phi \beta\|_2^2, \quad (6)$$

where $f_n = (f(\theta_1), \dots, f(\theta_n))^\top \in \mathbb{R}^n$ is the vector of evaluations of f , $\mathbb{1}_n = (1, \dots, 1)^\top \in \mathbb{R}^n$ and $\Phi \in \mathbb{R}^{n \times s}$ is the random matrix of control variates $\Phi = (\varphi(\theta_i)^\top)_{i=1}^n$.

Remark 2 (Complexity). The Monte Carlo estimate requires $\mathcal{O}(n\omega_f)$ computations where ω_f is the cost of evaluating the integrand f . The OLSMC estimate has a complexity of $\mathcal{O}(n\omega_f + \omega(\Phi))$ where $\omega(\Phi)$ is the complexity due to the additional step of fitting the optimal control variates. Typically, $\omega(\Phi) = \mathcal{O}(ns^2 + s^3)$ where $\mathcal{O}(ns^2)$ and $\mathcal{O}(s^3)$ operations are needed to compute and invert $\Phi^\top \Phi$, though it can be reduced by using approximate solution via optimization algorithms such as (stochastic) gradient descent algorithms (Zhang, 2004). When f is heavy to compute—which is the case for $f_{\mu_m, \nu_m}^{(p)}$ with discrete measures supported on a large number m of atoms—the extra term $\omega(\Phi)$ becomes negligible compared to $\mathcal{O}(n\omega_f)$.

Remark 3 (Regularization). As the number s of control variates may increase quickly with the dimension, Leluc et al. (2021) rely on the LASSO procedure (Tibshirani, 1996) in place of Equation (6) to achieve estimation and variable selection at the same time in the regime $n \ll s$.

¹It is well-defined if and only if $\mathbb{1}_n \notin \text{Span}(\Phi)$. If ill-defined, one can reduce the number of control variates such that $s < n$.

4. Sliced-Wasserstein Estimate with Spherical Harmonics as Control Variates

4.1. Spherical Harmonics

Originating from the solutions to the Laplace equation on the unit sphere, spherical harmonics (Atkinson & Han, 2012; Dai, 2013) constitute a mathematical framework with applications in computer graphics (Ramamoorthi & Hanrahan, 2001; Basri & Jacobs, 2003; Green, 2003) and machine learning (Cohen et al., 2018; Dutordoir et al., 2020), particularly in tasks involving spherical data representation.

Definition 2 (Polynomial spaces). Let \mathcal{P}_ℓ^d be the space of homogeneous polynomials of degree $\ell \geq 0$ on \mathbb{R}^d , i.e., $\mathcal{P}_\ell^d = \text{Span}\{x_1^{a_1} \cdots x_d^{a_d} \mid a_k \in \mathbb{N}, \sum_{k=1}^d a_k = \ell\}$. Let \mathcal{H}_ℓ^d be the space of real harmonic polynomials, homogeneous of degree ℓ on \mathbb{R}^d , i.e., $\mathcal{H}_\ell^d = \{Q \in \mathcal{P}_\ell^d \mid \Delta Q = 0\}$.

Definition 3 (Spherical harmonics). Spherical harmonics of degree (or level) $\ell \geq 0$ are defined as the restriction of elements in \mathcal{H}_ℓ^d to the unit sphere \mathbb{S}^{d-1} , i.e., restrictions to the sphere of harmonic homogeneous polynomials with d variables of degree ℓ .

Let N_ℓ^d denote the number of linearly independent spherical harmonics of degree $\ell \geq 0$ in dimension d , that is $N_\ell^d = \dim \mathcal{H}_\ell^d = \dim \mathcal{P}_\ell^d - \dim \mathcal{P}_{\ell-2}^d$ (Dai, 2013, Corollary 1.1.4) where it is agreed that $\dim \mathcal{P}_{\ell-2}^d = 0$ for $\ell = 0, 1$. We have

$$N_\ell^d = \frac{(2\ell + d - 2)(\ell + d - 3)!}{\ell!(d - 2)!}.$$

Note that, when restricted to the sphere, all polynomials are linear combinations of spherical harmonics. Moreover, spherical harmonics of different degrees are orthogonal with respect to $\langle f, g \rangle = \int_{\mathbb{S}^{d-1}} f(\theta)g(\theta) dP(\theta)$, where P is the uniform distribution on \mathbb{S}^{d-1} (Dai, 2013, Theorem 1.1.2). In practice, orthonormal bases can be constructed by applying the Gram–Schmidt process and an explicit basis of spherical harmonics can be written in terms of the Gegenbauer polynomials in spherical coordinates (Dai, 2013, Theorem 1.5.1). In the following, let $\varphi_{\ell,k}$ with $\ell \geq 0$ and $k = 1, \dots, N_\ell^d$ denote a set of spherical harmonics forming an orthonormal basis of \mathcal{H}_ℓ^d .

Proposition 1 (Hilbert basis). The spherical harmonics $\{\varphi_{\ell,k}, 0 \leq \ell, 1 \leq k \leq N_\ell^d\}$ form an orthonormal eigenbasis of the Hilbert space $L_2(\mathbb{S}^{d-1})$ so that for every $f \in L_2(\mathbb{S}^{d-1})$ we have

$$f = \sum_{\ell=0}^{\infty} \sum_{k=1}^{N_\ell^d} \hat{f}_{\ell,k} \varphi_{\ell,k} \quad \text{where} \quad \hat{f}_{\ell,k} = \int f \varphi_{\ell,k} dP.$$

Spherical harmonics thus constitute an eigenbasis of centered functions, and can therefore be used as control variates

with respect to the uniform distribution on the unit sphere.

For $\ell \geq 1$ and $k = 1, \dots, N_\ell^d$, we have

$$\mathbb{I}(\varphi_{\ell,k}) = \int_{\mathbb{S}^{d-1}} \varphi_{\ell,k}(\theta) d\mathbb{P}(\theta) = 0. \quad (7)$$

Orthonormality ensures that the empirical covariance matrix is well-conditioned when solving the OLS problem (6).

Remark 4 (Examples). In most applications, spherical harmonics in two and three variables are used. For $d = 2$, we have $\dim \mathcal{H}_\ell^2 = 2$ for all $\ell \geq 1$ and an orthogonal basis is given by the real and imaginary parts of $(x + iy)^\ell$. In polar coordinates $(x, y) = (r \cos(t), r \sin(t)) \in \mathbb{R}^2$, this basis is $\varphi_{\ell,1}(x, y) = r^\ell \cos(\ell t)$ and $\varphi_{\ell,2}(x, y) = r^\ell \sin(\ell t)$. By restriction on the circle \mathbb{S}^1 , spherical harmonics yield the classical Fourier expansions in cosine and sine functions and are commonly referred to as *circular harmonics*. For $d = 3$, we have $\dim \mathcal{H}_\ell^3 = 2\ell + 1$ and a classical orthogonal basis in spherical coordinates is written in terms of the associated Legendre polynomials (see Appendix D.1).

4.2. Spherical Harmonics Control Variates Estimate

With the idea of performing variance reduction when computing SW distances, we incorporate spherical harmonics as natural control variates on \mathbb{S}^{d-1} into the standard Monte Carlo estimate. We propose a novel procedure called *Spherical Harmonics Control Variates* (SHCV) estimate using the ordinary least squares formulation of Equation (6).

For $p \geq 1$, consider $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ and recall the integrand $f_{\mu,\nu}^{(p)} : \theta \mapsto \mathbb{W}_p^p(\theta_\#^* \mu, \theta_\#^* \nu)$ so that $\text{SW}_p^p(\mu, \nu) = \mathbb{I}(f_{\mu,\nu}^{(p)})$. Interestingly, this integrand is even so that for odd functions φ , the inner product vanishes, $\langle f_{\mu,\nu}^{(p)}, \varphi \rangle = 0$. The spherical harmonics expansion of $f_{\mu,\nu}^{(p)}$ (see Proposition 1) is only made of spherical harmonics with even degree. Thus, when finding the optimal control variates among spherical harmonics, we can omit the ones of odd degree as these are uncorrelated with our integrand.

Lemma 1 (Number of control variates). *In dimension d , the dimension of the linear space generated by the spherical harmonics of even degree up to maximal degree $2L$ is*

$$s_{L,d} = \sum_{\ell=1}^L N_{2\ell}^d = \binom{2L+d-1}{d-1} - 1.$$

For $d = 2$, the formula specializes to $s_{L,2} = 2L$ which is the number of circular harmonics of degree $2, 4, \dots, 2L$; For $d = 3$, the formula becomes $s_{L,3} = (2L+3)L$, which is indeed the sum of $(2\ell + 1)$ for $\ell = 2, 4, \dots, 2L$ (Remark 4).

Definition 4 (SHCV). *For $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, the SHCV estimate of maximum degree $2L$ based on particles $\theta_1, \dots, \theta_n \in \mathbb{S}^{d-1}$ and control variates $\varphi = (\varphi_j)_{j=1}^{s_{L,d}}$ is defined as the intercept of the OLS problem (6) with the*

Algorithm 2 Spherical Harmonics Control Variate Estimate

Require: $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, number of random projections n , $\mathbb{P} \in \mathcal{P}(\mathbb{S}^{d-1})$, spherical harmonics $\varphi = (\varphi_j)_{j=1}^{s_{L,d}}$

- 1: Sample random projections $\theta_1, \dots, \theta_n \sim \mathbb{P}$
- 2: Compute $f_n = (f_{\mu,\nu}^{(p)}(\theta_i))_{i=1}^n$, $\Phi = (\varphi(\theta_i)^\top)_{i=1}^n$
- 3: Solve $(\mathbb{I}_n^{\text{ols}}, \beta_n) \in \arg \min_{\alpha, \beta} \|f_n - \alpha \mathbb{1}_n - \Phi \beta\|_2^2$
- 4: Return $\text{SHCV}_n = \mathbb{I}_n^{\text{ols}}$

integrand $f_{\mu,\nu}^{(p)}$ as response and all spherical harmonics of even degree from 2 up to $2L$ as $n \times s_{L,d}$ covariate matrix $\Phi = (\varphi(\theta_i)^\top)_{i=1}^n$, i.e.,

$$\text{SHCV}_{n,L}^p(\mu, \nu) = \mathbb{I}_n^{\text{ols}}(f_{\mu,\nu}^{(p)}). \quad (8)$$

The procedure is described in Algorithm 2.

Remark 5 (Multiple integrals). Computational benefits occur when there are multiple integrands, since the SHCV estimate can be represented as a linear rule $w^\top f_n$ (Portier & Segers, 2019; Leluc et al., 2021), where the weight vector $w \in \mathbb{R}^n$ does not depend on the integrands. If K integrals need to be evaluated, the computation time of SHCV becomes $\mathcal{O}(Kn\omega_f + \omega(\Phi))$ compared to $\mathcal{O}(Kn\omega_f)$ for standard Monte Carlo, and the additional cost $\omega(\Phi)$ becomes negligible for large K . This is particularly relevant for SW-based kernels involving many pairwise computations $\text{SW}_2^2(\mu_i, \nu_j)$; see Algorithms 9 and 10 in Appendix B.4.

Remark 6 (Quadratic control variates). Any symmetric $d \times d$ matrix M yields a quadratic control variate $\varphi_M(\theta) = \theta^\top M \theta - \text{Tr}(M)/d$ as $\mathbb{I}(\varphi_M) = 0$. This is actually the approach used in Nguyen & Ho (2024) where the authors propose two different choices for M leading to two different estimators (see Appendix A). Note that, whenever $2L \geq 2$, these quadratic functions are included in our spherical harmonics space, making our method more general with potentially greater variance reduction (see Section 6).

5. Numerical Experiments

To assess the finite-sample performance of the SHCV estimate, we first present some synthetic data examples involving the computation of $\text{SW}_2^2(\mu, \nu)$ between multivariate Gaussian distributions. We consider different dimensions $d \in \{3; 5; 6; 10; 20\}$ and a varying number of random projections $n \in [10^2; 10^4]$. Then we focus on $\text{SW}_2^2(\mu_m, \nu_m)$ between 3D point clouds from the ShapeNetCore dataset of Chang et al. (2015). In both cases, the ground truth of the Sliced-Wasserstein distance is estimated with a very large number $n = 10^8$ of particles. The different results report the mean squared error (MSE) $\mathbb{E}[|\text{SW}_{2,n}^2(\mu, \nu) - \text{SW}_2^2(\mu, \nu)|^2]$ for the different estimates $\text{SW}_{2,n}^2(\mu, \nu)$ where the expectation is computed as an average over 100 independent runs. Finally, in the spirit of Kolouri et al. (2016) and Meunier

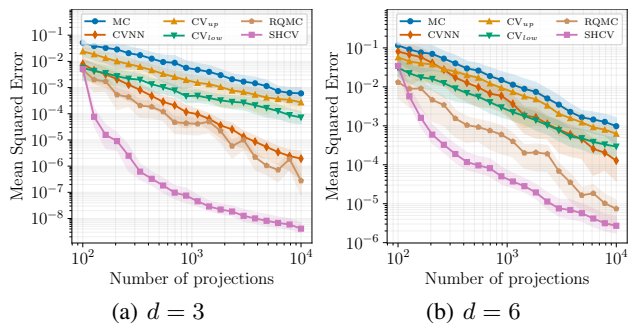


Figure 2. MSE curves for Gaussian distributions in dimension $d \in \{3; 6\}$ based on $n \in [10^2; 10^4]$ projections.

et al. (2022), we implement a Kernel Support Vector Machine (SVM) classifier for image classification.

For reproducibility, the code is available upon request and numerical details are available in the appendices. In particular, we rely on the efficient implementation of Dutordoir et al. (2020) to build the spherical harmonics.

Methods in competition. In the experiments, the SW distance is computed through different estimates. Along with the standard Monte Carlo estimate and the proposed SHCV estimate, we consider three other baselines based on control variates: the two estimates of Nguyen & Ho (2024) based on lower and upper bounds of a Gaussian approximation and the *control neighbors* estimate of Leluc et al. (2023) based on nearest neighbors estimates acting as control variates. For the sake of completeness, we also include a comparison with (Randomized) Quasi Monte Carlo procedures as in Nguyen et al. (2024) where deterministic QMC sets from $[0, 1]^d$ are mapped to the unit sphere and randomly rotated. The methods in competition are: the standard Monte Carlo estimate (MC), the lower-CV (CV_{low}) and upper-CV (CV_{up}) estimates, the control neighbors estimate (CVNN), the proposed estimate with spherical harmonics as control variates (SHCV), the Quasi Monte Carlo (QMC) and Random Quasi Monte Carlo estimates (RQMC).

5.1. Synthetic Data: Multivariate Gaussians

We compute $SW_2^2(\mu, \nu)$ between $\mu = \mathcal{N}(a, \mathbf{A})$ and $\nu = \mathcal{N}(b, \mathbf{B})$ with means $a, b \sim \mathcal{N}(\mathbf{1}_d, I_d)$ and covariance matrices $\mathbf{A} = \Sigma_a \Sigma_a^\top$ and $\mathbf{B} = \Sigma_b \Sigma_b^\top$ where all the entries of Σ_a, Σ_b are drawn according to $\mathcal{N}(0, 1)$. We consider here the *exact case* meaning that we directly evaluate the smooth integrand $f_{\mu, \nu}^{(2)}$ given in Remark 1. As the original lower-CV (CV_{low}) and upper-CV (CV_{up}) estimates are written to take discrete measures as inputs, we adapt them with the true means and covariance matrices. In Appendix B.2, we also consider the *sampled case* involving discrete measures μ_m, ν_m where we only have access to $m = 1000$ samples $x_1, \dots, x_m \sim \mu$ and $y_1, \dots, y_m \sim \nu$.

Method	$d = 5$		$d = 10$		$d = 20$	
	MSE	Time	MSE	Time	MSE	Time
MC	2.6e-2	06.5 ± 2.3	8.4e-3	06.7 ± 0.7	2.0e-2	06.5 ± 0.3
CV_{low}	4.7e-3	18.1 ± 4.5	3.0e-3	18.4 ± 4.3	4.4e-3	18.9 ± 3.3
CV_{up}	1.5e-2	20.1 ± 5.2	7.0e-3	16.8 ± 3.2	1.6e-2	14.6 ± 1.6
CVNN	7.7e-3	45.4 ± 12	1.0e-2	62.0 ± 14	2.6e-2	73.6 ± 12
QMC	4.9e-3	22.0 ± 7.1	4.1e-3	42.3 ± 9.7	1.3e-2	62.2 ± 6.1
RQMC	7.5e-4	23.6 ± 6.7	2.4e-3	52.4 ± 12	1.4e-2	84.3 ± 14
SHCV	4.8e-5	15.6 ± 4.5	1.7e-3	10.4 ± 4.9	2.7e-3	11.2 ± 5.2

Table 1. MSE and computing time (ms) for Gaussian distributions in dimension $d \in \{5; 10; 20\}$ based on $n = 500$ projections.

Figure 2 reports the MSE of the different methods in dimension $d \in \{3; 6\}$ with respect to the number of projections, while Figure 4 in Appendix B.1 does so with respect to computing time. Table 1 reports the MSE and computing times (in ms) in higher dimension $d \in \{5; 10; 20\}$. Figure 5 and Tables 4 and 5 in Appendix B.2 deal with the *sampled case* accordingly. In each case, the SHCV estimate gives the best MSE performance with respect to both the number of projections and computing time, beating other baselines by a factor up to 100.

5.2. Empirical Distributions of 3D Point Clouds

Similarly to Nguyen & Ho (2024), we use three point clouds taken at random from the dataset ShapeNetCore (Chang et al., 2015) corresponding to the objects *plane*, *lamp*, and *bed*, each composed of $m = 2048$ points in \mathbb{R}^3 . Each point cloud (x_1, \dots, x_m) is encoded as a discrete measure $\mu_m = m^{-1} \sum_{i=1}^m \delta_{x_i}$ with equal masses over the points in the cloud and the goal is to compute $SW(\mu_m, \nu_m)$ of order $p = 2$ between empirical distributions. Figure 3 reports the boxplots of the error distribution $\widehat{SW}_n(\mu_m, \nu_m) - SW(\mu_m, \nu_m)$ for the different SW estimates based on n random projections with $n \in \{100; 250; 500; 1000\}$ obtained over 100 independent runs. Once again, our SHCV estimate provides huge gains in terms of variance reduction and outperforms other baselines by a factor 100; see Table 6 in Appendix B.3.

5.3. Sliced-Wasserstein and Kernel SVM

SW distances are useful in the context of distribution regression (Szabó et al., 2016) where the goal is to learn a real-valued function defined on the space of probability distributions from a sequence of observations. Kolouri et al. (2016) first introduced a SW-based kernel for absolutely continuous distributions, which was later on extended to empirical distributions by Meunier et al. (2022). For any $\gamma > 0$, both $e^{-\gamma SW_2^2(\cdot, \cdot)}$ and $e^{-\gamma SW_1(\cdot, \cdot)}$ are valid kernels on probability measures. Using the SW kernel of Kolouri et al. (2016) defined by $k(\mu_i, \mu_j) = \exp(-\gamma SW_2^2(\mu_i, \mu_j))$ with $\gamma > 0$, we consider the image classification task on the *digits* dataset of Alpaydin & Kaynak (1998). This dataset contains $N = 1797$ images of size 8×8 and the goal is to




Figure 3. Boxplots of the error distribution $\widehat{SW}_n(\mu_m, \nu_m) - SW(\mu_m, \nu_m)$ for different SW estimates based on n random projections with $n \in \{100; 250; 500; 1000\}$ obtained over 100 independent runs.

predict the label $z \in \{0, 1, \dots, 9\}$ of the handwritten digit. Similarly to Meunier et al. (2022), we convert each image to a histogram so that the SW kernels have probability distributions as inputs. For image $i \in [N]$, let $(x_k^i, y_k^i)_{k=1}^{n_i}$ denote the coordinates of the n_i active pixels and $(\psi_k^i)_{k=1}^{n_i}$ the associated pixel intensities. The positions are renormalized to the grid $[0, 1]^2$ and we consider the weighted histogram $\mu_i = \sum_{k=1}^{n_i} \alpha_k^i \delta(x_k^i, y_k^i)$ where $\alpha_k^i = \psi_k^i / \sum_{l=1}^{n_i} \psi_l^i$ are the normalized intensities. We use a train/test split of size 80/20 giving $N_{\text{train}} = 1437$ training and $N_{\text{test}} = 360$ testing images, respectively, where each set is balanced.

We implement a Kernel SVM classifier where the SW_2 distance involved in the kernel expression is either computed with the standard Monte Carlo estimate (MC) or with the SHCV estimate (SHCV). The number of random projections n evolves in $\{25; 50; 75; 100\}$. To compare different scenarios, the bandwidth parameter $\gamma > 0$ of the Gaussian kernel is in $\{0.01; 1\}$ and the regularization parameter C of the SVM is optimized with a cross-validation strategy. Table 2 reports the means and standard deviations of the test accuracies of the different methods where the statistics are obtained over 10 independent runs.

n	$\gamma = 0.01$		$\gamma = 1$	
	MC	SHCV	MC	SHCV
25	86.0 \pm 0.6	87.3 \pm 0.3	92.7 \pm 0.6	93.3 \pm 0.6
50	86.7 \pm 0.5	87.6 \pm 0.3	93.2 \pm 0.4	94.0 \pm 0.3
75	86.8 \pm 0.2	87.8 \pm 0.2	93.3 \pm 0.4	94.3 \pm 0.2
100	86.9 \pm 0.3	87.7 \pm 0.1	93.3 \pm 0.3	94.3 \pm 0.2

Table 2. Kernel SVM: test accuracies with means and standard deviations. The statistics are obtained over 10 independent runs.

6. Theoretical Analysis

This section gathers several theoretical properties of the novel SHCV estimate. We first state some elementary properties of the underlying integration rule then provide a convergence rate in probability of the integration error.

6.1. Elementary Properties

By definition of the OLSMC estimate, the SHCV estimate is exact on the linear span of the control variates (Portier & Segers, 2019). Thus, if $f_{\mu, \nu}^{(p)}$ happens to be a polynomial function, then the SHCV estimate will be exact as soon as the maximal degree of the spherical harmonics is equal to or larger than the degree of $f_{\mu, \nu}^{(p)}$.

Proposition 2 (Exact rule). *If $f_{\mu, \nu}^{(p)}$ is a polynomial function with degree m , considering the SHCV estimate and control variates $\varphi = (\varphi_j)_{j=1}^{s_{L,d}}$, if $2L \geq m$ and $n > s_{L,d}$ then SHCV is exact, i.e., $SHCV_{n,L}^p(\mu, \nu) = SW_p^p(\mu, \nu)$.*

In the Gaussian case of Remark 1, for two proportional dispersion matrices \mathbf{A} and \mathbf{B} , say $\mathbf{B} = \gamma \mathbf{A}$ for some $\gamma > 0$, the formula of $f_{\mu, \nu}^{(2)}(\theta) = W_2^2(\theta_{\mu}^*, \theta_{\nu}^*)$ simplifies to a quadratic homogeneous polynomial. Therefore, the SHCV estimate is exact² as soon as the spherical harmonics of degree two are included as control variates, since the centered integrand is then contained in the linear space spanned by the control variates. Actually, this remains true in case $p = 2$ if the measures are related by an affine transformation.

Corollary 1 (Affine transform). *In case $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ are related by $X \sim \mu$ and $\alpha X + b \sim \nu$ where $\alpha \in (0, \infty)$ and $b \in \mathbb{R}^d$ then the SHCV estimate is exact.*

The SHCV estimator with spherical harmonics of maximal degree two or higher is guaranteed to have an asymptotic

²See numerical comparison in Appendix D.5 where the SHCV estimate is the only one achieving exact integration (zero error).

variance that is bounded above by the lower- and upper-CV methods of [Nguyen & Ho \(2024\)](#). Indeed, the latter two methods are control variate methods, each based on a *single* control variate, which happens to be a quadratic polynomial.

Corollary 2 (Mean invariance). *For $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, the error of the SHCV method is (exactly) invariant under changes of the mean vectors m_μ and m_ν of μ and ν respectively.*

This result implies that we can, without loss of generality, suppose that μ and ν are centered. The reason is that the effect of the two mean vectors being potentially different is the translation term $\langle \theta, m_Y - m_X \rangle$ coming from the optimal transport map that pushes $\theta_\#^* \mu$ to $\theta_\#^* \nu$, and this term produces an additional quadratic term $\langle \theta, m_Y - m_X \rangle^2$ in the function $f_{\mu, \nu}^{(2)}$. This additional quadratic polynomial in θ is integrated exactly by SHCV; its integral is $\|m_X - m_Y\|^2/d$. See [Remark 1](#) for the Gaussian case.

6.2. Asymptotic Error Bound

We provide a convergence rate in probability³ of the integration error $\text{SHCV}_{n,L}^p(\mu, \nu) - \text{SW}_p^p(\mu, \nu)$ for general $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ and $p \in [1, \infty)$, where $\text{SW}_p^p(\mu, \nu)$ is the SW distance with respect to the uniform distribution on the sphere and $\text{SHCV}_{n,L}^p(\mu, \nu)$ is our estimate in [Definition 4](#) based on n Monte Carlo particles on the sphere and using spherical harmonics of even degree up to $2L$ as control variates. The guarantee relies on two properties: first, a general convergence rate of the integration error of the OLSCV estimate in [Portier & Segers \(2019, Theorem 1\)](#), available in [Appendix C.2](#), and second, a bound on the uniform approximation error of a Lipschitz function on the sphere by its orthogonal projection on the space of polynomials ([Newman & Shapiro, 1964, Theorem 2](#)). The latter result applies since the integrand $f_{\mu, \nu}^{(p)}$ in [Equation \(3\)](#) is always at least Lipschitz: by [Han \(2023, Theorem 2.4\)](#), we have

$$\left| f_{\mu, \nu}^{(p)}(\theta) - f_{\mu, \nu}^{(p)}(\gamma) \right| \leq M_p(\mu, \nu) \|\theta - \gamma\| \quad (9)$$

for $\theta, \gamma \in \mathbb{S}^{d-1}$, where $M_p(\mu, \nu)$ is an explicit function of p and the p th order moments of μ and ν .

The modulus of continuity of $f \in \mathcal{C}(\mathbb{S}^{d-1})$ is $\omega(h) = \sup\{|f(\theta) - f(\gamma)| : \rho(\theta, \gamma) \leq h\}$ for $h > 0$, with $\rho(\theta, \gamma) = \arccos\langle \theta, \gamma \rangle$ the length of the arc connecting $\theta, \gamma \in \mathbb{S}^{d-1}$.

Theorem 1 ([Newman & Shapiro \(1964, Theorem 2\)](#)). *There exists an absolute constant $A > 0$ such that for each dimension $d \geq 2$, each function $f \in \mathcal{C}(\mathbb{S}^{d-1})$ and each degree L there exists a polynomial $P_{f,L}$ on \mathbb{S}^{d-1} of degree at most L such that: $\sup_{\theta \in \mathbb{S}^{d-1}} |f(\theta) - P_{f,L}(\theta)| \leq A \omega(d/L)$.*

The polynomial $P_{f,L}$ in [Theorem 1](#) that yields the uniform approximation is in general not the same as the or-

thogonal projection $\Pi_{f,L} = \sum_{\ell=0}^L \sum_{k=1}^{N_\ell^d} \hat{f}_{\ell,k} \varphi_{\ell,k}$ of f in $L_2(\mathbb{S}^{d-1})$ on the space of spherical harmonics of degree at most L , where we use the notation of [Proposition 1](#). Still, the $L_2(\mathbb{S}^{d-1})$ approximation error of $\Pi_{f,L}$ is bounded by the one of $P_{f,L}$, which is in turn bounded by the uniform approximation error of $P_{f,L}$. Applying this argument to the integrand $f_{\mu, \nu}^{(p)}$, we obtain the following convergence rate.

Theorem 2 (Convergence rate). *Let $d \geq 2$, $p \in [1, \infty)$ and $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ be fixed. For any degree sequence $L = L_n$ such that $L = o(n^{1/(2(d-1))})$ as $n \rightarrow \infty$, the integration error satisfies*

$$\left| \text{SHCV}_{n,L}^p(\mu, \nu) - \text{SW}_p^p(\mu, \nu) \right| = \mathcal{O}_{\mathbb{P}}(L^{-1} n^{-1/2}). \quad (10)$$

The proof of [Theorem 2](#) is given in [Appendix C.3](#). Writing $L = n^{1/(2(d-1))} / \ell_n$ where $\ell_n > 0$ diverges to infinity but can do so arbitrarily slowly (for instance logarithmically), the convergence rate in [Equation \(10\)](#) becomes

$$\left| \text{SHCV}_{n,L}^p(\mu, \nu) - \text{SW}_p^p(\mu, \nu) \right| = \mathcal{O}_{\mathbb{P}}(\ell_n n^{-\frac{1}{2} \frac{d}{d-1}}).$$

For $d = 3$, this yields the rate $n^{-3/4+o(1)}$ for the SHCV estimate, in comparison to the Monte Carlo rate $n^{-1/2}$.

If the integrand $f_{\mu, \nu}^{(p)}$ has higher-order derivatives, then a faster convergence rate can be expected because of approximation results in [Ragozin \(1971\)](#) improving [Theorem 1](#). This concerns for instance the SW distance between two Gaussian distributions, for which the integrand in [Remark 1](#) has derivatives of all orders. In contrast, for discrete measures, the integrand will not even be continuously differentiable, and the Lipschitz property is the best we can hope for. Both cases are illustrated in [Appendix D.4](#).

7. Conclusion

We have developed a novel method for reducing the variance of Monte Carlo estimation of the SW distance using spherical harmonics as control variates. The excellent practical performance of the SHCV estimate against state-of-the-art baselines is confirmed by theoretical properties and a convergence rate in probability for the integration error.

Using control variates with QMC sequences is usually *not* implemented in the same way as with Monte Carlo sequences ([Hickernell et al., 2005](#)) and would require a particular treatment which is beyond the scope of this paper. In statistical inference with parametric probability measures, note that SHCV is compatible with the computation of gradient $\nabla_\phi \text{SW}_p^p(\mu, \nu_\phi)$ and can be used for *generalized* SW flows ([Kolouri et al., 2019](#)). While the proposed SHCV estimate focuses on the uniform distribution on \mathbb{S}^{d-1} , it can be extended to more general probability distributions by combining control variates with importance sampling techniques as in [Leluc et al. \(2022\)](#).

³Notation $\mathcal{O}_{\mathbb{P}}(\cdot)$ means stochastic boundedness.

Acknowledgments

We thank Rémi Flamary from École Polytechnique for insightful discussion. The work of R. Leluc and A. Dieuleveut was supported by Hi!Paris (FLAG project), and ANR-REDEEM project. A. Zhuman gratefully acknowledges a research grant from the *National Bank of Belgium* and of the Research Council of the UCLouvain.

References

- Alpaydin, E. and Kaynak, C. Optical Recognition of Handwritten Digits. UCI Machine Learning Repository, 1998. DOI: <https://doi.org/10.24432/C50P49>.
- Atkinson, K. and Han, W. *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*. Lecture Notes in Mathematics. Springer, 2012.
- Basri, R. and Jacobs, D. W. Lambertian reflectance and linear subspaces. *IEEE transactions on pattern analysis and machine intelligence*, 25(2):218–233, 2003.
- Bayraktar, E. and Guo, G. Strong equivalence between metrics of Wasserstein type. *Electronic Communications in Probability*, 26(none):1 – 13, 2021. doi: 10.1214/21-ECP383. URL <https://doi.org/10.1214/21-ECP383>.
- Bhatia, R., Jain, T., and Lim, Y. On the Bures–Wasserstein distance between positive definite matrices. *Expositiones Mathematicae*, 37(2):165–191, 2019.
- Bonneel, N., Rabin, J., Peyré, G., and Pfister, H. Sliced and radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51:22–45, 2015.
- Bonnotte, N. *Unidimensional and evolution methods for optimal transportation*. PhD thesis, Université Paris Sud-Paris XI; Scuola normale superiore (Pise, Italie), 2013.
- Chang, A. X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Cohen, T. S., Geiger, M., Köhler, J., and Welling, M. Spherical CNNs. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Hkbd5xZRb>.
- Dai, F. *Approximation theory and harmonic analysis on spheres and balls*. Springer, 2013.
- Deshpande, I., Zhang, Z., and Schwing, A. G. Generative modeling using the sliced Wasserstein distance. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3483–3491, 2018.
- Deshpande, I., Hu, Y.-T., Sun, R., Pyrros, A., Siddiqui, N., Koyejo, S., Zhao, Z., Forsyth, D., and Schwing, A. G. Max-sliced Wasserstein distance and its use for gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10648–10656, 2019.
- Dowson, D. and Landau, B. The Fréchet distance between multivariate normal distributions. *Journal of multivariate analysis*, 12(3):450–455, 1982.
- Dutordoir, V., Durrande, N., and Hensman, J. Sparse Gaussian processes with spherical harmonic features. In *International Conference on Machine Learning*, pp. 2793–2802. PMLR, 2020.
- Feragen, A., Lauze, F., and Hauberg, S. Geodesic exponential kernels: When curvature and linearity conflict. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3032–3042, 2015.
- Glynn, P. W. and Szechtman, R. Some new perspectives on the method of control variates. In *Monte Carlo and Quasi-Monte Carlo Methods 2000: Proceedings of a Conference held at Hong Kong Baptist University, Hong Kong SAR, China, November 27–December 1, 2000*, pp. 27–49. Springer, 2002.
- Green, R. Spherical harmonic lighting: The gritty details. In *Archives of the game developers conference*, volume 56, pp. 4, 2003.
- Han, R. Sliced Wasserstein distance between probability measures on Hilbert spaces, 2023. arXiv preprint arXiv:2307.05802.
- Hickernell, F. J., Lemieux, C., and Owen, A. B. Control variates for Quasi-Monte Carlo. *Statistical Science*, 20(1):1–31, 2005.
- Kolouri, S., Zou, Y., and Rohde, G. K. Sliced Wasserstein kernels for probability distributions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5258–5267, 2016.
- Kolouri, S., Pope, P. E., Martin, C. E., and Rohde, G. K. Sliced Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018.
- Kolouri, S., Nadjahi, K., Simsekli, U., Badeau, R., and Rohde, G. Generalized sliced Wasserstein distances. *Advances in neural information processing systems*, 32, 2019.
- Leluc, R., Portier, F., and Segers, J. Control variate selection for Monte Carlo integration. *Statistics and Computing*, 31(4):50, 2021.

- Leluc, R., Portier, F., Segers, J., and Zhuman, A. A Quadrature Rule combining Control Variates and Adaptive Importance Sampling. In *Advances in Neural Information Processing Systems*, volume 35, pp. 11842–11853. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/4d4e8614a37f0aff841ba87ed1a898c1-Paper-Conference.pdf.
- Leluc, R., Portier, F., Segers, J., and Zhuman, A. Speeding up Monte Carlo integration: Control neighbors for optimal convergence. *arXiv preprint arXiv:2305.06151*, 2023.
- Liutkus, A., Simsekli, U., Majewski, S., Durmus, A., and Stöter, F.-R. Sliced-Wasserstein flows: Nonparametric generative modeling via optimal transport and diffusions. In *International Conference on Machine Learning*, pp. 4104–4113. PMLR, 2019.
- Meunier, D., Pontil, M., and Ciliberto, C. Distribution regression with sliced Wasserstein kernels. In *International Conference on Machine Learning*, pp. 15501–15523. PMLR, 2022.
- Mezzadri, F. How to generate random matrices from the classical compact groups. *Notices of the American Mathematical Society*, 54(5):592–604, 2007.
- Nadjahi, K., De Bortoli, V., Durmus, A., Badeau, R., and Şimşekli, U. Approximate Bayesian computation with the sliced-wasserstein distance. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5470–5474. IEEE, 2020a.
- Nadjahi, K., Durmus, A., Chizat, L., Kolouri, S., Shahrampour, S., and Simsekli, U. Statistical and topological properties of sliced probability divergences. *Advances in Neural Information Processing Systems*, 33:20802–20812, 2020b.
- Nadjahi, K., Durmus, A., Jacob, P. E., Badeau, R., and Simsekli, U. Fast approximation of the sliced-Wasserstein distance using concentration of random projections. *Advances in Neural Information Processing Systems*, 34:12411–12424, 2021.
- Newman, D. J. and Shapiro, H. S. Jackson’s Theorem in Higher Dimensions. In *Conference on Approximation Theory (Oberwolfach, 1963)*, pp. 208–219, 1964.
- Nguyen, K. and Ho, N. Sliced Wasserstein Estimation with Control Variates. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=StYc4hQAEi>.
- Nguyen, K., Bariletto, N., and Ho, N. Quasi-Monte Carlo for 3D Sliced Wasserstein. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=Wd47f7HEXg>.
- Nietert, S., Goldfeld, Z., Sadhu, R., and Kato, K. Statistical robustness, and computational guarantees for sliced Wasserstein distances. *Advances in Neural Information Processing Systems*, 35:28179–28193, 2022.
- Oates, C. J., Girolami, M., and Chopin, N. Control functionals for Monte Carlo integration. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 79(3):695–718, 2017.
- Owen, A. B. *Monte Carlo theory, methods and examples*. <https://artowen.su.domains/mc/>, 2013.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Peyré, G., Cuturi, M., et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Portier, F. and Segers, J. Monte Carlo integration with a growing number of control variates. *Journal of Applied Probability*, 56(4):1168–1186, 2019.
- Rabin, J., Peyré, G., Delon, J., and Bernot, M. Wasserstein barycenter and its application to texture mixing. In *Scale Space and Variational Methods in Computer Vision: Third International Conference, SSVM 2011, Ein-Gedi, Israel, May 29–June 2, 2011, Revised Selected Papers 3*, pp. 435–446. Springer, 2012.
- Ragozin, D. L. Constructive Polynomial Approximation on Spheres and Projective Spaces. *Transactions of the American Mathematical Society*, 162:157–170, 1971.
- Ramamoorthi, R. and Hanrahan, P. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pp. 497–500, 2001.
- South, L. F., Oates, C. J., Mira, A., and Drovandi, C. Regularized zero-variance control variates. *Bayesian Analysis*, 18(3):865–888, 2023.
- Szabó, Z., Sriperumbudur, B. K., Póczos, B., and Gretton, A. Learning theory for distribution regression. *The Journal of Machine Learning Research*, 17(1):5272–5311, 2016.

- Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.
- Villani, C. et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods*, 17(3):261–272, 2020.
- Zhang, T. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 116, 2004.

Supplementary Material: Sliced-Wasserstein Estimation with Spherical Harmonic as Control Variates

Appendix A gathers the pseudo-code of all the algorithms used in the experiments while Appendix B presents additional details and results for the numerical experiments. Appendix C gathers the proofs of the theoretical results and Appendix D is dedicated to additional results on spherical harmonics with (1) a case-study on the regularity of the integrand $f_{\mu,\nu}^{(p)}$ and (2) a numerical check of the exact integration rule.

A. Algorithms for Sliced-Wasserstein Estimates

This section presents algorithmic details of the different Sliced-Wasserstein estimates in the numerical experiments:

- Algorithm 3 corresponds to the standard Monte Carlo estimate where random projections $\theta_1, \dots, \theta_n \sim \mathcal{P} = \mathcal{U}(\mathbb{S}^{d-1})$ are used to compute the Wasserstein distance of projected distributions with $f_{\mu,\nu}^{(p)}(\theta) = \mathbb{W}_p^p(\theta_\#^* \mu, \theta_\#^* \nu)$.
- Algorithm 4 describes the proposed SHCV estimate which relies on the OLS procedure with spherical harmonics φ_j for $j = 1, \dots, s$ as control variates.
- Algorithm 5 and Algorithm 6 describe the two estimates of Nguyen & Ho (2024) which rely on Gaussian approximation with lower and upper control variates.
- Algorithm 7 presents the control neighbors technique of Leluc et al. (2023) and consists in building a control functional estimate with the 1-nearest neighbor estimate of the integrand f .
- Algorithm 8 implements a (Randomized) Quasi-Monte Carlo estimate. A low-discrepancy sequence on the unit cube $[0, 1]^d$ is mapped on \mathbb{S}^{d-1} to obtain a deterministic QMC point set on the unit sphere. Next, rotations from the orthogonal group are applied to randomize the samples.

Algorithm 3 Sliced Wasserstein Monte Carlo

- Require:** $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, number of random projections n
- 1: Sample random projections $\theta_1, \dots, \theta_n \sim \mathcal{U}(\mathbb{S}^{d-1})$
 - 2: Compute $f_n = (f_{\mu,\nu}^{(p)}(\theta_i))_{i=1}^n$
 - 3: Return average $\mathbf{MC}_n = (\mathbf{1}_n/n)^\top f_n$
-

Algorithm 4 Spherical Harmonics Control Variate Estimate

- Require:** $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, number of projections n , spherical harmonics $\varphi = (\varphi_j)_{j=1}^s$
- 1: Sample random projections $\theta_1, \dots, \theta_n \sim \mathcal{U}(\mathbb{S}^{d-1})$
 - 2: Compute $f_n = (f_{\mu,\nu}^{(p)}(\theta_i))_{i=1}^n$ and $\Phi = (\varphi(\theta_i)^\top)_{i=1}^n$
 - 3: $(\mathbf{I}_n^{\text{ols}}, \beta_n) \in \arg \min_{\alpha, \beta} \|f_n - \alpha \mathbf{1}_n - \Phi \beta\|_2^2$
 - 4: Return $\mathbf{SHCV}_n = \mathbf{I}_n^{\text{ols}}$
-

Algorithm 5 Lower-CV Sliced Wasserstein (Nguyen & Ho, 2024)

- Require:** $\mu = \sum_{i=1}^{m_X} \alpha_i \delta_{x_i}$ and $\nu = \sum_{j=1}^{m_Y} \beta_j \delta_{y_j}$, number of random projections n
- 1: Sample random projections $\theta_1, \dots, \theta_n \sim \mathcal{U}(\mathbb{S}^{d-1})$
 - 2: Compute $f_n = (f_{\mu,\nu}^{(p)}(\theta_i))_{i=1}^n$
 - 3: Compute $\mathbf{MC}_n = (\mathbf{1}_n/n)^\top f_n$
 - 4: Compute $\bar{x} = \sum_{i=1}^{m_X} \alpha_i x_i$, and $\bar{y} = \sum_{j=1}^{m_Y} \beta_j y_j$
 - 5: $C = (c_1, \dots, c_n)$ with $c_k = (\theta_k^\top (\bar{x} - \bar{y}))^2$
 - 6: Compute $b = \|\bar{x} - \bar{y}\|_2^2 / d$
 - 7: Compute $\gamma_{\text{low}} = \frac{\langle f_n - \mathbf{MC}_n, C - b \rangle}{\|C - b\|_2^2}$ and control variate $\mathcal{C} = (\mathbf{1}_n^\top (C - b)) / n$
 - 8: Return $\mathbf{CV}_{\text{low}} = \mathbf{MC}_n - \gamma_{\text{low}} \mathcal{C}$
-

Algorithm 6 Upper-CV Sliced-Wasserstein (Nguyen & Ho, 2024)

Require: $\mu = \sum_{i=1}^{m_X} \alpha_i \delta_{x_i}$ and $\nu = \sum_{j=1}^{m_Y} \beta_j \delta_{y_j}$, number of random projections n

- 1: Sample random projections $\theta_1, \dots, \theta_n \sim \mathcal{U}(\mathbb{S}^{d-1})$
- 2: Compute $f_n = (f_{\mu, \nu}^{(p)}(\theta_i))_{i=1}^n$
- 3: Compute $\text{MC}_n = (\mathbb{1}_n/n)^\top f_n$
- 4: Compute $\bar{x} = \sum_{i=1}^{m_X} \alpha_i x_i$, and $\bar{y} = \sum_{j=1}^{m_Y} \beta_j y_j$
- 5: $C = (c_1, \dots, c_n)$ with $c_k = (\theta_k^\top (\bar{x} - \bar{y}))^2 + \sum_{i=1}^{m_X} \alpha_i (\theta_k^\top (x_i - \bar{x}))^2 + \sum_{j=1}^{m_Y} \beta_j (\theta_k^\top (y_j - \bar{y}))^2$
- 6: Compute $b = \left(\|\bar{x} - \bar{y}\|_2^2 + \sum_{i=1}^{m_X} \alpha_i \|(x_i - \bar{x})\|_2^2 + \sum_{j=1}^{m_Y} \beta_j \|(y_j - \bar{y})\|_2^2 \right) / d$
- 7: Compute $\gamma_{\text{up}} = \frac{\langle f_n - \text{MC}_n, C - b \rangle}{\|C - b\|_2^2}$ and control variate $\mathcal{C} = (\mathbb{1}_n^\top (C - b)) / n$
- 8: Return $\text{CV}_{\text{up}} = \text{MC}_n - \gamma_{\text{up}} \mathcal{C}$

Algorithm 7 Control Neighbors Sliced Wasserstein (Leluc et al., 2023)

Require: $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, number of random projections n , probability $P \in \mathcal{P}(\mathbb{S}^{d-1})$

- 1: Sample random projections $\theta_1, \dots, \theta_n \sim P$
- 2: Compute $f_n = (f_{\mu, \nu}^{(p)}(\theta_i))_{i=1}^n$
- 3: Compute $\text{MC}_n = (\mathbb{1}_n/n)^\top f_n$
- 4: Compute nearest neighbor evaluations $\hat{f}_n = (\hat{f}_n^{(1)}(\theta_1), \dots, \hat{f}_n^{(n)}(\theta_n))$
- 5: Return $\text{CVNN}_n = \text{MC}_n - [\mathbb{1}_n^\top \{\hat{f}_n - \text{I}(\hat{f}_n)\}] / n$

Algorithm 8 Randomized Quasi-Monte Carlo Sliced Wasserstein (Nguyen et al., 2024)

Require: $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, number of projections n , Φ the cdf of $\mathcal{N}(0, 1)$

- 1: Generate low-discrepancy sequence U_1, \dots, U_n on $[0, 1]^d$
- 2: Compute Gaussian mapping $v_i = \Phi^{-1}(U_i) / \|\Phi^{-1}(U_i)\|_2$, where the standard normal quantile function Φ^{-1} is applied component-wise.
- 3: Generate a random rotation $R \sim \mathcal{U}(O_d(\mathbb{R}))$ where $O_d(\mathbb{R}) = \{R \in \mathbb{R}^{d \times d} \mid R^\top R = I_d\}$
- 4: Generate random projections $\theta_i = R v_i$
- 5: Compute $f_n = (f_{\mu, \nu}^{(p)}(\theta_i))_{i=1}^n$
- 6: Return $\text{RQMC}_n = (\mathbb{1}_n/n)^\top f_n$

Remarks and Practical Details on the Algorithms.

• In Algorithm 5, with the notation $z = \bar{x} - \bar{y}$, the control variate is the centered version of the function $\theta \mapsto |\theta^\top z|^2$, which has known integral $\int_{\mathbb{S}^{d-1}} |\theta^\top z|^2 dP(\theta) = \|z\|^2 / d = b$ (Step 6 in the algorithm) and where P denotes the uniform distribution on \mathbb{S}^{d-1} . The Lower-CV estimate is thus an OLSMC estimate with control variate $\varphi(\theta) = |\theta^\top z|^2 - \|z\|^2 / d$, a polynomial in θ of degree two.

• Algorithm 6 is actually an OLSMC estimate with a single control variate given by the centered version of $\theta \mapsto \theta^\top M \theta$ where $M = M_0 + M_X + M_Y$ with $M_0 = (\bar{x} - \bar{y})(\bar{x} - \bar{y})^\top$, $M_X = \sum_{i=1}^{m_X} \alpha_i (x_i - \bar{x})(x_i - \bar{x})^\top$ and similarly for M_Y . The integral with respect to the uniform distribution on the sphere is $\int_{\mathbb{S}^{d-1}} \theta^\top M \theta dP(\theta) = \text{Tr}(M) / d = b$ in Step 6 of the algorithm. The control variate is thus the centered quadratic polynomial $\varphi(\theta) = \theta^\top M \theta - \text{Tr}(M) / d$.

• In Step 4 of Algorithm 7, $\hat{f}_n^{(i)}(\theta_i)$ is a leave-one-out nearest-neighbor estimate computed with the Euclidean metric. In Step 5, computing $\text{I}(\hat{f}_n)$ requires sampling M additional particles θ_i from P and we set $M = n^{1+2/d}$ as recommended by the theory in Leluc et al. (2023).

• In Algorithm 8, random orthogonal matrices $R \in O_d(\mathbb{R})$ are generated using the function `ortho_group` of the Python package `scipy` (Virtanen et al., 2020). This function returns random orthogonal matrices drawn from the Haar distribution using a careful QR decomposition⁴ as in Mezzadri (2007).

⁴QR decomposition refers to the factorization of a matrix X into a product $X = QR$ of an orthonormal matrix Q and an upper triangular matrix R .

B. Numerical Details and Additional Results

To assess the finite-sample performance of the SHCV estimate, we first present some synthetic data examples involving the computation of $SW_2^2(\mu, \nu)$ between multivariate Gaussian distributions. We consider different dimensions $d \in \{3; 5; 6; 10; 20\}$ and a varying number of random projections $n \in [10^2; 10^4]$. We consider the two cases of *exact* distributions in Section B.1 and *sampled* distributions in Section B.2. Then in Section B.3, we focus on the Sliced-Wasserstein distance $SW_2(\mu_m, \nu_m)$ between 3D point clouds from the ShapeNetCore dataset of Chang et al. (2015). In both cases, the true value of the Sliced-Wasserstein distance is estimated with a very large number $n = 10^8$ of particles. The different results report the mean squared error $\mathbb{E}[|SW_{2,n}^2(\mu, \nu) - SW_2^2(\mu, \nu)|^2]$ for the different estimates $SW_{2,n}^2(\mu, \nu)$ where the expectation is computed as an average over 100 independent runs. The experiments were performed on a laptop Intel Core i7-10510U CPU 1.80GHz \times 8.

Methods in competition. Along with the standard Monte Carlo estimate (Algorithm 3) and the proposed SHCV estimate (Algorithm 4), we consider three other baselines based on control variates: the two estimates of Nguyen & Ho (2024) based on lower and upper bounds of a Gaussian approximation (Algorithm 5 and Algorithm 6) and the *control neighbors* estimate of Leluc et al. (2023) based on nearest neighbors estimates acting as control variates (Algorithm 7). For completeness, we also include a comparison with (Randomized) Quasi Monte Carlo procedures as in Nguyen et al. (2024) where deterministic QMC sets from $[0, 1]^d$ are mapped to the unit sphere and randomly rotated (Algorithm 8). The methods in competition are: the standard Monte Carlo estimate (MC), the lower-CV (CV_{low}) and upper-CV (CV_{up}) estimates, the control neighbors estimate (CVNN), the proposed estimate with spherical harmonics as control variates (SHCV), the Quasi Monte Carlo (QMC) and Random Quasi Monte Carlo estimates (RQMC).

B.1. Multivariate Gaussian: Exact Case with Distributions μ, ν and Smooth Integrands

We compute $SW_2^2(\mu, \nu)$ between $\mu = \mathcal{N}(a, \mathbf{A})$ and $\nu = \mathcal{N}(b, \mathbf{B})$ with means $a, b \sim \mathcal{N}(\mathbf{1}_d, I_d)$ and covariance matrices $\mathbf{A} = \Sigma_a \Sigma_a^\top$ and $\mathbf{B} = \Sigma_b \Sigma_b^\top$ where all the entries of Σ_a, Σ_b are drawn according to $\mathcal{N}(0, 1)$. The integrand of interest is $f_{\mu, \nu}^{(2)}(\theta) = |\theta^\top(a - b)|^2 + (\sqrt{\theta^\top \mathbf{A} \theta} - \sqrt{\theta^\top \mathbf{B} \theta})^2$. Figure 4 reports the MSE of the different methods with respect to both the number of projections and the computing time in dimension $d \in \{3; 6\}$, while Table 3 reports the MSE and computing times (in ms) in dimension $d \in \{5; 10; 20\}$.




Figure 4. MSE and computing time for Gaussian distributions, dimension $d \in \{3; 6\}$, obtained over 100 replications.

Method	$d = 5$		$d = 10$		$d = 20$	
	MSE	Time	MSE	Time	MSE	Time
MC	2.6e-2	06.5 \pm 2.3	8.4e-3	06.7 \pm 0.7	2.0e-2	06.5 \pm 0.3
CV _{low}	4.7e-3	18.1 \pm 4.5	3.0e-3	18.4 \pm 4.3	4.4e-3	18.9 \pm 3.3
CV _{up}	1.5e-2	20.1 \pm 5.2	7.0e-3	16.8 \pm 3.2	1.6e-2	14.6 \pm 1.6
CVNN	7.7e-3	45.4 \pm 12	1.0e-2	62.0 \pm 14	2.6e-2	73.6 \pm 12
QMC	4.9e-3	22.0 \pm 7.1	4.1e-3	42.3 \pm 9.7	1.3e-2	62.2 \pm 6.1
RQMC	7.5e-4	23.6 \pm 6.7	2.4e-3	52.4 \pm 12	1.4e-2	84.3 \pm 14
SHCV	4.8e-5	15.6 \pm 4.5	1.7e-3	10.4 \pm 4.9	2.7e-3	11.2 \pm 5.2

Table 3. Mean Squared Error (MSE) and computing time (ms) for Gaussian distributions in dimension $d \in \{5; 10; 20\}$ and $n = 500$.

B.2. Multivariate Gaussian: Sampled Case with Discrete Distributions μ_m, ν_m

We compute $SW_2^2(\mu_m, \nu_m)$ between empirical distributions $\mu_m = m^{-1} \sum_{i=1}^m \delta_{x_i}$ and $\nu_m = m^{-1} \sum_{j=1}^m \delta_{y_j}$,

$$x_1, \dots, x_m \sim \mu = \mathcal{N}(a, \mathbf{A}), \quad y_1, \dots, y_m \sim \nu = \mathcal{N}(b, \mathbf{B})$$

with means $a, b \sim \mathcal{N}_d(\mathbf{1}_d, I_d)$ and covariance matrices $\mathbf{A} = \Sigma_a \Sigma_a^\top$ and $\mathbf{B} = \Sigma_b \Sigma_b^\top$ where all the entries of Σ_a, Σ_b are drawn according to $\mathcal{N}(0, 1)$. This time, instead of evaluating the integrand $f_{\mu, \nu}^{(2)}$ directly, we assume that one only has access to data samples from the source and the target. We use $m = 1000$ data samples from the source $x_1, \dots, x_m \sim \mu$ and the target $y_1, \dots, y_m \sim \nu$. Figure 5 reports the mean squared errors of the different methods with respect to both the number of projections and the computing time in dimension $d \in \{3; 6\}$. Tables 4 and 5 report the MSE and computing times (in ms) in dimension $d \in \{5; 10; 20\}$.




Figure 5. MSE for sampled Gaussian distributions supported on $m = 1000$ points, dimension $d \in \{3; 6\}$, obtained over 100 replications.

Method	$d = 5$		$d = 10$		$d = 20$	
	$n = 500$	$n = 1000$	$n = 500$	$n = 1000$	$n = 500$	$n = 1000$
MC	1.45e-3	8.05e-4	9.45e-4	5.46e-4	1.47e-3	6.67e-4
CV _{low}	2.67e-4	1.46e-4	3.45e-4	1.63e-4	3.82e-4	1.73e-4
CV _{up}	8.44e-4	4.82e-4	7.51e-4	4.76e-4	1.09e-3	4.73e-4
CVNN	4.29e-4	1.79e-4	1.12e-3	4.77e-4	2.14e-3	1.44e-3
QMC	2.91e-4	7.97e-5	2.37e-4	8.35e-5	6.60e-4	1.42e-4
RQMC	5.80e-5	1.91e-5	2.75e-4	1.20e-4	1.17e-3	3.64e-4
SHCV	2.68e-6	5.63e-7	1.93e-4	6.86e-5	2.95e-4	1.16e-4

Table 4. Mean Squared Error (MSE) for sampled Gaussian distributions supported on $m = 1000$ points in dimension $d \in \{5; 10; 20\}$ with $n \in \{500; 1000\}$ particles, obtained over 100 independent runs.

Method	$d = 5$		$d = 10$		$d = 20$	
	$n = 500$	$n = 1000$	$n = 500$	$n = 1000$	$n = 500$	$n = 1000$
MC	81.1 \pm 3.5	152 \pm 0.9	80.7 \pm 4.4	152 \pm 1.2	81.1 \pm 1.8	155 \pm 2.3
CV _{low}	79.7 \pm 1.1	151 \pm 1.1	80.1 \pm 1.4	150 \pm 2.2	80.0 \pm 1.0	155 \pm 4.0
CV _{up}	83.0 \pm 1.2	157 \pm 1.7	83.0 \pm 1.7	157 \pm 1.6	83.1 \pm 1.5	160 \pm 3.0
CVNN	110 \pm 2.2	208 \pm 1.5	122 \pm 1.6	244 \pm 2.1	127 \pm 1.4	275 \pm 3.6
QMC	100 \pm 1.2	174 \pm 1.3	113 \pm 1.4	187 \pm 1.0	129 \pm 1.4	214 \pm 3.5
RQMC	96.3 \pm 2.2	190 \pm 1.4	113 \pm 1.2	188 \pm 2.4	130 \pm 1.0	214 \pm 3.8
SHCV	89.0 \pm 6.3	175 \pm 3.4	89.0 \pm 4.5	176 \pm 2.4	88.1 \pm 2.8	180 \pm 5.6

Table 5. Computing time (ms) for sampled Gaussian distributions supported on $m = 1000$ points in dimension $d \in \{5; 10; 20\}$ with $n \in \{500; 1000\}$ particles, obtained over 100 independent runs.

B.3. Empirical Distributions of 3D Point Clouds

We use three point clouds taken at random from the ShapeNetCore dataset of Chang et al. (2015) corresponding to the objects plane, lamp and bed, each composed of $m = 2048$ points of \mathbb{R}^3 (see Figure 6 below for visualization). The SHCV estimate uses spherical harmonics up to degree $2L = 16$ leading to $s = 152$ control variates. Figure 7 below reports the boxplots of the error distribution $(\widehat{SW}_n(\mu_m, \nu_m) - SW(\mu_m, \nu_m))$ for the different SW estimates based on n random projections with $n \in \{100; 250; 500; 1000\}$ obtained over 100 independent runs while Table 6 presents the mean squared errors $\mathbb{E}[|\widehat{SW}_n(\mu_m, \nu_m) - SW(\mu_m, \nu_m)|^2]$ for the different SW estimates with $n \in \{100; 500; 1000\}$. Once again, our SHCV estimate provides the best performance with huge gains in terms of variance reduction.




Figure 6. Visualization of 3D point clouds of plane, lamp, bed.




Figure 7. Boxplots of the error distribution $(\widehat{SW}_n(\mu_m, \nu_m) - SW(\mu_m, \nu_m))$ for different SW estimates based on n random projections with $n \in \{100; 250; 500; 1000\}$ obtained over 100 independent runs.

Method	PLANE/LAMP			LAMP/BED			PLANE/BED		
	$n = 100$	$n = 500$	$n = 1000$	$n = 100$	$n = 500$	$n = 1000$	$n = 100$	$n = 500$	$n = 1000$
MC	$5.1e-5$	$1.1e-5$	$6.0e-6$	$2.0e-4$	$3.8e-5$	$1.5e-5$	$1.3e-4$	$2.6e-5$	$1.2e-5$
CV_{low}	$4.8e-5$	$9.3e-6$	$5.1e-6$	$9.3e-5$	$1.4e-5$	$6.5e-6$	$4.4e-5$	$9.8e-6$	$5.2e-6$
CV_{up}	$1.1e-5$	$2.2e-6$	$1.3e-6$	$1.1e-5$	$2.3e-6$	$1.0e-6$	$6.6e-5$	$1.5e-5$	$8.7e-6$
CVNN	$1.3e-5$	$5.2e-7$	$1.8e-7$	$2.4e-5$	$1.1e-6$	$3.5e-7$	$3.2e-5$	$1.1e-6$	$3.3e-7$
QMC	$1.4e-5$	$6.9e-7$	$2.8e-8$	$2.4e-5$	$3.8e-7$	$2.9e-7$	$1.2e-4$	$1.2e-7$	$4.9e-7$
RQMC	$7.7e-6$	$2.2e-7$	$8.1e-8$	$1.5e-5$	$2.7e-7$	$1.3e-7$	$2.6e-5$	$1.0e-6$	$3.0e-7$
SHCV	$5.3e-6$	$2.1e-8$	$2.1e-8$	$2.5e-6$	$1.5e-8$	$1.4e-8$	$1.3e-5$	$3.7e-8$	$3.6e-8$

Table 6. Mean Squared Error $\mathbb{E}[|\widehat{SW}_n(\mu_m, \nu_m) - SW(\mu_m, \nu_m)|^2]$ for different SW estimates based on n random projections with $n \in \{100; 500; 1000\}$ obtained over 100 independent runs.

B.4. Kernel Support Vector Machines and Sliced-Wasserstein Distances

Kernel Support Vector Machines. Kernel Support Vector Machines (SVMs) are a powerful class of supervised learning algorithms used for classification and regression tasks. At their core, SVMs aim to find a hyperplane that best separates the data into different classes while maximizing the margin between these classes. However, in cases where the classes are not linearly separable, Kernel SVMs come into play. The key idea behind Kernel SVMs is to map the input data $x \in \mathbb{R}^q$ into a higher-dimensional space via $\phi(x)$, where a linear separation might be achievable. The kernel function $k(\cdot, \cdot)$ computes the similarity between pairs of data points in the higher-dimensional space without explicitly calculating the transformed data. Given training data $x_i \in \mathbb{R}^q$ with $i = 1, \dots, n$ divided into two classes and a vector of labels $y \in \{-1; +1\}$, the goal is to find the parameters $\omega \in \mathbb{R}^q$ and $b \in \mathbb{R}$ such that the prediction given by $\text{sign}(\omega^\top \phi(x_i) + b)$ is correct for most samples.

In Scikit-learn (Pedregosa et al., 2011), the ‘SVC class’ provides an implementation of Kernel SVMs. The objective function (also known as the loss function) for the SVM is derived from the problem of finding the optimal hyperplane that maximizes the margin between different classes. The formulation includes a regularization term C to balance the desire for a large margin with the goal of minimizing classification errors and the optimization problem of interest is given by

$$\min_{\omega, b, \zeta} \frac{1}{2} \omega^\top \omega + C \sum_{i=1}^n \zeta_i \quad \text{subject to } y_i (\omega^\top \phi(x_i) + b) \geq 1 - \zeta_i \text{ and } \zeta_i \geq 0, i = 1, \dots, n.$$

The parameter C , which is shared among all Support Vector Machine (SVM) kernels, balances the compromise between misclassifying training examples and maintaining a simple decision surface. A lower C value results in a smoother decision surface, while a higher C value strives to accurately classify all training examples. This problem is usually solved with the dual formulation through the Kernel matrix $K(x_i, x_j) = \langle \phi(x_i), \phi(x_j) \rangle$.

Sliced-Wasserstein Gaussian Kernels. SW distances inherit topological properties from the original Wasserstein distance such as the weak convergence of probability measures (Nadjahi et al., 2020b). SW distances are useful in the context of distribution regression (Szabó et al., 2016) where the goal is to learn a real-valued function defined on the space of probability distributions from a sequence of observations. The classical Wasserstein distance cannot be used to build a positive definite kernel on the space of distributions as the induced Wasserstein space has curvature (Feragen et al., 2015). To overcome this problem, a SW-based kernel for absolutely continuous distributions was first introduced in Kolouri et al. (2016) and later on extended to empirical distributions by Meunier et al. (2022). For any $\gamma > 0$, both $e^{-\gamma \text{SW}_2^2(\cdot, \cdot)}$ and $e^{-\gamma \text{SW}_1(\cdot, \cdot)}$ are valid kernels on probability measures.

Using the SW kernel of Kolouri et al. (2016) defined by $k(\mu_i, \mu_j) = \exp(-\gamma \text{SW}_2^2(\mu_i, \mu_j))$ with $\gamma > 0$, we consider the image classification task on the digits dataset of Alpaydin & Kaynak (1998). This dataset contains $N = 1797$ images of size 8×8 and the goal is to predict the label $z \in \{0, 1, \dots, 9\}$ of the handwritten digit. Similarly to Meunier et al. (2022), we convert each image to an histogram so that the SW kernels have probability distributions as inputs. We use a train/test split of size 80/20 giving $N_{\text{train}} = 1437$ training and $N_{\text{test}} = 360$ testing images respectively where each set is balanced. The Kernel SVM classification requires to compute the Gram matrices $K = (k(\mu_i, \mu_j))$ with $(i, j) \in [N_{\text{train}}] \times [N_{\text{train}}]$ and with $(i, j) \in [N_{\text{train}}] \times [N_{\text{test}}]$. Computational benefits occur with the SHCV estimate when many SW distances need to be computed. Using the particles $\theta_1, \dots, \theta_n \sim P$ and the spherical harmonics $\varphi = (\varphi_j)_{j=1}^s$ to compute $\Pi = \Phi(\Phi^\top \Phi)^{-1} \Phi^\top$ with $\Phi = (\varphi(\theta_i)^\top)_{i=1}^n$, we need to compute the weights $w = [(I - \Pi)\mathbb{1}_n] / [\mathbb{1}_n^\top (I - \Pi)\mathbb{1}_n]$ only once in order to obtain a linear integration rule $w_n^\top f_n^{(i,j)}$ with $f_n^{(i,j)} = (f_{\mu_i, \nu_j}^{(p)}(\theta_1), \dots, f_{\mu_i, \nu_j}^{(p)}(\theta_n))^\top$ to estimate $\text{SW}_p^2(\mu_i, \nu_j)$.

Algorithm 9 SW-Kernel Monte Carlo

Require: $\mathcal{D}_X = (\mu_1, \dots, \mu_{N_X}), \mathcal{D}_Y = (\nu_1, \dots, \nu_{N_Y})$,
 nb random projections $n \in \mathbb{N}, P \in \mathcal{P}(\mathbb{S}^{d-1})$

- 1: Initialize Kernel matrix $K \in \mathbb{R}^{N_X \times N_Y}$
- 2: Sample random projections $\theta_1, \dots, \theta_n \sim P$
- 3: **For** $i = 1, \dots, N_X$ **do**
- 4: **For** $j = 1, \dots, N_Y$ **do**
- 5: $f_n^{(i,j)} = (f_{\mu_i, \nu_j}^{(2)}(\theta_k))_{k=1}^n$
- 6: $\text{MC}_n = (\mathbb{1}_n/n)^\top f_n^{(i,j)}$
- 7: $K[i, j] = \exp(-\gamma \text{MC}_n)$
- 8: Return Kernel matrix K

Algorithm 10 SW-Kernel SHCV

Require: $\mathcal{D}_X = (\mu_1, \dots, \mu_{N_X}), \mathcal{D}_Y = (\nu_1, \dots, \nu_{N_Y})$,
 $n \in \mathbb{N}, P \in \mathcal{P}(\mathbb{S}^{d-1})$, control variates $\varphi = (\varphi_j)_{j=1}^s$

- 1: Sample random projections $\theta_1, \dots, \theta_n \sim P$
- 2: Compute $\Phi = (\varphi(\theta_i)^\top)_{i=1}^n$ and SHCV weights w_n
- 3: **For** $i = 1, \dots, N_X$ **do**
- 4: **For** $j = 1, \dots, N_Y$ **do**
- 5: $f_n^{(i,j)} = (f_{\mu_i, \nu_j}^{(2)}(\theta_k))_{k=1}^n$
- 6: $\text{SHCV}_n = w_n^\top f_n^{(i,j)}$
- 7: $K[i, j] = \exp(-\gamma \text{SHCV}_n)$
- 8: Return Kernel matrix K

C. Mathematical Proofs

C.1. Proofs of Lemma 1 and Corollary 1

Proof of Lemma 1. Let \mathcal{P}_ℓ^d be the space of homogeneous polynomials of degree ℓ . Then $\dim \mathcal{P}_\ell^d = \binom{\ell+d-1}{d-1}$, the number of multi-indices of the form $(\alpha_1, \dots, \alpha_d)$ for nonnegative integer α_j such that $\sum_{k=1}^d \alpha_k = \ell$. From this space, we need to remove the polynomials that are a multiple of $x_1^2 + \dots + x_d^2$, since, on the sphere, this factor trivial. Any such multiple is of the form $(x_1^2 + \dots + x_d^2)p(x)$ with $p \in \mathcal{P}_{\ell-2}^d$. For $\ell \geq 2$, we have (Atkinson & Han, 2012, Proposition 2.16) $N_\ell^d = \dim \mathcal{P}_\ell^d - \dim \mathcal{P}_{\ell-2}^d$. But then we have
$$\sum_{\ell=2,4,\dots,2k} N_\ell^d = \sum_{\ell=2,4,\dots,2k} \left[\binom{\ell+d-1}{d-1} - \binom{\ell-2+d-1}{d-1} \right] = \binom{2k+d-1}{d-1} - 1.$$

Proof of Corollary 1. For $\theta \in \mathbb{S}^{d-1}$, we have $\langle \theta, X \rangle \sim \theta_\#^* \mu$ and $\alpha \langle \theta, X \rangle + \langle \theta, b \rangle \sim \theta_\#^* \nu$. The optimal transport map between $\theta_\#^* \mu$ and $\theta_\#^* \nu$ is $x \mapsto \alpha x + \langle \theta, b \rangle$ and $f_{\mu,\nu}^{(2)}(\theta) = \mathbb{E}_\mu[\langle \theta, (1-\alpha)X - b \rangle^2]$, which is a quadratic polynomial in θ .

C.2. Asymptotic Bound for OLSMC

Consider the general set-up of Section 3, in particular the OLSMC estimate in Equation (6).

Theorem 3 (Portier & Segers (2019)). *Let (Θ, \mathcal{F}, P) be a probability space. Let $\varphi = (\varphi_1, \dots, \varphi_s)^\top \in (L_2(P))^s$ be a vector of orthonormal control variates, where $s = s_n$ depends on the Monte Carlo sample size n . Define $q_n(\theta) = \varphi(\theta)^\top \varphi(\theta)$ for $\theta \in \Theta$. If $\sup_{\theta \in \Theta} q(\theta) = o(n/s)$ as $n \rightarrow \infty$, then $I_n^{\text{ols}}(f) - I(f) = \mathcal{O}_\mathbb{P}(\sigma_n/\sqrt{n})$ as $n \rightarrow \infty$, where $\sigma_n^2 = I(\varepsilon_n^2)$, $\varepsilon_n = f - \hat{f}_s$, and $\hat{f}_s = \mathbb{E}[f] + \sum_{i=1}^s \langle f, \varphi_i \rangle \varphi_i$.*

C.3. Proof of Theorem 2

To apply Theorem 3, we consider $q_n(\theta) = \varphi(\theta)^\top \varphi(\theta)$ for $\theta \in \mathbb{S}^{d-1}$. Since $\varphi_1, \dots, \varphi_s$ are the orthonormal control variates,

$$q_n(\theta) = \sum_{i=1}^s \varphi_i^2(\theta) = \sum_{\ell=2,4,\dots,2L} \sum_{k=1}^{N_\ell^d} \varphi_{\ell,k}^2(\theta).$$

As a corollary to Theorem 2.9 in Atkinson & Han (2012), we have $\sum_{k=1}^{N_\ell^d} \varphi_{\ell,k}^2(\theta) = N_\ell^d / |\mathbb{S}^{d-1}|$. In view of Lemma 1,

$$q_n(\theta) = |\mathbb{S}^{d-1}|^{-1} \sum_{\ell=2,4,\dots,2L} N_\ell^d = |\mathbb{S}^{d-1}|^{-1} s_{L,d} \quad (11)$$

with $s = s_{L,d} = \binom{2L+d-1}{d-1} - 1$. Applying Stirling's formula gives $s_{L,d} = \mathcal{O}(L^{d-1})$ as $L \rightarrow \infty$. Hence the condition $\sup_{\theta \in \mathbb{S}^{d-1}} q(\theta) = o(n/s)$ in Theorem 3 becomes $s_{L,d}^2 = o(n)$, which is fulfilled as soon as

$$L = o\left(n^{1/(2(d-1))}\right), \quad n \rightarrow \infty. \quad (12)$$

Since the function $\theta \mapsto f(\theta) = f_{\mu,\nu}^{(p)}(\theta) = W_p^p(\theta_\#^* \mu, \theta_\#^* \nu)$ on \mathbb{S}^{d-1} is Lipschitz (Han, 2023, Theorem 2.4), the modulus of continuity of f evaluated at $h = d/(2L)$ satisfies

$$\omega(d/(2L)) \leq \text{Lip}_\rho(f) d/(2L) \leq \text{Lip}_{\|\cdot\|}(f) d/(2L), \quad (13)$$

where $\text{Lip}_\rho(f) = \sup_{\theta, \gamma \in \mathbb{S}^{d-1}} |f(\theta) - f(\gamma)|/\rho(\theta, \gamma)$ and $\text{Lip}_{\|\cdot\|}(f) = \sup_{\theta, \gamma \in \mathbb{S}^{d-1}} |f(\theta) - f(\gamma)|/\|\theta - \gamma\|$. More precisely, from Theorem 2.4 in Han (2023), we know that

$$\text{Lip}_{\|\cdot\|}(f) \leq M_p(\mu, \nu) = p2^{p-1} \max\{M_p(\mu), M_p(\nu)\}^{(p-1)/p} \left((M_p(\mu))^{1/p} + (M_p(\nu))^{1/p} \right) \quad (14)$$

with $M_p(\mu) = \int_{\mathbb{R}^d} \|x\|^p d\mu(x)$ and $M_p(\nu) = \int_{\mathbb{R}^d} \|x\|^p d\nu(x)$. Hence, by Theorem 1, the residual function $\varepsilon_n = f - \Pi_{f,L}$ with $\Pi_{f,L} = \sum_{\ell=0}^L \sum_{k=1}^{N_{2\ell}^d} \langle f, \varphi_{2\ell,k} \rangle \varphi_{2\ell,k}$ satisfies

$$\sigma_n = \|\varepsilon_n\|_2 \leq \|f - \Pi_{f,L}\|_2 \leq \|f - \Pi_{f,L}\|_\infty \leq \frac{AM_p(\mu, \nu)d}{2L}. \quad (15)$$

Therefore, by Theorem 3 along with (12) and (15), we have, since $I(f_{\mu,\nu}^{(p)}) = \text{SW}_p^p(\mu, \nu)$, the rate

$$\left| \text{SHCV}_{n,L}^p(\mu, \nu) - \text{SW}_p^p(\mu, \nu) \right| = \mathcal{O}_\mathbb{P}(L^{-1}n^{-1/2}).$$

D. Additional Results on Spherical Harmonics

D.1. Spherical Harmonics in Small Dimensions

Circular Harmonics on \mathbb{S}^1 . On the unit circle, the spherical harmonics are often called *circular harmonics* and correspond to the Fourier basis defined in circular coordinates $(x, y) = (\cos(t), \sin(t))$ by $\psi_0(\theta) = 1/\sqrt{2\pi}$ and for $k > 0$

$$\psi_{2k}(\theta) = \frac{\sin(kt)}{\sqrt{\pi}}, \quad \psi_{2k-1}(\theta) = \frac{\cos((k-1)t)}{\sqrt{\pi}}.$$

The circular harmonics are $\varphi_{1,0} = \psi_0$ and for $\ell \geq 1$, $\varphi_{\ell,1} = \psi_{2\ell-1}$ and $\varphi_{\ell,2} = \psi_{2\ell}$.

Spherical Harmonics on \mathbb{S}^2 . On the sphere \mathbb{S}^2 there exist $N_\ell^3 = 2\ell + 1$ linearly independent complex spherical harmonics $\psi_{\ell,k} : \mathbb{S}^2 \rightarrow \mathbb{C}$ of degree ℓ . Using the spherical coordinates $(x, y, z) = (\cos(u)\sin(v), \sin(u)\sin(v), \cos(v))$ with $0 \leq u < 2\pi$ and $0 \leq v < \pi$, the Laplace spherical harmonics are generally defined for all $-\ell \leq k \leq \ell$ by

$$\psi_{\ell,k}(u, v) = \sqrt{\frac{(2\ell+1)(\ell-k)!}{4\pi(\ell+k)!}} P_\ell^k(\cos(v)) e^{iku},$$

where P_ℓ^k is the associated Legendre polynomial defined by $P_\ell^k(x) = (-1)^k (1-x^2)^{k/2} (d^k/dx^k) P_\ell(x)$ with P_ℓ the Legendre polynomial of degree ℓ and $P_\ell^{-k}(x) = (-1)^k [(\ell-k)!/(\ell+k)!] P_\ell^k(x)$. This definition ensures orthonormality $\int_{v=0}^\pi \int_{u=0}^{2\pi} \psi_{\ell,k} \psi_{\ell',k'}^* d\Omega = \delta_{\ell,\ell'} \delta_{k,k'}$ with $d\Omega = \sin(v) du dv$. A real basis of spherical harmonics $\varphi_{\ell,k} : \mathbb{S}^2 \rightarrow \mathbb{R}$ can be defined in terms of their complex analogues by setting $\varphi_{\ell,0} = \psi_{\ell,0}$ and taking the real and imaginary parts as $\varphi_{\ell,k} = \sqrt{2}(-1)^k \Im(\psi_{\ell,|k|})$ if $k < 0$ and $\varphi_{\ell,k} = \sqrt{2}(-1)^k \Re(\psi_{\ell,k})$ if $k > 0$.

D.2. Number of Spherical Harmonics

Figure 8(a) presents the evolution of the number N_ℓ^d of spherical harmonics of degree ℓ in dimension d with $d \in [3, 10]$ and $\ell \in [1, 10]$ while Figures 8(b) and 8(c) report the evolution of the cumulative sum $\sum_{\ell=1}^L N_\ell^d$ and even cumulative sum $\sum_{\ell=1}^L N_{2\ell}^d$ for maximum degree $L \in [1, 10]$ and $2L \in \{2; 4; 6; 8; 10; 12\}$ respectively.

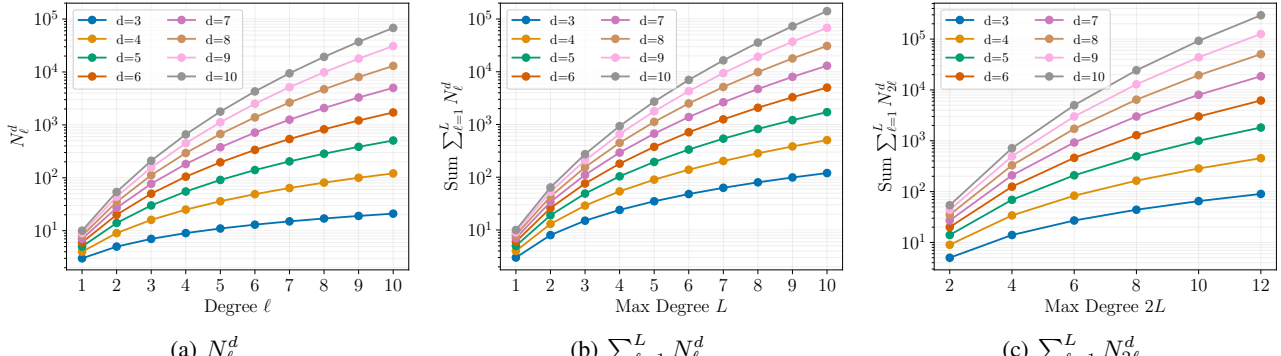


Figure 8. Evolution of the number N_ℓ^d of Spherical Harmonics and (even) cumulative sums according to the degree.

Table 7 reports the maximal degree of spherical harmonics and number of control variates for the considered dimensions in the experiments (See Lemma 1). For $d \in \{10; 20\}$ we only consider a subset of spherical harmonics of degree 4 to obtain a total number of $s = 300$ control variates.

dimension d	3	5	6	10	20
max degree $2L$	16	6	4	4	4
number s of control variates	152	209	125	714	8854

Table 7. Parameter setting of control variates.

D.3. Spherical Harmonics and Gegenbauer Polynomials

For $\alpha > 0$, Gegenbauer polynomials of degree ℓ , denoted by $C_\ell^\alpha : [-1, 1] \rightarrow \mathbb{R}$ are orthogonal polynomials with respect to the weight function $z \mapsto (1 - z^2)^{\alpha-1/2}$. They are special cases of Jacobi polynomials and can be written as

$$C_\ell^\alpha(z) = \sum_{k=0}^{\lfloor \ell/2 \rfloor} \frac{(-1)^k \Gamma(\ell - k + \alpha)}{\Gamma(\alpha) k! (\ell - 2k)!} (2z)^{\ell-2k}, \quad C_\ell^\alpha(1) = \frac{\Gamma(2\alpha + \ell)}{\Gamma(2\alpha) \ell!}.$$

In the following we fix ℓ , set $N = N_\ell^d$, and fix $\{\varphi_1, \dots, \varphi_N\}$ as an orthonormal basis for \mathcal{H}_ℓ^d . Denote by $Z_\ell(\cdot, \cdot)$ the reproducing kernel of \mathcal{H}_ℓ^d . In terms of orthonormal basis we have

$$\forall \theta, \eta \in \mathbb{S}^{d-1}, \quad Z_\ell(\theta, \eta) = \sum_{k=1}^N \varphi_k(\theta) \varphi_k(\eta).$$

Interestingly, the spherical harmonics are linked to the Gegenbauer polynomials through the following addition formula

Theorem 4. (Dai, 2013, Theorem 1.2.6) For $\ell \geq 0$ and $\theta, \eta \in \mathbb{S}^{d-1}$, $d \geq 3$

$$Z_\ell(\theta, \eta) = \sum_{k=1}^N \varphi_k(\theta) \varphi_k(\eta) = \frac{\ell + \alpha}{\alpha} C_\ell^\alpha(\langle \theta, \eta \rangle), \quad \alpha = \frac{d-2}{2}.$$

Corollary 3. (Dai, 2013, Corollary 1.2.8) For $\alpha = \frac{d-2}{2}$, the Gegenbauer polynomials C_ℓ^α satisfy the orthogonality relation

$$\frac{|\mathbb{S}^{d-2}|}{|\mathbb{S}^{d-1}|} \int_{-1}^1 C_\ell^\alpha(t) C_{\ell'}^\alpha(t) (1-t^2)^{\alpha-1/2} dt = h_\ell^\alpha \delta_{\ell, \ell'}, \quad h_\ell^\alpha = \frac{\alpha}{\ell + \alpha} C_\ell^\alpha(1).$$

Let $\{\theta_1, \dots, \theta_N\}$ be a collection of points on \mathbb{S}^{d-1} . The following remarks motivate the definition of a fundamental set. We let $M_1 := \varphi_1(\theta_1)$ and for $k = 2, 3, \dots, N$, define matrices

$$M_k := \begin{bmatrix} \varphi_1(\theta_1) & \dots & \varphi_1(\theta_k) \\ \vdots & \dots & \vdots \\ \varphi_k(\theta_1) & \dots & \varphi_k(\theta_k) \end{bmatrix}, \quad M_k(\theta) := \left[\begin{array}{c|c} & \varphi_1(\theta) \\ \hline & \vdots \\ M_{k-1} & \varphi_{k-1}(\theta) \\ \hline \varphi_k(\theta_1), \dots, \varphi_k(\theta_{k-1}) & \varphi_k(\theta) \end{array} \right]$$

The product of M_N and its transpose M_N^\top can be summed, on applying the addition formula, as $M_N^\top M_N = [Z_\ell(\theta_i, \theta_j)]_{i,j=1}^N$, which shows, in particular, $\det [Z_\ell(\theta_i, \theta_j)]_{i,j=1}^N = (\det M_N)^2 \geq 0$

Definition 5. (Dai, 2013, Definition 1.3.1) A collection of points $\{\theta_1, \dots, \theta_N\}$ in \mathbb{S}^{d-1} is called a fundamental system of degree ℓ on the sphere \mathbb{S}^{d-1} if

$$\det [C_\ell^\alpha(\langle \theta_i, \theta_j \rangle)]_{i,j=1}^N > 0, \quad \alpha = \frac{d-2}{2}$$

Theorem 5. (Dai, 2013, Theorem 1.3.3) If $\{\theta_1, \dots, \theta_N\}$ is a fundamental system of points on the sphere, then $\{C_\ell^\alpha(\langle \cdot, \theta_i \rangle) | i = 1, \dots, N\}$, $\alpha = (d-2)/2$ is a basis of $\mathcal{H}_\ell^d|_{\mathbb{S}^{d-1}}$

This theorem lies at the heart of the practical implementation used in Dutordoir et al. (2020) which constructs a fundamental system of vectors for efficient evaluation of spherical harmonics. The mathematical aspects of this process involve concepts such as the Gegenbauer polynomial, Cholesky decomposition, optimization, and normalization. The goal is to construct a set of vectors efficiently spanning the space required for evaluating spherical harmonics in the given dimension and degree. The procedure works as follows: Start by defining the Gegenbauer polynomial, denoted as $C_\ell^\alpha(\theta)$, of a specific degree ℓ and a dimension-dependent parameter $\alpha = (d-2)/2$. Select the first vector arbitrarily, often chosen as the north pole or the unit vector along the last dimension. Compute the Cholesky decomposition of the Gegenbauer-gram matrix associated with the current vectors and iteratively add new vectors to expand the fundamental system. For each new vector, optimize the determinant of the Gegenbauer-gram matrix, measuring the span of the space generated by the vectors. The optimization process involves finding a new vector that maximizes the determinant of the augmented matrix. This is typically done using optimization techniques such as BFGS. After finding the optimal vector, normalize it to ensure it has unit length.

D.4. Simple Case Study for Discrete Measures and Gaussian Measures

To investigate the approximation power of the integrand $\theta \mapsto f_{\mu,\nu}^{(p)}(\theta) = W_p^p(\theta_{\sharp}^* \mu_2, \theta_{\sharp}^* \nu_2)$ via spherical harmonics, we consider a case study in dimension $d = 2$ on the unit circle \mathbb{S}^1 with circular harmonics $(\varphi_{\ell,k})$ (see Appendix D.1). Using an Ordinary Least Squares regression, the polynomial approximation $\widehat{f}_{2L}^{(p)}$ of $f_{\mu,\nu}^{(p)}$ can be written as

$$\widehat{f}_{2L}(\theta) = \alpha + \sum_{\ell=2,4,\dots,2L} \left(\beta_{\ell}^{(1)} \varphi_{\ell,1}(\theta) + \beta_{\ell}^{(2)} \varphi_{\ell,2}(\theta) \right)$$

where α and $\beta = (\beta_2^{(1)}, \beta_2^{(2)}, \dots)$ are estimated by the OLS procedure. We investigate the smoothness properties of the integrand $f_{\mu,\nu}^{(p)}$ by comparing different measures μ and ν . We first consider the case of discrete measures μ_2 and ν_2 each supported on two atoms. In this simple case, the integrand $f_{\mu,\nu}^{(p)}(\theta)$ is only Lipschitz and we study the approximation error of our method. Then we consider the case of Gaussian measures $\mu = \mathcal{N}_2(\mathbf{m}_X, \Sigma_X)$ and $\nu = \mathcal{N}_2(\mathbf{m}_Y, \Sigma_Y)$. In that case, the integrand of interest is infinitely smooth.

Discrete measures. We consider the discrete measures μ_2 and ν_2 associated to the points $x_1 = (0, 0)$, $x_2 = (1, 0)$ and $y_1 = (0, 0)$, $y_2 = (1, 1)$ respectively. We have $\mu_2 = \frac{1}{2}\delta_{(0,0)} + \frac{1}{2}\delta_{(1,0)}$ and $\nu_2 = \frac{1}{2}\delta_{(0,0)} + \frac{1}{2}\delta_{(1,1)}$. The integrand $f_{\mu,\nu}^{(p)}$ can be computed exactly as

$$W_p^p(\theta_{\sharp}^* \mu_2, \theta_{\sharp}^* \nu_2) = \begin{cases} \frac{1}{2} |\sin(\theta)|^p, & \theta \in [0, \pi/2] \cup [3\pi/4, 3\pi/2] \cup [7\pi/4, 2\pi], \\ \frac{1}{2} (|\cos(\theta)|^p + |\cos(\theta) + \sin(\theta)|^p), & \theta \in [\pi/2, 3\pi/4] \cup [3\pi/2, 7\pi/4]. \end{cases}$$

By focusing on the case of order $p = 2$, we have

$$\begin{aligned} I_1 &= \int_0^{\pi/2} \sin^2(\theta) d\theta = \frac{\pi}{4}, & I_2 &= \int_{\pi/2}^{3\pi/4} \cos^2(\theta) + (\cos(\theta) + \sin(\theta))^2 d\theta = \frac{3\pi}{8} - \frac{3}{4}, \\ I_3 &= \int_{3\pi/4}^{7\pi/4} \sin^2(\theta) d\theta = \frac{3\pi}{8} - \frac{1}{4}, & I_4 &= \int_{3\pi/2}^{7\pi/4} \cos^2(\theta) + (\cos(\theta) + \sin(\theta))^2 d\theta = \frac{3\pi}{8} - \frac{3}{4}, \\ I_5 &= \int_{7\pi/4}^{2\pi} \sin^2(\theta) d\theta = \frac{\pi}{8} - \frac{1}{4}. \end{aligned}$$

Therefore,

$$SW_2^2(\mu_2, \nu_2) = \frac{1}{2\pi} \int_0^{2\pi} W_2^2(\theta_{\sharp}^* \mu_2, \theta_{\sharp}^* \nu_2) d\theta = \frac{1}{4\pi} (I_1 + I_2 + I_3 + I_4 + I_5) = \frac{3}{8} - \frac{1}{2\pi}.$$

Recall that $\text{SHCV}_{n,L}^p(\mu_2, \nu_2) = \mathbb{I}_n^{\text{ols}}$ where $(\mathbb{I}_n^{\text{ols}}, \beta_n) \in \arg \min_{\alpha, \beta} \|f_n - \alpha \mathbb{1}_n - \Phi \beta\|_2^2$.

Figure 9(a) reports the evolution of $W_2^2(\theta_{\sharp}^* \mu_2, \theta_{\sharp}^* \nu_2)$ (black line) and its polynomial approximation $\widehat{f}_{2L}(\theta)$ with different degrees of harmonics $2L \in \{2; 4; 8; 10\}$. Figure 9(b) displays the approximation error, i.e., the difference ε_n between $W_2^2(\theta_{\sharp}^* \mu_2, \theta_{\sharp}^* \nu_2)$ and its polynomial approximation $\widehat{f}_{2L}(\theta)$.

Gaussian measures. Let $\mu = \mathcal{N}_d(\mathbf{m}_X, \Sigma_X)$ and $\nu = \mathcal{N}_d(\mathbf{m}_Y, \Sigma_Y)$. Then, for $\theta_{\sharp}^* \mu = \mathcal{N}_1(\theta^{\top} \mathbf{m}_X, \theta^{\top} \Sigma_X \theta)$ and $\theta_{\sharp}^* \nu = \mathcal{N}_1(\theta^{\top} \mathbf{m}_Y, \theta^{\top} \Sigma_Y \theta)$, we have $W_2^2(\theta_{\sharp}^* \mu, \theta_{\sharp}^* \nu) = |\theta^{\top} (\mathbf{m}_X - \mathbf{m}_Y)|^2 + (\sqrt{\theta^{\top} \Sigma_X \theta} - \sqrt{\theta^{\top} \Sigma_Y \theta})^2$. We consider $\mathbf{m}_X = (0, 0)^{\top}$, $\mathbf{m}_Y = (1, 1)^{\top}$, $\Sigma_X = \begin{bmatrix} 1 & 0.2 \\ 0.2 & 1 \end{bmatrix}$, $\Sigma_Y = \begin{bmatrix} 10 & 3 \\ 3 & 10 \end{bmatrix}$

Figure 10(a) reports the evolution of $W_2^2(\theta_{\sharp}^* \mu_2, \theta_{\sharp}^* \nu_2)$ (black line) and its polynomial approximation $\widehat{f}_{2L}(\theta)$ with different degrees of harmonics $2L \in \{2; 4; 6\}$. Observe that in this case the integrand is very smooth so that the polynomial approximations seem to overlap. Figure 10(b) displays the approximation error, i.e., the difference ε_n between $W_2^2(\theta_{\sharp}^* \mu_2, \theta_{\sharp}^* \nu_2)$ and its polynomial approximation $\widehat{f}_{2L}(\theta)$.

Figure 11(a) presents the boxplots of the error distribution $(\text{SHCV}_{n,L}^2 - SW_2^2)$ both for discrete measures μ_2, ν_2 and for Gaussian measures with different degrees of harmonics $2L \in \{2; 4; 6; 8\}$ and Monte Carlo sample size $n = 10^4$. The Sliced-Wasserstein distance for Gaussian measures is obtained by the naive Monte Carlo method with $n = 10^8$. The boxplots are obtained over 100 replications.




Figure 9. $W_2^2(\theta_*^* \mu_2, \theta_*^* \nu_2)$, polynomial approximation $\hat{f}_{2L}(\theta)$ and errors $\varepsilon = W_2^2(\theta_*^* \mu_2, \theta_*^* \nu_2) - \hat{f}_{2L}(\theta)$ for two discrete measures.

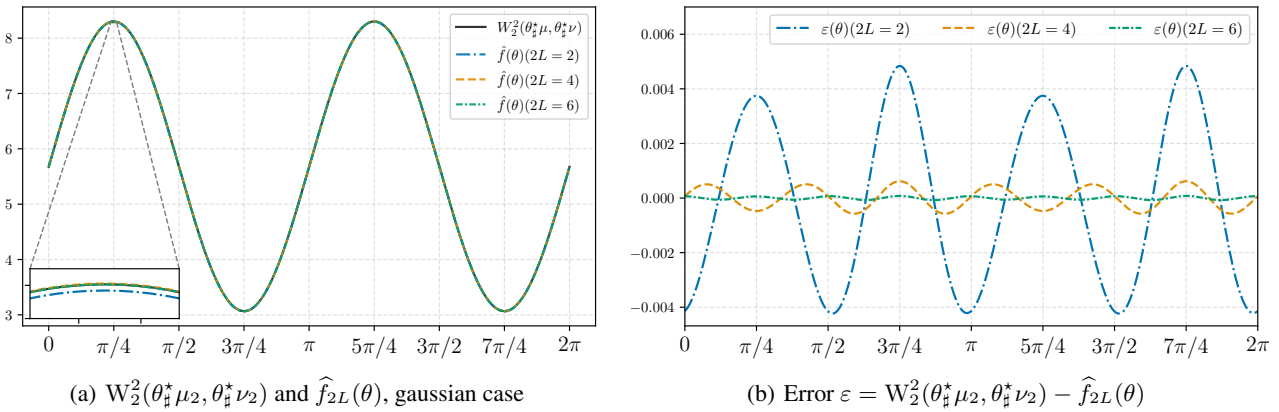


Figure 10. $W_2^2(\theta_*^* \mu_2, \theta_*^* \nu_2)$, polynomial approximation $\hat{f}_{2L}(\theta)$ and errors $\varepsilon = W_2^2(\theta_*^* \mu_2, \theta_*^* \nu_2) - \hat{f}_{2L}(\theta)$ for two Gaussian measures.




Figure 11. Error distribution of $(\text{SHCV}_{n,L}^2 - \text{SW}_2^2)$.

D.5. Numerical Check for Exact Integration Rule

We compute $SW_2^2(\mu, \nu)$ for $\mu = \mathcal{N}(a, \mathbf{A})$ and $\nu = \mathcal{N}(b, \mathbf{B})$ with means $a, b \sim \mathcal{N}(\mathbb{1}_d, I_d)$ and proportional (see Section 6.1) covariance matrices $\mathbf{A} = \Sigma_a \Sigma_a^\top$ and $\mathbf{B} = \gamma \mathbf{A}$ where all the entries of Σ_a are drawn according to $\mathcal{N}(0, 1)$. We consider dimensions $d \in \{3; 5\}$ and a number of projections $n \in [10^2; 10^4]$. Observe that in this case, the integrand of interest is a quadratic polynomial in θ given by

$$\begin{aligned} f_{\mu, \nu}^{(2)}(\theta) &= |\theta^\top (a - b)|^2 + (\sqrt{\theta^\top \mathbf{A} \theta} - \sqrt{\theta^\top \mathbf{B} \theta})^2 \\ &= |\theta^\top (a - b)|^2 + (1 - \sqrt{\gamma})^2 (\theta^\top \mathbf{A} \theta), \end{aligned}$$

and $SW_2^2(\mu, \nu)$ is given in close-form by integrating over $\theta \in \mathbb{S}^{d-1}$ such that

$$SW_2^2(\mu, \nu) = \frac{\|a - b\|_2^2}{d} + \frac{(1 - \sqrt{\gamma})^2}{d} \text{Tr}(\mathbf{A}).$$

Figure 12 reports the mean squared errors for the different estimates $\widehat{SW}_n(\mu, \nu)$ where the expectation is computed as an average over 100 independent runs with parameter $\gamma = 2$. Observe that, while the standard Monte Carlo estimate and other control variate-based methods achieve a small MSE in the range $[10^{-7}; 10^{-3}]$, our SHCV estimate obtains a MSE of 10^{-30} which is virtually zero up to machine precision. This numerical comparison empirically validates the exact integration rule of Proposition 2.




Figure 12. MSE for Gaussian distributions with proportional dispersion matrices, dimension $d \in \{3; 5\}$, obtained over 100 replications.