
Minimax-Optimal Two-Sample Test with Sliced Wasserstein

Binh Thuan Tran

LAMA, LIGM, Université Gustave Eiffel
binh-thuan.tran3@univ-eiffel.fr

Nicolas Schreuder

CNRS, LIGM
nicolas.schreuder@cnrs.fr

Abstract

We study the problem of nonparametric two-sample testing using the sliced Wasserstein (SW) distance. While prior theoretical and empirical work indicates that the SW distance offers a promising balance between strong statistical guarantees and computational efficiency, its theoretical foundations for hypothesis testing remain limited. We address this gap by proposing a permutation-based SW test and analyzing its performance. The test inherits finite-sample Type I error control from the permutation principle. Moreover, we establish non-asymptotic power bounds and show that the procedure achieves the minimax separation rate $n^{-1/2}$ over multinomial and bounded-support alternatives, matching the optimal guarantees of kernel-based tests while building on the geometric foundations of Wasserstein distances. Our analysis further quantifies the trade-off between the number of projections and statistical power. Finally, numerical experiments demonstrate that the test combines finite-sample validity with competitive power and scalability, and—unlike kernel-based tests, which require careful kernel tuning—it performs consistently well across all scenarios we consider.

1 INTRODUCTION

As machine-generated data becomes ubiquitous—from synthetic images to large language model outputs and scientific simulations—it is necessary to assess whether generated samples are statistically indistinguishable from real data samples. This question fits naturally into the framework of two-sample testing: one sample is drawn from a data distribution μ and the other from a generative model distribution ν . Formally, the goal

is to test

$$\mathcal{H}_0 : \mu = \nu \quad \text{against} \quad \mathcal{H}_1 : \mu \neq \nu.$$

Two-sample testing is a fundamental problem in statistics, with a long history and a wide range of practical applications (Lehmann and Romano, 2005). It has played a central role in areas such as clinical laboratory science (Miles et al., 2004; Zhang et al., 2021), finance (Horváth et al., 2014), bioinformatics (Borgwardt et al., 2006), neuroscience (Stelzer et al., 2013), and video content analysis (Liu et al., 2018). Recently, it has attracted growing interest in emerging machine learning and applied domains, including generative modeling (Li et al., 2017), model equality testing for large language models (Gao et al., 2024), and physics (Chakravarti et al., 2023).

Two primary classes of two-sample tests have been developed: *parametric* and *nonparametric*. Parametric tests, such as the classical t -test (Student, 1908) and Hotelling’s two-sample T^2 test (Hotelling, 1931), rely on strong distributional assumptions (e.g., Gaussianity) and may lose power when these assumptions fail. Nonparametric tests avoid such assumptions and are thus more broadly applicable.

Among nonparametric methods, kernel-based tests have become especially prominent. In particular, the kernel two-sample test of Gretton et al. (2012), based on the Maximum Mean Discrepancy (MMD)—a special case of Integral Probability Metrics (Müller, 1997)—is widely used for its flexibility and computational tractability. However, kernel-based tests depend critically on kernel choice and hyperparameter tuning, which ideally should reflect the geometry of the underlying data (Schrab et al., 2023). This limitation motivates exploring alternatives such as optimal transport (OT), which incorporates geometric structure through the choice of ground cost. The Wasserstein distance, in particular, is a natural candidate for two-sample testing, as it captures geometric discrepancies between distributions (Villani, 2008). However, the standard plug-in estimator suffers from the curse of dimensionality, with sharp error rates of order $n^{-1/d}$ (Fournier

and Guillin, 2015; Weed and Bach, 2019).

To mitigate this curse of dimensionality, we focus on the sliced Wasserstein distance (Rabin et al., 2011), which projects measures onto one-dimensional subspaces and attains a parametric $n^{-1/2}$ convergence rate while retaining geometric interpretability (Nad-jahi, 2021; Nietert et al., 2022). Despite its widespread use in generative modeling and learning, as well as promising empirical performance (Grossi et al., 2025), theoretical guarantees for its application in two-sample testing remain largely unexplored. In particular, finite-sample validity, non-asymptotic power guarantees, and minimax optimality results for tests based on the sliced Wasserstein distances are currently lacking in the literature. This gap limits the theoretical understanding and rigorous adoption of sliced Wasserstein-based two-sample tests in practice.

In this paper, we address these gaps by proposing a permutation-based two-sample test using the sliced Wasserstein distance. We prove that this test enjoys finite-sample Type I error control. Furthermore, we establish non-asymptotic lower bounds on the test’s power and demonstrate its minimax optimality against multinomial and bounded-support alternatives. We also analyze the computational–statistical trade-off inherent in the test, showing how the number of projections affects both runtime and statistical performance. A key technical contribution is our analysis of the SW distance between permuted distributions: to the best of our knowledge, this is the first work to control permutation quantiles for Wasserstein-based statistics. This analysis underpins the finite-sample guarantees for the SW test and extends naturally to the general Wasserstein distance. Finally, we complement our theoretical analysis with experiments, showing that our test achieves strong empirical performance and, unlike kernel-based methods whose power depends heavily on kernel choice, performs consistently well across diverse benchmarks without the need for parameter tuning or kernel aggregation.

The remainder of the paper is organized as follows. Section 2 reviews the necessary background on two-sample testing and sliced Wasserstein distance. Section 3 introduces the permutation-based sliced Wasserstein two-sample test. Section 4 presents our theoretical results on finite-sample validity, power analysis, and minimax optimality. Section 5 reports experimental evaluations on synthetic and real datasets. We conclude in Section 6 with a discussion of future research directions. All proofs are deferred to the Appendix.

2 BACKGROUND

In this section, we present the two-sample testing problem setup, define the sliced Wasserstein distance, and introduce minimax optimality for testing.

2.1 Problem setup

We consider two independent samples drawn from distributions μ and ν on $\mathbb{R}^d$:

$$\mathcal{Y}_n = (Y_1, \dots, Y_n) \stackrel{\text{i.i.d.}}{\sim} \mu, \quad \mathcal{Z}_m = (Z_1, \dots, Z_m) \stackrel{\text{i.i.d.}}{\sim} \nu,$$

both supported on a common ball $B_D = \{x \in \mathbb{R}^d : \|x\| \leq D\}$ for some $D > 0$, and assume $n \leq m$.

The goal is to test the hypotheses

$$H_0 : \mu = \nu \quad \text{versus} \quad H_1 : d(\mu, \nu) \geq \epsilon,$$

for some discrepancy metric d and fixed $\epsilon > 0$. Many existing tests fit this framework by choosing d as a specific probability discrepancy, such as the total variation distance or the MMD. Later, we will focus on the sliced Wasserstein distance.

A test is a measurable function

$$\phi_{n,m} : (\mathbb{R}^d)^n \times (\mathbb{R}^d)^m \rightarrow \{0, 1\},$$

where $\phi_{n,m} = 1$ indicates rejection of the hypothesis H_0 . Typically, $\phi_{n,m}$ is based on a test statistic Ψ measuring discrepancy between samples, rejecting H_0 if Ψ exceeds a threshold.

A test’s performance is characterized by its Type I (false positive) and Type II (false negative) errors, defined as follows. Let $\epsilon > 0$ and define

$$\mathcal{P}_1(\epsilon) := \{(\mu, \nu) : d(\mu, \nu) \geq \epsilon\}.$$

A level- α test controls the Type I error uniformly by $\alpha \in (0, 1)$, and has power at least $1 - \beta$ if its Type II error is bounded by $\beta \in (0, 1)$ uniformly over $\mathcal{P}_1(\epsilon)$:

$$\sup_{\mu=\nu} \mathbb{P}_{\mu \times \nu}(\phi_{n,m}=1) \leq \alpha, \quad \sup_{(\mu,\nu) \in \mathcal{P}_1(\epsilon)} \mathbb{P}_{\mu \times \nu}(\phi_{n,m}=0) \leq \beta,$$

where probabilities are taken over the samples and any additional randomness (e.g., random projections).

In Section 5, we establish that our test is of level α and characterize the minimal separation between distributions detectable with power at least $1 - \beta$.

2.2 Sliced Wasserstein Distance

For $p \geq 1$, the p -Wasserstein distance (Villani, 2008, Section 6) between probability measures μ, ν in the

space of probability measures with finite p -th moments $\mathcal{P}_p(\mathbb{R}^d)$, is defined as

$$W_p(\mu, \nu) := \left(\inf_{\pi \in \Gamma(\mu, \nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^p d\pi(x, y) \right)^{1/p},$$

where $\Gamma(\mu, \nu)$ is the set of couplings with marginals μ and ν . The p -Wasserstein distance W_p is a natural candidate for two-sample testing: it vanishes if and only if the distributions coincide and captures geometric discrepancies. However, its minimax separation rate scales as $n^{-c/d}$ for some $c > 1$, making it impractical in moderate to high dimensions (Ba et al., 2011; Chewi et al., 2024, Section 2.5).

To address this limitation, we use the sliced Wasserstein (SW) distance (Rabin et al., 2011), defined next.

Definition 1 (SW distance) For $p \geq 1$ and $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$, the sliced Wasserstein distance is

$$SW_p(\mu, \nu) := \left(\int_{\mathbb{S}^{d-1}} W_p^p(\Pi_{\#}^{\theta} \mu, \Pi_{\#}^{\theta} \nu) \sigma(d\theta) \right)^{1/p}, \quad (1)$$

where $\Pi^{\theta}(x) = \langle \theta, x \rangle$ is the projection onto direction θ on the unit sphere $\mathbb{S}^{d-1}$, and σ is the uniform measure on the unit sphere.

By averaging one-dimensional Wasserstein distances over random projections, SW_p benefits from a dimension-free sample complexity rate while preserving geometric interpretability and has been successfully applied in large-scale statistical problems (see Sections 1 and 3). In order to estimate the SW distance, we replace the unknown measures μ and ν by their empirical counterpart

$$\hat{\mu}_n := \frac{1}{n} \sum_{i=1}^n \delta_{Y_i} \quad \text{and} \quad \hat{\nu}_m := \frac{1}{m} \sum_{i=1}^m \delta_{Z_i}.$$

Moreover, we approximate the expectation with respect to the uniform distribution on the sphere σ by Monte Carlo sampling: we draw $L \geq 1$ independent directions $\Theta = (\theta_1, \dots, \theta_L)$ from σ and compute the tractable statistic

$$\widehat{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m) := \frac{1}{L} \sum_{\ell=1}^L W_p^p(\Pi_{\#}^{\theta_{\ell}} \hat{\mu}_n, \Pi_{\#}^{\theta_{\ell}} \hat{\nu}_m). \quad (2)$$

For brevity, we write $\widehat{SW}_p^p$ for the estimator in (2).

2.3 Minimax optimality

Minimax theory provides a benchmark for testing by characterizing the smallest detectable difference for level- α tests (Ingster, 1993; Baraud, 2002). Define

$$\Phi_{n+m, \alpha} := \left\{ \phi_{n,m} : \sup_{\mu=\nu} \mathbb{P}_{\mu \times \nu}(\phi_{n,m} = 1) \leq \alpha \right\}$$

as the set of tests controlling Type I error at level α based on n samples from μ and m samples from ν . Given $\epsilon > 0$, the minimax risk is then defined as the minimal worst-case Type II error over $\mathcal{P}_1(\epsilon)$ among all level- α tests,

$$R_{n+m, \epsilon}^{\dagger} := \inf_{\phi_{n,m} \in \Phi_{n+m, \alpha}} \sup_{(\mu, \nu) \in \mathcal{P}_1(\epsilon)} \mathbb{P}_{\mu \times \nu}(\phi_{n,m} = 0).$$

The minimax separation $\epsilon_{n,m}^{\dagger}$ is the smallest discrepancy $\epsilon > 0$ such that a level- α test exists with Type II error at most β :

$$\epsilon_{n,m}^{\dagger} := \inf \left\{ \epsilon > 0 : R_{n+m, \epsilon}^{\dagger} \leq \beta \right\}, \quad (3)$$

for fixed $\beta \in (0, 1 - \alpha)$. Intuitively, $\epsilon_{n,m}^{\dagger}$ characterizes the detection boundary below which no α -level test can reliably distinguish alternatives from the null.

For a given separation $\epsilon > 0$, the maximum Type II error of a test $\phi_{n,m} \in \Phi_{n+m, \alpha}$ over the class of alternative distributions $\mathcal{P}_1(\epsilon)$ is

$$R_{n+m, \epsilon}(\phi_{n,m}) := \sup_{(\mu, \nu) \in \mathcal{P}_1(\epsilon)} \mathbb{P}_{\mu \times \nu}(\phi_{n,m} = 0).$$

Then, its minimum separation is

$$\tilde{\epsilon}_{\phi, n, m} := \inf \{ \epsilon > 0 : R_{n+m, \epsilon}(\phi_{n,m}) \leq \beta \}.$$

A test $\phi_{n,m}$ is minimax rate-optimal if its minimum separation $\tilde{\epsilon}_{\phi, n, m}$ is equivalent to $\epsilon_{n,m}^{\dagger}$ up to constant factors. We show in Section 5 that the sliced Wasserstein test we consider is minimax rate-optimal.

3 RELATED WORK

The problem of two-sample testing has been extensively studied for over a century, with two main classes of approaches: *parametric* and *non-parametric*.

Classical parametric tests. Early approaches to the two-sample problem include the Kolmogorov–Smirnov test (Kolmogorov, 1933; Smirnov, 1948), Student’s t -test for mean comparison (Student, 1908), and Hotelling’s T^2 test as a multivariate generalization of the t -test (Hotelling, 1931). These procedures are computationally simple and well understood, but they are either restricted to the univariate setting (e.g., the KS test) or to testing specific moments such as the mean (e.g., t -test, Hotelling’s T^2), which limits their applicability in high-dimensional or nonparametric problems.

Kernel-based methods. Kernel two-sample testing has emerged as a popular nonparametric alternative to classical parametric tests, designed to handle

complex and high-dimensional data. The maximum mean discrepancy (MMD) test of Gretton et al. (2012) is widely used for its tractability and power, and it is closely related to energy distance (Székely et al., 2004) through the equivalence established by Sejdinovic et al. (2013). Because test performance is sensitive to kernel choice, later work has proposed kernel aggregation for adaptivity (Biggs et al., 2023; Schrab et al., 2023). In parallel, scalable variants have been developed for large-scale problems, including random Fourier feature approximations (Choi and Kim, 2024; Mukherjee and Sriperumbudur, 2025), Nyström subsampling (Chatalic et al., 2025), and coresets-based methods (Domingo-Enrich et al., 2023). While conceptually related to OT-based approaches—since both compare distributions via a discrepancy measure—kernel tests are grounded in reproducing kernel Hilbert space embeddings, whereas OT-based methods exploit the geometry induced by transport costs.

Optimal transport-based methods. Parallel to kernel testing, there has been growing interest in tests based on optimal transport. Ramdas et al. (2017) survey connections between Wasserstein distance, MMD, and energy distance, framing entropic OT as interpolating between transport and kernel discrepancies. Building on this, several works analyze the statistical properties of Wasserstein-based tests (Imaizumi et al., 2022; González-Delgado et al., 2023) and propose projection-based approaches such as projected Wasserstein tests (Wang et al., 2021a,b). Sliced variants, which exploit the tractability of one-dimensional OT, have been studied extensively: the sliced Wasserstein distance was introduced by Rabin et al. (2011), with theoretical properties developed by Bonnotte (2013). Generalizations include generalized and max-sliced Wasserstein distances (Kolouri et al., 2019; Deshpande et al., 2019), subspace-robust Wasserstein (Paty and Cuturi, 2019), and recent work on statistical guarantees for max-sliced variants (Boedihardjo, 2025; Wang et al., 2025). Closest to our work, Hu and Lin (2025) propose a max-sliced Wasserstein test with bootstrap calibration and asymptotic validity, while finite-sample guarantees remain an open question.

Project-then-test methods. A related line of work uses random projections to reduce dimensionality before applying classical test statistics. For example, Lopes et al. (2011) propose projecting high-dimensional data onto random directions and then applying Hotelling’s T^2 test.

Calibration and optimality. Permutation testing plays a central role in our approach. Its finite-sample validity has been long established (Hoeffding, 1952;

Hemerik and Goeman, 2018), and more recent results show that permutation tests can achieve the minimax optimal detection boundary (Kim et al., 2022). This positions permutation calibration as a natural complement to Wasserstein-based discrepancies, allowing us to combine strong finite-sample guarantees with minimax optimality.

In summary, classical tests are restricted in scope, kernel tests provide flexible and efficient nonparametric alternatives, and OT-based tests capture geometric aspects of the distributions but often lack finite-sample guarantees. We propose a permutation-based sliced Wasserstein test that combines size control with non-asymptotic power bounds, achieves the minimax rate $n^{-1/2}$, and quantifies the trade-off between computational efficiency and statistical power.

4 PERMUTATION SW TEST

Before presenting our method in Algorithm 1, we detail the permutation procedure and the obtained test.

4.1 Permutation procedure

Recall $\mathcal{Y}_n = (Y_1, \dots, Y_n)$ and $\mathcal{Z}_m = (Z_1, \dots, Z_m)$ denote the two samples, and let $\Theta = (\theta_1, \dots, \theta_L)$ be the set of projection directions (drawn i.i.d. uniformly from the sphere). We form the pooled dataset $\mathcal{X}_{n+m} = (X_1, \dots, X_{n+m})$ as

$$X_i := Y_i \quad \text{for } 1 \leq i \leq n, \quad X_{n+i} := Z_i \quad \text{for } 1 \leq i \leq m.$$

For a permutation $\pi \in S_{n+m}$, the symmetric group on $\{1, \dots, n+m\}$, define the permuted dataset $\mathcal{X}_{n+m}^\pi = (X_{\pi(1)}, \dots, X_{\pi(n+m)})$. Construct empirical measures

$$\hat{\mu}_n^\pi := \frac{1}{n} \sum_{i=1}^n \delta_{X_{\pi(i)}}, \quad \hat{\nu}_m^\pi := \frac{1}{m} \sum_{i=n+1}^{n+m} \delta_{X_{\pi(i)}},$$

and compute the empirical sliced Wasserstein distance

$$\widehat{\text{SW}}_p^{p,\pi} := \widehat{\text{SW}}_p^p(\hat{\mu}_n^\pi, \hat{\nu}_m^\pi). \quad (4)$$

Let $N := n+m$, and denote by $F_{\widehat{\text{SW}}_p^{p,\pi}}$ the permutation empirical cumulative distribution function:

$$F_{\widehat{\text{SW}}_p^{p,\pi}}(t) := \frac{1}{|S_N|} \sum_{\pi \in S_N} \mathbf{1}\{\widehat{\text{SW}}_p^{p,\pi} \leq t\},$$

where S_N is the set of all permutations of $\{1, \dots, N\}$. We write the $1 - \alpha$ quantile of this distribution as

$$c_{1-\alpha,N} := \inf\{t : F_{\widehat{\text{SW}}_p^{p,\pi}}(t) \geq 1 - \alpha\}. \quad (5)$$

Decision making. Given the quantile $c_{1-\alpha, N}$, we reject the null hypothesis if $\widehat{\text{SW}}_p^p > c_{1-\alpha, N}$. This critical value ensures finite-sample Type I error control under the permutation-invariance (exchangeability) assumption, which holds for both two-sample and independence testing problems (see Section 5.1).

Remark 2 *Exact computation of the critical value (5) is generally infeasible for large samples, so it is commonly approximated via Monte Carlo simulations. This approximation can be made arbitrarily accurate by increasing the number of sampled permutations (see Appendix or Domingo-Enrich et al. (2023, Lem. 6)).*

Let r be the uniform distribution on S_N and let $\mathbb{Z}_B := (\pi_b)_{1 \leq b \leq B}$ be a collection of i.i.d. samples from r . Define $\pi_{B+1} := \text{id}$ as the identity permutation. Following the approach of Romano and Wolf (2005, Lemma 1), and in order to achieve the prescribed non-asymptotic test level, we then consider the statistics $\widehat{\text{SW}}_p^{p, \pi_b}$ for $b = 1, \dots, B+1$ and estimate the $(1-\alpha)$ quantile of the permutation distribution by

$$\hat{c}_{1-\alpha, N}^B := \inf \left\{ t : \frac{1}{B+1} \sum_{i=1}^{B+1} \mathbf{1}\{\widehat{\text{SW}}_p^{p, \pi_i} \leq t\} \geq 1-\alpha \right\}.$$

Letting $\widehat{\text{SW}}_p^{p, r \bullet 1} \leq \widehat{\text{SW}}_p^{p, r \bullet 2} \leq \dots \leq \widehat{\text{SW}}_p^{p, r \bullet (B+1)}$ denote the order statistics of $\{\widehat{\text{SW}}_p^{p, \pi_b}\}_{b=1}^{B+1}$, the quantile estimator equals $\widehat{\text{SW}}_p^{p, r \bullet \lceil (B+1)(1-\alpha) \rceil}$. The Sliced Wasserstein Test is then defined as

$$\Delta(\mathcal{Y}_n, \mathcal{Z}_m, \mathbb{Z}_B, \Theta) := \mathbf{1}(\widehat{\text{SW}}_p^p > \hat{c}_{1-\alpha, N}^B).$$

Algorithm 1 summarizes the testing procedure.

Remark 3 *While bootstrap or subsampling methods are often used to calibrate critical values (see, e.g., Hu and Lin (2025)), their asymptotic validity rely on well-behaved limiting distributions and do not guarantee uniform finite-sample size control, as required for the finite sample minimax framework that we consider.*

4.2 Computational complexity

With a one-time pre-computation of projections $O(LdN)$ and per-projection sorting $O(LN \log(N))$, and the main loop $O(LBN)$, our method achieves an overall time complexity of $O(LN(d + \log N + B))$. Moreover, storing the original and projected datasets, as well as the permuted statistics requires space complexity $O(Nd + NL + B)$.

Note that, choosing $L = N$ projections leads to quadratic time complexity (up to logarithmic factors) in the sample size, comparable to that of the MMD

Algorithm 1 SW-Permutation Test

Require: Datasets $\mathcal{Y}_n = (Y_1, \dots, Y_n)$ and $\mathcal{Z}_m = (Z_1, \dots, Z_m)$; significance level $\alpha \in (0, 1)$; number of permutations B ; number of projection L .

Ensure: Decision $\Delta \in \{0, 1\}$

1: Form pooled dataset $\mathcal{X}_{n+m} = (X_1, \dots, X_{n+m})$ as

$$X_i := \begin{cases} Y_i & 1 \leq i \leq n, \\ Z_{i-n} & n < i \leq n+m \end{cases}$$

2: Sample $\theta_1, \dots, \theta_L \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\mathbb{S}^{d-1})$

3: Sample $\pi_1, \dots, \pi_B \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(S_{n+m})$

4: Set $\pi_{B+1} := \text{id}$ (identity permutation)

5: **for** $b = 1$ to $B+1$ **do**

6: Define empirical measures

$$\hat{\mu}_n^{\pi_b} := \frac{1}{n} \sum_{i=1}^n \delta_{X_{\pi_b(i)}}, \quad \hat{\nu}_m^{\pi_b} := \frac{1}{m} \sum_{i=n+1}^{n+m} \delta_{X_{\pi_b(i)}}$$

7: Compute statistic $\widehat{\text{SW}}_p^{p, \pi_b} := \widehat{\text{SW}}_p^p(\hat{\mu}_n^{\pi_b}, \hat{\nu}_m^{\pi_b})$

8: **end for**

9: Set the critical $\hat{c}_{1-\alpha, n+m}^B$ as the $(1-\alpha)$ empirical quantile of $\{\widehat{\text{SW}}_p^{p, \pi_b}\}_{b=1}^{B+1}$

10: **return** $\Delta := \mathbf{1}(\widehat{\text{SW}}_p^{p, \pi_{B+1}} > \hat{c}_{1-\alpha, n+m}^B)$

two-sample test. Reducing the number of projections can mitigate this cost, trading off between statistical power and computational efficiency. A characterization of this trade-off can be obtained from Theorem 5.

5 THEORETICAL ANALYSIS

In this section, we analyze the proposed test by studying its level and power. Without loss of generality, we assume $n \leq m$. We let $B \geq 1$ denote the number of permutations and $L \geq 1$ the number of projections. For space constraints and to improve readability, all proofs are deferred to the Appendix and all constants are made explicit in the proof.

5.1 Level of the test

We begin by establishing that the proposed test controls the Type I error.

It is now a well-known fact that permutation-based tests control the Type I error at finite samples for any test statistic when the data are exchangeable under the null hypothesis (Hoeffding, 1952; Lehmann and Romano, 2005). This property is crucial since it ensures validity without asymptotic approximations. In the

two-sample setting, this holds under $H_0 : \mu = \nu$ since all observations are i.i.d. and the joint distribution is invariant under permutations of group labels. We formalize this below.

Theorem 4 (Type I error Control) *The test Δ defined in Algorithm 1 has non-asymptotic level α for any $\alpha \in (0, 1)$. That is*

$$\mathbb{P}_{\mu \times \mu \times r \times \sigma}(\Delta(\mathcal{Y}_n, \mathcal{Z}_m, \mathbb{Z}_B, \Theta) = 1) \leq \alpha,$$

where the probability is taken over the samples, projection directions, and permutation randomness.

5.2 Power of the test

While Type I error control follows from standard permutation arguments, analyzing the test's power is more challenging. The main difficulty is that the critical value is defined as a data-dependent permutation quantile of the sliced Wasserstein statistic. For MMD, existing analyses exploit its U - or V -statistic structure together with associated concentration inequalities (Schrab et al., 2023; Domingo-Enrich et al., 2023), but the sliced Wasserstein distance does not admit such a representation, preventing the direct use of these techniques. To address this issue, we leverage properties of the sliced Wasserstein distance under permutations, in particular, a permutation bounded-differences (McDiarmid-type) inequality and an optimal matching bound from Bobkov and Ledoux (2021, Corollary 5). These ingredients yield the non-asymptotic power bound stated in Theorem 5.

Theorem 5 (Power Control) *Let $\beta \in (0, 1)$ and $\frac{1}{B+1} \leq \alpha < 1$. The test described in Algorithm 1 has power at least $1 - \beta$ provided that*

$$\text{SW}_p^p(\mu, \nu) \geq \frac{C(D, p, \alpha, \beta, B)}{\sqrt{L \wedge n}},$$

where the constant $C(D, p, \alpha, \beta, B)$ is made explicit in the proof.

This result shows that the SW test achieves the parametric separation rate $n^{-1/2}$ up to constants. To match the sample size, the number of projections L should scale proportionally to n , ensuring that projection variance does not limit power.

5.3 Minimax optimality

To assess the statistical optimality of the proposed test, we now study its minimax separation rate. Establishing minimax lower bounds in nonparametric testing typically involves constructing least favorable alternatives and adapting classical arguments, which is

the approach we take here. These are the first minimax lower bounds for two-sample testing with respect to the sliced Wasserstein distance. Consider the class

$$\mathcal{P}_{\mathbb{R}^d}(D) := \{\mu \in \mathcal{P}(\mathbb{R}^d) : \text{diam}(\text{supp}(\mu)) \leq D\},$$

comprising all probability measures on $\mathbb{R}^d$ with support diameter at most $D > 0$. Taking $L \geq cn$ for some absolute constant $c > 0$, Theorem 5 implies an upper bound on the minimax separation rate (see Eq. (3))

$$\epsilon_{n,m}^\dagger \leq Cn^{-1/2}.$$

We now establish lower bounds in two scenarios to assess the tightness of the upper bound. First, we provide a lower bound for the simpler case of multinomial distributions. We will then consider more generally distributions with bounded support. The following result is obtained via an adaptation of the classical lower bound technique from Ingster (1987).

Proposition 6 *Let $\alpha \in (\frac{1}{B+1}, 1)$ and $\beta \in (0, 1 - \alpha)$. For any integer $d \geq 1$, let $[d] := \{1, \dots, d\}$ and denote by $\mathcal{P}_{\text{Multi}}^{(d)}$ the class of multinomial distributions on $[d]$. For the two-sample testing problem over $\mathcal{P}_{\text{Multi}}^{(d)} \times \mathcal{P}_{\text{Multi}}^{(d)}$, the minimax separation rate satisfies*

$$\epsilon_{n,m}^\dagger \geq \frac{C(d, \alpha, \beta)}{\sqrt{n}},$$

where the constant $C(d, \alpha, \beta)$ is made explicit in the proof.

The above result establishes that no test can achieve faster than $n^{-1/2}$ separation in the simple multinomial setting.

Turning to the more general class of distributions with bounded support $\mathcal{P}_{\mathbb{R}^d}(D)$, we obtain the following lower bound using Le Cam's two-point method (Le Cam, 1973, 2012).

Proposition 7 *Let $\alpha, \beta \in (0, 1)$ with $\alpha + \beta < 0.5$. For the two-sample testing problem over $\mathcal{P}_{\mathbb{R}^d}(D)$, the minimax separation rate satisfies*

$$\epsilon_{n,m}^\dagger \geq \frac{C(D, \alpha, \beta, p)}{\sqrt{n}}.$$

Thus, the lower bound extends beyond discrete settings to general bounded-support distributions.

Together with Theorem 5, these results imply that the permutation SW test achieves the minimax optimal separation rate $n^{-1/2}$ over both multinomial and bounded-support distributions, matching known optimal rates for MMD-based tests (Kim et al., 2022, Position 4.4; Kim and Schrab, 2023, Section E.10.1).

6 NUMERICAL EXPERIMENTS

To complement our theoretical findings, we empirically evaluate the performance of the SW test described in Algorithm 1. We focus on the balanced case $n = m$ and assess performance across various numbers of projections L . We compare against several baselines: the Projected Wasserstein (PW) test (Wang et al., 2021a) and the MMD test (Gretton et al., 2012). For MMD, we consider linear, Gaussian, and Laplace kernels, with bandwidths for the latter two selected via the median heuristic (Garreau et al., 2017). For the PW test, we use the authors’ official implementation with default settings (projection dimension 3 and a 50/50 split for training and testing), as no practical guidance on these parameters was provided. All experiments were run on a machine with an Intel Core Ultra 9 185H CPU and 64 GB of RAM.

As discussed in Section 3, several scalable variants of MMD have been developed for large-scale problems. However, these approximations do not increase statistical power beyond standard MMD, so we exclude them from our benchmarks and focus on contrasting how SW- and MMD-based tests distinguish distributions. Moreover, we also exclude MMD aggregation methods: while aggregation improves adaptivity and can outperform individual kernels, our goal is to compare discrepancies rather than kernel selection strategies, so we restrict attention to standard kernels.

In all experiments, we set the significance level to $\alpha = 0.05$ and use $B = 200$ permutations. Empirical power is estimated from 150 independent repetitions, parallelized with Joblib. Error bars represent 95% confidence intervals computed from the normal approximation of the empirical proportion (mean ± 1.96 standard errors).

6.1 Tests on Synthetic Datasets

Our first experiment evaluates power under a Gaussian covariance shift: samples are drawn from $\mu = \mathcal{N}(0, I_{60})$ and $\nu = \mathcal{N}(0, \Sigma)$, where $\Sigma = \text{diag}(\delta^2, \delta^2, 1, \dots, 1)$ with shift magnitude $\delta = 2.7$.

The second scenario examines geometric differences between distributions: the uniform distribution on the 5D unit sphere versus the uniform distribution on the 5D unit ball. This setting tests performance when one distribution lies on a lower-dimensional manifold embedded within the other’s support.

As shown in Figure 1, the PW test achieves the highest power under covariance shifts, consistent with Wang et al. (2021b), while the SW test performs competitively with Gaussian and Laplace MMD. In the ball-vs.-sphere scenario, the MMD test with a Laplace

kernel attains the best performance, but the SW test substantially outperforms both Gaussian MMD and PW, even with a limited number of projection directions. The weaker performance of Gaussian MMD likely stems from over-smoothing under the median heuristic, which blurs sharp radial differences, whereas the Laplace kernel’s slower decay preserves more of this contrast. By contrast, MMD with a linear kernel performs worst across both settings, since with $k(x, y) = x^\top y$ it reduces to a mean-difference test, which vanishes in these cases. Overall, while the SW test is not always the most powerful, it delivers consistently strong performance across benchmarks without requiring parameter tuning, demonstrating robustness to diverse distributional differences.

6.2 Tests on MNIST Handwritten Digits

We apply the two-sample testing procedures to the MNIST dataset (LeCun et al., 2012). Let μ_6 denote the uniform distribution over images of the digit 6 and μ_9 the corresponding distribution for digit 9. We define the mixture $\nu = 0.85\mu_6 + 0.15\mu_9$. Following Wang et al. (2021b), we preprocess the dataset by applying a sigmoid transformation to each image so that all pixel values lie within $[0, 1]$. The SW test effectively discriminates between μ_6 and ν by capturing localized intensity differences in 1D projections.

This experiment highlights the SW test’s ability to detect subtle distributional differences in high-dimensional image data without requiring kernel parameter tuning, confirming its practical applicability.

6.3 Effect of the Number of Projections

While the SW test generally performs well on our three benchmark datasets, it underperforms relative to MMD tests with Gaussian, Laplace, and linear kernels in an additional Gaussian mean-shift experiment for a range of projection numbers $L \in 0.5n, n, 2n$ (see Appendix). To better understand this behavior, we examine the effect of increasing the number of projections L on both power and computation time. Fixing $n = m = 140$ in the mean-shift setting, Figure 2 shows that power improves steadily with larger L , though at the cost of increased computation time (Figure 4).

Fortunately, the computational bottleneck of permutation-based SW testing is mitigated by its projection structure: since it computes many independent 1D projections, the workload is naturally amenable to GPU parallelization. This allows scaling the number of projections with only moderate increases in wall-clock time, given sufficient GPU resources. In contrast, MMD computations hinge on the kernel matrix, whose size is fixed once the




Figure 1: Power across three scenarios: Gaussian covariance shift, ball vs. sphere, and MNIST mixture.




 Figure 2: SW Test power vs. number of projections (fixed sample sizes $n = m = 140$)


Figure 3: Computation time (log scale) of SW tests


Table 1: Type I error for the SW test

Test Statistics	Type I Error
SW ($L = 0.5n$)	0.04982 ± 0.00543
SW ($L = n$)	0.04844 ± 0.00501
SW ($L = 2n$)	0.04961 ± 0.00461

sample size is given, limiting further scaling through additional parallelism. In practice, we find that SW implementations can better exploit hardware parallelism and thus often scale better than naive MMD implementations under similar resource constraints.

To validate this, we compared computation times using two sets of 140 samples from a 60-dimensional mean-shifted Gaussian distribution. We measured the average execution time over 100 repetitions on a local CPU, Google Colab’s T4, and A100 GPUs. Results in Figure 3 confirm substantial speedups for the SW test on GPUs, though larger-scale studies would be valuable to fully assess its scalability in practice.

Finally, to illustrate Theorem 4, we assessed the empirical Type I error of the SW test. In each of 2000 repetitions, two independent samples of size 50 were drawn from $\mathcal{N}(0, I_{60})$. The results, reported in Table 1 for three choices of L , show that the permutation-based SW test effectively controls the Type I error rate even at small sample sizes.


 Figure 4: Computation time for ball vs. sphere ¹

¹To leverage available CPU resources, we perform independent repetitions in parallel across all CPU cores and report the mean wall-clock time per repetition. Due to heterogeneous core performance, the average time can exceed that observed when using a single core. All methods are timed on the same inputs within each repetition, so their relative ranking is unaffected.

In summary, our experiments demonstrate that the SW test offers a robust, parameter-free alternative to kernel-based methods, balancing statistical power and computational efficiency, especially when leveraging parallel hardware.

7 Conclusion

We proposed a permutation-based two-sample test using the sliced Wasserstein distance, establishing finite-sample validity, non-asymptotic power bounds, and minimax optimality over multinomial and bounded-support alternatives. To our knowledge, this is the first Wasserstein-based test with finite-sample guarantees. Our analysis quantified the trade-off between the number of projections and statistical power, and experiments showed that the test achieves consistently strong performance across benchmarks without parameter tuning, while remaining scalable on parallel hardware. Future directions include extending minimax optimality results to broader smoothness classes and exploring adaptive or non-uniform slicing strategies (e.g., generalized projections (Kolouri et al., 2019)).

Acknowledgements

The authors would like to thank Pierre Bizeul for pointing out an interesting reference and Mohamed Hebiri for helpful discussions.

References

- Albert, M. (2015). *Tests of independence by bootstrap and permutation: an asymptotic and non-asymptotic study. Application to neurosciences*. PhD thesis, Université Nice Sophia Antipolis.
- Albert, M. (2019). Concentration inequalities for randomly permuted sums. In *High Dimensional Probability VIII: The Oaxaca Volume*, pages 341–383. Springer.
- Albert, M., Bouret, Y., Fromont, M., and Reynaud-Bouret, P. (2015). Bootstrap and permutation tests of independence for point processes.
- Arias-Castro, E., Pelletier, B., and Saligrama, V. (2018). Remember the curse of dimensionality: The case of goodness-of-fit testing in arbitrary dimension. *Journal of Nonparametric Statistics*, 30(2):448–471.
- Ba, K. D., Nguyen, H. L., Nguyen, H. N., and Rubinfeld, R. (2011). Sublinear time algorithms for earth mover’s distance. *Theory of Computing Systems*, 48(2):428–442.
- Balakrishnan, N. and Nevzorov, V. B. (2004). *A primer on statistical distributions*. John Wiley & Sons.
- Baraud, Y. (2002). Non-asymptotic minimax rates of testing in signal detection. *Bernoulli*, 8(5):577 – 606.
- Biggs, F., Schrab, A., and Gretton, A. (2023). Mmd-fuse: Learning and combining kernels for two-sample testing without data splitting. *Advances in Neural Information Processing Systems*, 36:75151–75188.
- Bobkov, S. G. and Ledoux, M. (2021). A simple fourier analytic proof of the akt optimal matching theorem. *The Annals of Applied Probability*, 31(6):2567–2584.
- Boedihardjo, M. T. (2025). Sharp bounds for max-sliced wasserstein distances. *Foundations of Computational Mathematics*, pages 1–32.
- Bonnotte, N. (2013). *Unidimensional and evolution methods for optimal transportation*. PhD thesis, ENS Cachan.
- Borgwardt, K. M., Gretton, A., Rasch, M. J., Kriegel, H.-P., Schölkopf, B., and Smola, A. J. (2006). Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57.
- Canonne, C. L. (2022). A short note on an inequality between kl and tv. *arXiv preprint arXiv:2202.07198*.
- Chakravarti, P., Kuusela, M., Lei, J., and Wasserman, L. (2023). Model-independent detection of new physics signals using interpretable semisupervised classifier tests. *The Annals of Applied Statistics*, 17(4):2759–2795.
- Chatalic, A., Letizia, M., Schreuder, N., and Rosasco, L. (2025). A scalable nystrom-based kernel two-sample test with permutations. *arXiv preprint arXiv:2502.13570*.
- Chewi, S., Niles-Weed, J., and Rigollet, P. (2024). Statistical optimal transport. *arXiv preprint arXiv:2407.18163*, 3.
- Choi, I. and Kim, I. (2024). Computational-statistical trade-off in kernel two-sample testing with random fourier features. *arXiv preprint arXiv:2407.08976*.
- Deshpande, I., Hu, Z., Song, C., and et al. (2019). Max-sliced wasserstein distance and its use for gans. In *CVPR*.
- Domingo-Enrich, C., Dwivedi, R., and Mackey, L. (2023). Compress then test: Powerful kernel testing in near-linear time. In Ruiz, F., Dy, J., and van de Meent, J.-W., editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 1174–1218. PMLR.

- Dubhashi, D. P. and Ranjan, D. (1996). Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25).
- Dümbgen, L. (1998). Symmetrization and decoupling of combinatorial random elements. *Statistics & probability letters*, 39(4):355–361.
- El-Yaniv, R. and Pechyony, D. (2009). Transductive rademacher complexity and its applications. *Journal of Artificial Intelligence Research*, 35:193–234.
- Fournier, N. and Guillin, A. (2015). On the rate of convergence in wasserstein distance of the empirical measure. *Probability theory and related fields*, 162(3):707–738.
- Gao, I., Liang, P., and Guestrin, C. (2024). Model equality testing: Which model is this API serving? *arXiv preprint arXiv:2410.20247*.
- Garreau, D., Jitkrittum, W., and Kanagawa, M. (2017). Large sample analysis of the median heuristic. *arXiv preprint arXiv:1707.07269*.
- González-Delgado, J. A., Weed, J., and Del Barrio, E. (2023). Two-sample tests for distributions supported on the flat torus. *Annals of Statistics*.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. J. (2012). A kernel two-sample test. In *Journal of Machine Learning Research*, volume 13, pages 723–773.
- Grossi, S., Letizia, M., and Torre, R. (2025). Refereeing the referees: evaluating two-sample tests for validating generators in precision sciences. *Machine Learning: Science and Technology*, 6(1):015052.
- Hemerik, J. and Goeman, J. J. (2018). Exact testing with random permutations. *TEST*, 27(4):811–825.
- Hoeffding, W. (1952). The large-sample power of tests based on permutations of observations. *Annals of Mathematical Statistics*, 23(2):169–192.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30.
- Horváth, L., Reeder, R., and Kokoszka, P. (2014). Estimation of the mean of functional time series and a two-sample problem. *Quality control and applied statistics*, 59(1):51–52.
- Hotelling, H. (1931). The generalization of student’s ratio. *Annals of Mathematical Statistics*, 2(3):360–378.
- Hu, X. and Lin, Z. (2025). Two-sample distribution tests in high dimensions via max-sliced wasserstein distance and bootstrapping. *Biometrika*, 112(2):asaf001.
- Imaizumi, M., Ota, H., and Hamaguchi, K. (2022). Hypothesis testing for distributions supported on general metric spaces with wasserstein distances. *arXiv preprint arXiv:2201.10456*.
- Ingster, Y. I. (1987). Minimax testing of nonparametric hypotheses on a distribution density in the l_p metrics. *Theory of Probability & Its Applications*, 31(2):333–337.
- Ingster, Y. I. (1993). Asymptotically minimax hypothesis testing for nonparametric alternatives. i, ii, iii. *Math. Methods Statist*, 2(2):85–114.
- Jameson, G. (2015). A simple proof of stirling’s formula for the gamma function. *The Mathematical Gazette*, 99(544):68–74.
- Kim, I., Balakrishnan, S., and Wasserman, L. (2022). Minimax optimality of permutation tests. *The Annals of Statistics*, 50(1):225–251.
- Kim, I. and Schrab, A. (2023). Differentially private permutation tests: Applications to kernel methods. *arXiv preprint arXiv:2310.19043*.
- Kolmogorov, A. (1933). Sulla determinazione empirica di una legge di distribuzione. *Giornale dell’Istituto Italiano degli Attuari*, 4:83–91.
- Kolouri, S., Rohde, G. K., and Hoffmann, H. (2019). Generalized sliced wasserstein distances. In *Advances in Neural Information Processing Systems*.
- Le Cam, L. (1973). Convergence of estimates under dimensionality restrictions. *The Annals of Statistics*, pages 38–53.
- Le Cam, L. (2012). *Asymptotic methods in statistical decision theory*. Springer Science & Business Media.
- LeCun, Y., Cortes, C., and Burges, C. J. C. (2012). The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>. [Online; accessed 25-April-2018].
- Lehmann, E. L. and Romano, J. P. (2005). *Testing statistical hypotheses*. Springer.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. (2017). Mmd gan: Towards deeper understanding of moment matching network. *Advances in neural information processing systems*, 30.
- Liu, Y., Li, C.-L., and Póczos, B. (2018). Classifier two sample test for video anomaly detections. In *BMVC*, page 71.
- Lopes, M. E., Jacob, L., and Wainwright, M. J. (2011). More powerful two-sample tests via random projections. *Neural Information Processing Systems*.
- McDiarmid, C. et al. (1989). On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188.
- Miles, R. R., Roberts, R. F., Putnam, A. R., and Roberts, W. L. (2004). Comparison of serum and

- heparinized plasma samples for measurement of chemistry analytes. *Clinical chemistry*, 50(9):1704–1706.
- Mukherjee, S. and Sriperumbudur, B. K. (2025). Minimax optimal kernel two-sample tests with random features. *arXiv preprint arXiv:2502.20755*.
- Müller, A. (1997). Integral probability metrics and their generating classes of functions. *Advances in applied probability*, 29(2):429–443.
- Nadjahi, K. (2021). *Sliced-Wasserstein distance for large-scale machine learning: theory, methodology and extensions*. PhD thesis, Institut polytechnique de Paris.
- Nietert, S., Goldfeld, Z., Sadhu, R., and Kato, K. (2022). Statistical, robustness, and computational guarantees for sliced wasserstein distances. *Advances in Neural Information Processing Systems*, 35:28179–28193.
- Ohana, R., Nadjahi, K., Rakotomamonjy, A., and Ralaivola, L. (2023). Shedding a pac-bayesian light on adaptive sliced-wasserstein distances. In *International Conference on Machine Learning*, pages 26451–26473. PMLR.
- Paty, F.-P. and Cuturi, M. (2019). Subspace robust wasserstein distances. In *International Conference on Machine Learning*, pages 5072–5081.
- Politis, D. N., Romano, J. P., and Wolf, M. (1999). Subsampling in the iid case. In *Subsampling*, pages 39–64. Springer.
- Rabin, J., Peyré, G., Delon, J., and Bernot, M. (2011). Wasserstein barycenter and sliced wasserstein distances. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer.
- Ramdas, A., García Trillos, N., and Cuturi, M. (2017). On wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47.
- Rodríguez-Vítores, D., del Barrio, E., and Loubes, J.-M. (2025). An improved central limit theorem for the empirical sliced wasserstein distance. *arXiv preprint arXiv:2503.18831*.
- Romano, J. P. (1988). A bootstrap revival of some nonparametric distance tests. *Journal of the American Statistical Association*, 83(403):698–708.
- Romano, J. P. and Wolf, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association*, 100(469):94–108.
- Schrab, A., Kim, I., Albert, M., Laurent, B., Guedj, B., and Gretton, A. (2023). Mmd aggregated two-sample test. *Journal of Machine Learning Research*, 24(194):1–81.
- Sejdinovic, D., Sriperumbudur, B., Gretton, A., and Fukumizu, K. (2013). Equivalence of distance-based and rkhs-based statistics in hypothesis testing. *The annals of statistics*, pages 2263–2291.
- Serfling, R. J. (2009). *Approximation theorems of mathematical statistics*. John Wiley & Sons.
- Shekhar, S., Kim, I., and Ramdas, A. (2022). A permutation-free kernel two-sample test. *Advances in Neural Information Processing Systems*, 35:18168–18180.
- Shekhar, S., Kim, I., and Ramdas, A. (2023). A permutation-free kernel independence test. *Journal of Machine Learning Research*, 24(369):1–68.
- Smirnov, N. (1948). Table for estimating the goodness of fit of empirical distributions. *Annals of Mathematical Statistics*, 19(2):279–281.
- Soch, J. et al. (2025). Statproofbook/statproofbook.github.io: The book of statistical proofs (version 2024).
- Stelzer, J., Chen, Y., and Turner, R. (2013). Statistical inference and multiple testing correction in classification-based multi-voxel pattern analysis (mvpa): random permutations and cluster size control. *Neuroimage*, 65:69–82.
- Student (1908). The probable error of a mean. *Biometrika*, 6(1):1–25.
- Székel, G. J., Rizzo, M. L., et al. (2004). Testing for equal distributions in high dimension. *InterStat*, 5(16.10):1249–1272.
- Talagrand, M. (1995). Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathématiques de l’Institut des Hautes Etudes Scientifiques*, 81(1):73–205.
- Tibshirani, R. J. and Efron, B. (1993). An introduction to the bootstrap. *Monographs on statistics and applied probability*, 57(1):1–436.
- Tolstikhin, I. (2017). Concentration inequalities for samples without replacement. *Theory of Probability & Its Applications*, 61(3):462–481.
- Tsybakov, A. (2008). *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer New York.
- Van Der Vaart, A. W. and Wellner, J. A. (1996). Weak convergence. In *Weak convergence and empirical processes: with applications to statistics*, pages 16–28. Springer.
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.
- Villani, C. (2008). *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg.

- Wang, J., Boedihardjo, M., and Xie, Y. (2025). Statistical and computational guarantees of kernel max-sliced wasserstein distances. In *Forty-second International Conference on Machine Learning*.
- Wang, J., Gao, R., and Xie, Y. (2021a). Two-sample test using projected wasserstein distance. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 3320–3325. IEEE.
- Wang, J., Gao, R., and Xie, Y. (2021b). Two-sample test with kernel projected wasserstein distance. *arXiv preprint arXiv:2102.06449*.
- Wasserman, L. (2014). Minimax theory. Lecture notes, Carnegie Mellon University. Section 11: Minimax Lower Bounds.
- Weed, J. and Bach, F. (2019). Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *Bernoulli*, 25(4A):2620–2648.
- Wei, F. and Dudley, R. M. (2011). Dvoretzky–kiefer–wolfowitz inequalities for the two-sample case. *arXiv preprint arXiv:1107.5356*.
- Wendel, J. G. (1948). Note on the gamma function. *The American Mathematical Monthly*, 55(9):563.
- Xu, X. and Huang, Z. (2022). Central limit theorem for the sliced 1-wasserstein distance and the max-sliced 1-wasserstein distance. *arXiv preprint arXiv:2205.14624*.
- Zaremba, W., Gretton, A., and Blaschko, M. (2013). B-test: A non-parametric, low variance kernel two-sample test. *Advances in neural information processing systems*, 26.
- Zhang, J., Merikangas, K. R., Li, H., and Shou, H. (2021). Two-sample tests for repeated measurements of histogram objects with applications to wearable device data. *arXiv preprint arXiv:2106.12768*.

Supplementary Material

Minimax-Optimal Two-Sample Test with Sliced Wasserstein

OVERVIEW

This supplementary material contains additional discussions, detailed proofs, and an extra experiment that were omitted from the main text due to page constraints. The remaining sections are organized as follows.

- In Section 8, we recall several key concentration inequalities that serve as important tools for our proofs.
- In Section 9, we derive intermediate technical lemmas for permuted sliced Wasserstein statistics using results from optimal transport theory. A core contribution of this section is the derivation of a novel concentration inequality and an expectation bound for the permuted sliced Wasserstein statistic. These results are fundamental for controlling the random permutation threshold in our power analysis.
- In Section 10, we present a result characterizing the sample complexity of the sliced Wasserstein distance.
- Our main theoretical contributions are presented in Sections 11 and 12. These sections provide the complete proofs of our main theorems concerning the test’s performance (Theorems 4 and 5) and its minimax optimality (Propositions 6 and 7).
- In Section 13, we elaborate on why we chose a permutation-based framework for two-sample testing with the sliced Wasserstein distance, highlighting the intractability and practical challenges that prevent the use of the statistic’s asymptotic null distribution.
- Section 14 offers a technical review of existing methods for analyzing permutation tests. We explain why common techniques developed for U-statistics and other test statistics are not directly applicable to the sliced Wasserstein distance, thereby underscoring the novelty of our analytical approach.
- Finally, Section 15 provides results for the Gaussian mean shift experiment mentioned in Section 6.

Additional Notation. Throughout this supplementary material, we use an additional set of notation, described below.

- For $x \in \mathbb{R}$, $\lceil x \rceil$ denotes the smallest integer greater than or equal to x , and $\lfloor x \rfloor$ denotes the largest integer less than or equal to x .
- We write $d_{\text{TV}}(P, Q)$ for the total variation (TV) distance between P and Q , and $D_{\text{KL}}(P\|Q)$ for the Kullback–Leibler (KL) divergence. More details on those metrics can be found in, e.g., Tsybakov (2008, Section 2.4).
- For any integer $K \geq 1$, we denote by $[K]$ the set of the first K integers $\{1, \dots, K\}$.
- We use $\mathbb{P}_{\mu \times \nu}$ to denote probability with respect to $Y_1, \dots, Y_n \sim \mu$ and $Z_1, \dots, Z_m \sim \nu$. When additional randomness is present, we extend the notation accordingly, e.g., $\mathbb{P}_{\mu \times \nu \times r \times \sigma}$.
- Finally, given two random variables U and V , we write $U \stackrel{(d)}{=} V$ to denote that U and V have the same distribution.

8 SOME USEFUL CONCENTRATION INEQUALITIES

We begin this section by recalling two classical concentration inequalities: McDiarmid's inequality (McDiarmid et al., 1989) and Hoeffding's inequality (Hoeffding, 1963).

Lemma 8 (McDiarmid's Inequality) *Let $N \geq 1$. Let $X_1, \dots, X_N$ be independent random variables taking values in a set $\mathcal{X}$, and let $f : \mathcal{X}^N \rightarrow \mathbb{R}$. Assume that f satisfies the bounded difference property, i.e., there exist constants $c_1, \dots, c_N \geq 0$ such that for all $i \in \{1, \dots, N\}$ and for all $x_1, \dots, x_N, x'_i \in \mathcal{X}$,*

$$|f(x_1, \dots, x_i, \dots, x_N) - f(x_1, \dots, x'_i, \dots, x_N)| \leq c_i.$$

Then, for any $t > 0$,

$$\mathbb{P}(|f(X_1, \dots, X_N) - \mathbb{E}[f(X_1, \dots, X_N)]| \geq t) \leq 2 \exp\left(-\frac{2t^2}{\sum_{i=1}^N c_i^2}\right).$$

Lemma 9 (Hoeffding's Inequality) *Let $N \geq 1$. Let $X_1, \dots, X_N$ be independent random variables such that $a_i \leq X_i \leq b_i$ for $i = 1, \dots, N$. Then, for any $t > 0$,*

$$\mathbb{P}\left[\left|\frac{1}{N} \sum_{i=1}^N X_i - \mathbb{E}\left(\frac{1}{N} \sum_{i=1}^N X_i\right)\right| \geq t\right] \leq 2 \exp\left(-\frac{2N^2 t^2}{\sum_{i=1}^N (b_i - a_i)^2}\right).$$

We shall also rely on a permutation-based version of the classical McDiarmid's inequality to control the empirical quantile $\widehat{c}_{1-\alpha, N}^B$ introduced in Algorithm 1. A statement of this result can be found in Tolstikhin (2017, Theorem 6) and El-Yaniv and Pechyony (2009, Lemma 2). To prepare for the statement of that result, we first introduce the following definition.

Definition 10 ((n, m)-symmetric function) *Let $N \geq 1$ and let n be a positive integer such that $n < N$. Set $m := N - n$. A function $f : S_N \rightarrow \mathbb{R}$, defined on the symmetric group over $\{1, \dots, N\}$, is called (n, m)-symmetric if it remains invariant under the change of order of the first n coordinates and/or last m coordinates of any permutation $\pi \in S_N$.*

We are now ready to state a permutation-based version of McDiarmid's inequality.

Lemma 11 (Permutation McDiarmid's inequality) *Let n and N be positive integers such that $n < N$. Let π be a random permutation drawn uniformly from the symmetric group over $\{1, \dots, N\}$, and let $f(\pi)$ be an ($n, N - n$)-symmetric function such that there exists a constant $b > 0$ satisfying*

$$|f(\pi) - f(\pi^{i,j})| \leq b,$$

for all $\pi, i \in \{1, \dots, n\}, j \in \{n+1, \dots, N\}$, where $\pi^{i,j}$ is the permutation obtained by transposing the i -th and j -th entries of π . Then, for any $\epsilon > 0$,

$$\mathbb{P}_\pi(f(\pi) - \mathbb{E}_\pi[f(\pi)] \geq \epsilon) \leq \exp\left(-\frac{2\epsilon^2}{nb^2} \cdot \frac{N - \frac{1}{2}}{N - n} \cdot \left(1 - \frac{1}{2 \max\{n, N - n\}}\right)\right).$$

9 PROPERTIES OF THE PERMUTED SLICED WASSERSTEIN STATISTIC

Our main objective in this section is to obtain a deterministic upper bound the (random) permuted test statistics $\widehat{\text{SW}}_p^{p, \pi}$ introduced in (4). To this end, we exploit the properties of the sliced Wasserstein distance under permutation via a permutation-based version of McDiarmid's inequality (recalled in Lemma 11). It yields a high-probability bound relating the permuted statistics to its expectation. We then control the expectation term by applying the optimal matching bound from Bobkov and Ledoux (2021, Corollary 5).

We recall that all results are established under the assumption that the samples $\mathcal{Y}_n = (y_1, \dots, y_n)$ and $\mathcal{Z}_m = (z_1, \dots, z_m)$ satisfy $\|y_i\| \leq D$ and $\|z_j\| \leq D$ for all $i = 1, \dots, n$ and $j = 1, \dots, m$. Moreover, $\Theta = (\theta_1, \dots, \theta_L)$ denotes the projection directions drawn from the uniform distribution σ on the unit sphere $\mathbb{S}^{d-1}$.

We first establish a high-probability concentration inequality for the permuted empirical sliced Wasserstein distance $\widehat{\text{SW}}_p^{p, \pi}$ introduced in (4). It is the first step towards proving our main theorem in Section 11.2.

Lemma 12 Let $p \geq 1$. Let π be a permutation drawn uniformly from the symmetric group on $\{1, \dots, n+m\}$, and let $\pi^{i,j}$ denote the permutation obtained from π by exchanging its i -th and j -th elements, as defined in Lemma 11. Then, for any $\epsilon > 0$,

$$\mathbb{P}_\pi \left(\widehat{\text{SW}}_p^{p,\pi} - \mathbb{E}_\pi \left[\widehat{\text{SW}}_p^{p,\pi} \right] \geq \epsilon \mid \mathcal{Y}_n, \mathcal{Z}_m, \Theta \right) \leq \exp \left(- \frac{2\epsilon^2}{n \left(\frac{(2D)^p}{n} + \frac{(2D)^p}{m} \right)^2} \cdot \frac{n+m-\frac{1}{2}}{m} \left(1 - \frac{1}{2 \max\{m, n\}} \right) \right).$$

Proof Given the samples $\mathcal{Y}_n, \mathcal{Z}_m$ and a set of projection directions $\Theta = (\theta_1, \dots, \theta_L)$ on the unit sphere $\mathbb{S}^{d-1}$, for brevity, we set

$$f(\pi) := \widehat{\text{SW}}_p^{p,\pi}, \quad \pi \in S_{n+m},$$

where $\widehat{\text{SW}}_p^{p,\pi}$ is the permuted empirical quantity introduced in (4).

By construction, f is (n, m) -symmetric. In order to apply Lemma 11 we need to get an upper bound on the sensitivity of f to transpositions involving elements from opposite halves of the permutation.

Without loss of generality, suppose that $\pi = (1, 2, 3, \dots, n+m)$ and $\pi^{i,j} = (1, 2, \dots, i-1, j, i+1, \dots, n, \dots, j-1, i, j+1, \dots, n+m)$ for $1 \leq i \leq n$ and $n+1 \leq j \leq n+m$.

Let $\ell \in \{1, \dots, L\}$. Denote the projected samples on the direction θ_ℓ as

$$\tilde{\mathcal{X}} := (\tilde{x}_1, \dots, \tilde{x}_n, \tilde{x}_{n+1}, \dots, \tilde{x}_{n+m}) = (\langle y_1, \theta_\ell \rangle, \dots, \langle y_n, \theta_\ell \rangle, \langle z_1, \theta_\ell \rangle, \dots, \langle z_m, \theta_\ell \rangle),$$

and accordingly,

$$\tilde{Y} := (\tilde{x}_1, \dots, \tilde{x}_n), \quad \tilde{Z} := (\tilde{x}_{n+1}, \dots, \tilde{x}_{n+m}).$$

By definition of the p -Wasserstein distance and the permutation π , we have

$$\begin{aligned} W_p^p \left(\Pi_{\#}^{\theta_\ell} \hat{\mu}_n^\pi, \Pi_{\#}^{\theta_\ell} \hat{\nu}_m^\pi \right) &= \min_{\substack{\gamma \geq 0, \\ \sum_k \gamma_{kl} = \frac{1}{m}, \\ \sum_l \gamma_{kl} = \frac{1}{n}}} \left(\sum_{k=1}^n \sum_{l=1}^m \gamma_{kl} |\tilde{x}_{\pi(k)} - \tilde{x}_{\pi(n+l)}|^p \right) \\ &= \min_{\substack{\gamma \geq 0, \\ \sum_k \gamma_{kl} = \frac{1}{m}, \\ \sum_l \gamma_{kl} = \frac{1}{n}}} \left(\sum_{k=1}^n \sum_{l=1}^m \gamma_{kl} |\tilde{y}_k - \tilde{z}_l|^p \right) \\ &= \sum_{k,l} \gamma_{kl}^* |\tilde{y}_k - \tilde{z}_l|^p, \end{aligned}$$

where γ^* denotes the minimizer of the above Monge problem.

We now consider the permutation $\pi^{i,j}$. Noting that γ^* is a valid coupling, we have

$$\begin{aligned} W_p^p \left(\Pi_{\#}^{\theta_\ell} \hat{\mu}_n^{\pi^{i,j}}, \Pi_{\#}^{\theta_\ell} \hat{\nu}_m^{\pi^{i,j}} \right) &= \min_{\substack{\gamma \geq 0, \\ \sum_k \gamma_{kl} = \frac{1}{m}, \\ \sum_l \gamma_{kl} = \frac{1}{n}}} \left(\sum_{k=1}^n \sum_{l=1}^m \gamma_{kl} |\tilde{x}_{\pi^{i,j}(k)} - \tilde{x}_{\pi^{i,j}(n+l)}|^p \right) \\ &\leq \sum_{k=1}^n \sum_{l=1}^m \gamma_{kl}^* |\tilde{x}_{\pi^{i,j}(k)} - \tilde{x}_{\pi^{i,j}(n+l)}|^p \\ &\leq \sum_{k=1}^n \sum_{l=1}^m \gamma_{kl}^* |\tilde{y}_k - \tilde{z}_l|^p + \sum_{l=1}^m \gamma_{ik}^* |\tilde{x}_{\pi^{i,j}(i)} - \tilde{x}_{\pi^{i,j}(n+l)}|^p + \sum_{k=1}^n \gamma_{kj}^* |\tilde{x}_{\pi^{i,j}(k)} - \tilde{x}_{\pi^{i,j}(j)}|^p \\ &\leq W_p^p \left(\Pi_{\#}^{\theta_\ell} \hat{\mu}_n^\pi, \Pi_{\#}^{\theta_\ell} \hat{\nu}_m^\pi \right) + \frac{(2D)^p}{n} + \frac{(2D)^p}{m}, \end{aligned}$$

where the last inequality follows from the triangle inequality and the fact that every projected sample satisfies

$$|\tilde{x}_h| = |\langle x_h, \theta_\ell \rangle| \leq \|x_h\| \|\theta_\ell\| \leq D,$$

for $1 \leq h \leq n + m$. As a result, we obtain

$$f(\pi) - f(\pi^{i,j}) \leq \frac{(2D)^p}{n} + \frac{(2D)^p}{m}.$$

Following the same steps, we can prove that

$$f(\pi^{i,j}) - f(\pi) \leq \frac{(2D)^p}{n} + \frac{(2D)^p}{m},$$

which is sufficient to conclude that

$$|f(\pi) - f(\pi^{i,j})| \leq \frac{(2D)^p}{n} + \frac{(2D)^p}{m}.$$

Applying Lemma 11 with $b = \frac{(2D)^p}{n} + \frac{(2D)^p}{m}$ and $f(\pi) = \widehat{\text{SW}}_p^{p,\pi}$ completes the proof. $\blacksquare$

Having successfully related the permuted test statistic $\widehat{\text{SW}}_p^{p,\pi}$ (introduced in (4)) to its expectation, we now turn to controlling this expectation term. To proceed, we next present several auxiliary lemmas that will be used later in the proof.

Our argument begins with the following result, stated by Bobkov and Ledoux (2021, Corollary 5).

Lemma 13 (Corollary 5 in Bobkov and Ledoux (2021)) *Let n and N be integers such that $1 \leq n \leq N$. Let $x_1, \dots, x_N \in [0, 1]^d$. Denote by $\mathcal{G}_n$ the collection of all subsets $\tau \subset \{1, \dots, N\}$ of cardinality $|\tau| = n$, equipped with the uniform probability measure π_n . With every $\tau \in \mathcal{G}_n$, we associate the empirical measure*

$$\mu_\tau := \frac{1}{n} \sum_{j \in \tau} \delta_{x_j}, \quad \tilde{\mu} := \frac{1}{N} \sum_{j=1}^N \delta_{x_j}.$$

Then the empirical measures μ_τ satisfy

$$\mathbb{E}_{\pi_n} [W_1(\mu_\tau, \tilde{\mu})] \leq \begin{cases} \sqrt{\frac{2}{n}} & \text{if } d = 1, \\ 8 \sqrt{\frac{1 + \log(2n)}{n}} & \text{if } d = 2, \\ \frac{13\sqrt{d}}{n^{1/d}} & \text{if } d \geq 3. \end{cases}$$

Remark 14 *The result of Corollary 5 in Bobkov and Ledoux (2021) is stated for empirical measures supported on the unit cube $[0, 1]^d$. We now extend this result to the case where the data points lie in a general cube $[-D, D]^d$ for some $D > 0$. To relate these two settings, consider the affine transformation*

$$T: [0, 1]^d \rightarrow [-D, D]^d, \quad T(x) = 2D(x - \mathbf{1}_d),$$

where $\mathbf{1}_d := (1, 1, \dots, 1) \in \mathbb{R}^d$. Given data points $x_1, \dots, x_N \in [0, 1]^d$, we define

$$x'_i = T(x_i), \quad i = 1, \dots, N,$$

so that each x'_i lies in the cube $[-D, D]^d$.

Let μ'_τ and $\tilde{\mu}'$ denote the empirical measures defined analogously to μ_τ and $\tilde{\mu}$ in Lemma 13 but based on points $(x'_i)_{i=1}^N$, that is

$$\mu'_\tau = \frac{1}{n} \sum_{j \in \tau} \delta_{x'_j}, \quad \tilde{\mu}' = \frac{1}{N} \sum_{j=1}^N \delta_{x'_j}.$$

By construction, we have $T_{\#}\mu_{\tau} = \mu'_{\tau}$ and $T_{\#}\tilde{\mu} = \tilde{\mu}'$. Moreover, for any $x, y \in [0, 1]^d$, we have

$$\|T(x) - T(y)\|_2 = 2D\|x - y\|_2.$$

Consequently, we have the scaling relation

$$W_1(\mu'_{\tau}, \tilde{\mu}') = 2D W_1(\mu_{\tau}, \tilde{\mu}).$$

Applying Corollary 5 in Bobkov and Ledoux (2021) to μ_{τ} and $\tilde{\mu}$ then yields

$$\mathbb{E}_{\pi_n}[W_1(\mu'_{\tau}, \tilde{\mu}')] = \mathbb{E}_{\pi_n}[(2D)W_1(\mu_{\tau}, \tilde{\mu})] \leq \begin{cases} 2D\sqrt{\frac{2}{n}} & \text{if } d = 1, \\ 16D\sqrt{\frac{1 + \log(2n)}{n}} & \text{if } d = 2, \\ \frac{26D\sqrt{d}}{n^{1/d}} & \text{if } d \geq 3. \end{cases}$$

Lemma 15 (Relation between 1-Wasserstein and p -Wasserstein distances) *Let P, Q , and M be three probability measures supported on the compact set $[-D, D]^d \subset \mathbb{R}^d$. Let $p \geq 1$, then the following inequality holds:*

$$W_p^p(P, Q) \leq (4D\sqrt{d})^{p-1} (W_1(P, Q) + W_1(Q, M)).$$

Proof By applying the triangle inequality for the p -Wasserstein distance, we get

$$W_p(P, Q) \leq W_p(P, M) + W_p(Q, M).$$

Using the inequality $(a + b)^p \leq 2^{p-1}(a^p + b^p)$, we obtain

$$W_p^p(P, Q) \leq 2^{p-1} (W_p^p(P, M) + W_p^p(Q, M)).$$

For $x, y \in [-D, D]^d$, we have

$$\|x - y\|_2 \leq 2D\sqrt{d}.$$

Moreover, for $p \geq 1$, we have

$$\|x - y\|_2^p \leq (2D\sqrt{d})^{p-1} \|x - y\|_2.$$

Hence, since all three measures are supported on $[-D, D]^d$, by the definition of Wasserstein distances, we obtain

$$W_p^p(P, M) \leq (2D\sqrt{d})^{p-1} W_1(P, M), \quad W_p^p(Q, M) \leq (2D\sqrt{d})^{p-1} W_1(Q, M).$$

Combining all the inequalities gives the stated result. ■

Remark 16 *We recall that our standing assumption is that all samples $\mathcal{Y}_n = (y_1, \dots, y_n)$ and $\mathcal{Z}_m = (z_1, \dots, z_m)$ satisfy $\|y_i\| \leq D$ and $\|z_j\| \leq D$ for all $i = 1, \dots, n$ and $j = 1, \dots, m$. Meanwhile, the results in Lemma 13 and Lemma 15 are stated for the case where all data points lie in the cube $[-D, D]^d$. Since the Euclidean ball $\{x \in \mathbb{R}^d : \|x\| \leq D\}$ is contained in this cube, those results remain valid in our setting.*

Using the results of Lemma 13 and Lemma 15, along with their accompanying remarks, the following proposition establishes an upper bound for the expectation term conditional on the given samples and the projection directions.

Proposition 17 *Let n and m be positive integers such that $n \leq m$, $p \geq 1$ and set $N := n + m$. Given a collection of points $x_1, \dots, x_N$ with $\|x_i\| \leq D$ for all $1 \leq i \leq N$, and a set of projection directions $\Theta = (\theta_1, \dots, \theta_L)$ on the unit sphere $\mathbb{S}^{d-1}$. Let π be a random permutation drawn uniformly from the symmetric group over $\{1, \dots, N\}$. Then, with $p \geq 1$, we have*

$$\mathbb{E}_{\pi} \left[\widehat{\text{SW}}_p^{p, \pi} \right] = \frac{1}{L} \sum_{\ell=1}^L \mathbb{E}_{\pi} \left[W_p^p \left(\Pi_{\#}^{\theta_{\ell}} \hat{\mu}_n^{\pi}, \Pi_{\#}^{\theta_{\ell}} \hat{\nu}_m^{\pi} \right) \right] \leq (4D)^p \frac{\sqrt{2}}{\sqrt{n}}.$$

Here, the expectation $\mathbb{E}_{\pi}$ is taken with respect to the random permutation π , conditional on the samples and the projection directions.

Proof Fix a projection direction θ_ℓ , $1 \leq \ell \leq L$. We want to upper bound $\mathbb{E}_\pi \left[W_p^p \left(\Pi_{\#}^{\theta_\ell} \hat{\mu}_n^\pi, \Pi_{\#}^{\theta_\ell} \hat{\nu}_m^\pi \right) \right]$.

We introduce the empirical measure

$$\tilde{\mu} := \frac{1}{N} \sum_{j=1}^N \delta_{\langle x_j, \theta_\ell \rangle}.$$

Note that each projected sample satisfies $|\langle x_j, \theta_\ell \rangle| \leq D$ for all $1 \leq j \leq N$ by the Cauchy-Schwarz inequality.

Adapting the notations from Lemma 13, we first define $\mathcal{G}_n$ and $\mathcal{G}_m$ as the collections of all subsets of $\{1, \dots, n+m\}$ of cardinalities n and m , respectively. We denote by π_n and π_m the uniform distributions on $\mathcal{G}_n$ and $\mathcal{G}_m$. To every $\tau^n \in \mathcal{G}_n$ and $\tau^m \in \mathcal{G}_m$, we associate two empirical measures

$$\mu_{\tau^n} := \frac{1}{n} \sum_{j \in \tau^n} \delta_{\langle \theta_\ell, x_j \rangle}, \quad \mu_{\tau^m} := \frac{1}{m} \sum_{j \in \tau^m} \delta_{\langle \theta_\ell, x_j \rangle}.$$

It follows that

$$\begin{aligned} \mathbb{E}_\pi \left[W_p^p \left(\Pi_{\#}^{\theta_\ell} \hat{\mu}_n^\pi, \Pi_{\#}^{\theta_\ell} \hat{\nu}_m^\pi \right) \right] &= \frac{1}{(n+m)!} \sum_{\pi \in S_{n+m}} W_p^p \left(\frac{1}{n} \sum_{i=1}^n \delta_{\langle x_{\pi(i)}, \theta_\ell \rangle}, \frac{1}{m} \sum_{i=1}^m \delta_{\langle x_{\pi(n+i)}, \theta_\ell \rangle} \right) \\ &= \frac{m!n!}{(n+m)!} \sum_{K \subset \{x_1, \dots, x_{n+m}\}, |K|=n} W_p^p \left(\frac{1}{n} \sum_{x \in K} \delta_{\langle x, \theta_\ell \rangle}, \frac{1}{m} \sum_{x \in K^c} \delta_{\langle x, \theta_\ell \rangle} \right) \\ &\stackrel{(i)}{\leq} \frac{(4D)^{p-1}}{\binom{n+m}{n}} \sum_{K \subset \{x_1, \dots, x_{n+m}\}, |K|=n} \left[W_1 \left(\frac{1}{n} \sum_{x \in K} \delta_{\langle x, \theta_\ell \rangle}, \tilde{\mu} \right) + W_1 \left(\frac{1}{m} \sum_{x \in K^c} \delta_{\langle x, \theta_\ell \rangle}, \tilde{\mu} \right) \right] \\ &= (4D)^{p-1} [\mathbb{E}_{\pi_m} (W_1(\mu_{\tau^m}, \tilde{\mu})) + \mathbb{E}_{\pi_n} (W_1(\mu_{\tau^n}, \tilde{\mu}))] \\ &\stackrel{(ii)}{\leq} 2D \cdot (4D)^{p-1} \left(\frac{\sqrt{2}}{\sqrt{n}} + \frac{\sqrt{2}}{\sqrt{m}} \right) \\ &\stackrel{(iii)}{\leq} 4D \cdot (4D)^{p-1} \frac{\sqrt{2}}{\sqrt{n}}, \end{aligned}$$

where (i) follows from Lemma 15, and (ii) and (iii) follow from Lemma 13 and from the assumption that $n \leq m$, respectively.

Finally, by linearity of the expectation,

$$\mathbb{E}_\pi \left[\widehat{SW}_p^{p, \pi} \right] = \frac{1}{L} \sum_{\ell=1}^L \mathbb{E}_\pi \left[W_p^p \left((\Pi^{\theta_\ell})_{\#} \hat{\mu}_n^\pi, (\Pi^{\theta_\ell})_{\#} \hat{\nu}_m^\pi \right) \right] \leq (4D)^p \frac{\sqrt{2}}{\sqrt{n}},$$

and the proof is concluded. ■

10 SLICED WASSERSTEIN SAMPLE COMPLEXITY

Let μ and ν be probability distributions on $\mathbb{R}^d$ whose supports lie in a common ball centered at the origin with radius D . Their corresponding empirical measures are denoted by $\hat{\mu}_n$ and $\hat{\nu}_m$. The quantities $SW_p^p(\mu, \nu)$ and $\widehat{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m)$ are defined in Eqs. (1) and (2), respectively.

To the best of our knowledge, recent works have established bounds on quantities such as $\mathbb{E} [SW_p(\hat{\mu}_n, \mu)]$ (see, e.g., Nietert et al. (2022, Theorem 1)), $\mathbb{E} [SW_p^p(\hat{\mu}_n, \hat{\nu}_m) - SW_p^p(\mu, \nu)]$ (see, e.g., Ohana et al. (2023, Theorem 2)), and $\mathbb{E} [|\widehat{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m) - SW_p^p(\mu, \nu)|]$ (see, e.g., Nietert et al. (2022, Proposition 5)), as well as high-probability bounds for $|SW_p^p(\mu, \nu) - \widehat{SW}_p^p(\mu, \nu)|$ (see, e.g., Xu and Huang (2022, Proposition 4)). These results, however, do not exactly yield the type of bound required for the subsequent step—specifically, in the proof of Lemma 19.

Under the assumption that both distributions μ and ν are supported on a common bounded ball, we derive the following high-probability bound for $|\widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m) - \text{SW}_p^p(\mu, \nu)|$ by applying standard concentration inequalities—namely McDiarmid’s inequality and Hoeffding’s inequality (recalled in Lemma 8 and Lemma 9, respectively). The resulting statement is presented in the following lemma.

Lemma 18 *Let $\beta \in (0, 1)$ and $p \geq 1$. Suppose that μ and ν are probability distributions supported on the centered ball of radius $D > 0$ in $\mathbb{R}^d$. We consider two independent samples drawn from μ and ν :*

$$\mathcal{Y}_n := (Y_1, \dots, Y_n) \stackrel{i.i.d.}{\sim} \mu, \quad \mathcal{Z}_m := (Z_1, \dots, Z_m) \stackrel{i.i.d.}{\sim} \nu,$$

with an assumption that $n \leq m$. Let $\Theta := (\theta_1, \dots, \theta_L)$ denote L i.i.d. projection directions drawn from the uniform distribution σ on the unit sphere $\mathbb{S}^{d-1}$. Then, with probability at least $1 - \frac{\beta}{2}$, the following inequality holds:

$$|\widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m) - \text{SW}_p^p(\mu, \nu)| < (2D)^p \left(\sqrt{\frac{\log(8/\beta)}{2L}} + \sqrt{\frac{\log(8/\beta)}{n}} \right).$$

Proof We decompose the total error into two contributions:

$$|\widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m) - \text{SW}_p^p(\mu, \nu)| \leq A + B,$$

where

$$A = |\widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m) - \text{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m)|, \quad B = |\text{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m) - \text{SW}_p^p(\mu, \nu)|.$$

Control of A. Given the datasets $\mathcal{Y}_n, \mathcal{Z}_m$, the Monte Carlo estimator $\widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m)$ is an average of L i.i.d. bounded random variables

$$R_\ell = \text{W}_p^p(\Pi_{\#}^{\theta_\ell} \hat{\mu}_n, \Pi_{\#}^{\theta_\ell} \hat{\nu}_m), \quad \ell = 1, \dots, L.$$

Since $\|\theta_\ell\|_2 = 1$ and the supports of μ and ν are contained in the centered ball with radius D , the push-forward measures $(\Pi^{\theta_\ell})_{\#} \hat{\mu}_n$ and $(\Pi^{\theta_\ell})_{\#} \hat{\nu}_m$ are supported on the interval $[-D, D]$. Consequently,

$$0 \leq R_\ell \leq (2D)^p, \quad \forall 1 \leq \ell \leq L.$$

Then, Hoeffding’s inequality (recalled in Lemma 9) yields that, for any $t_A > 0$,

$$\mathbb{P}_\sigma(A \geq t_A \mid \mathcal{Y}_n, \mathcal{Z}_m) \leq 2 \exp\left(-\frac{2Lt_A^2}{(2D)^{2p}}\right).$$

Since the right-hand side does not depend on the samples, the same bound holds unconditionally, namely

$$\mathbb{P}_{\mu, \nu, \sigma}(A \geq t_A) \leq 2 \exp\left(-\frac{2Lt_A^2}{(2D)^{2p}}\right).$$

Control of B. Fix the projection directions $\theta_1, \dots, \theta_L$.

The function $(y_1, \dots, y_n, z_1, \dots, z_m) \mapsto \text{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m)$ satisfies the bounded difference property, whose definition is recalled in Lemma 8. Indeed, changing a single observation y_i modifies $\hat{\mu}_n$ by at most $1/n$ of its mass. Since all points are supported in a set of diameter $2D$, the value of $\text{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m)$ changes by at most $(2D)^p/n$. Similarly, replacing one sample z_j alters $\hat{\nu}_m$ by $1/m$ of its mass, and hence $\text{SW}_p^p(\hat{\mu}_n, \hat{\nu}_m)$ changes by at most $(2D)^p/m$. McDiarmid’s inequality (recalled in Lemma 8) guarantees that, for any $t_B > 0$,

$$\mathbb{P}_{\mu, \nu}(B \geq t_B \mid \theta_1, \dots, \theta_L) \leq 2 \exp\left(-\frac{2t_B^2}{(2D)^{2p}(1/n + 1/m)}\right).$$

Since the right-hand side is independent of the projection directions, the bound also holds without conditioning, that is,

$$\mathbb{P}_{\mu, \nu, \sigma}(B \geq t_B) \leq 2 \exp\left(-\frac{2t_B^2}{(2D)^{2p}(1/n + 1/m)}\right).$$

Moreover, since $n \leq m$, then

$$\mathbb{P}_{\mu, \nu, \sigma}(B \geq t_B) \leq 2 \exp\left(-\frac{nt_B^2}{(2D)^{2p}}\right).$$

A simple union bound gives

$$\mathbb{P}_{\mu, \nu, \sigma}(A + B \geq t_A + t_B) \leq \mathbb{P}_{\mu, \nu, \sigma}(A \geq t_A) + \mathbb{P}_{\mu, \nu, \sigma}(B \geq t_B).$$

Hence,

$$\mathbb{P}_{\mu, \nu, \sigma}\left(\left|\widehat{\text{SW}}_p^p(\widehat{\mu}_n, \widehat{\nu}_m) - \text{SW}_p^p(\mu, \nu)\right| \geq t_A + t_B\right) \leq 2 \exp\left(-\frac{2Lt_A^2}{(2D)^{2p}}\right) + 2 \exp\left(-\frac{nt_B^2}{(2D)^{2p}}\right).$$

We obtain the stated result by setting

$$t_A = (2D)^p \sqrt{\frac{\log(8/\beta)}{2L}} \quad \text{and} \quad t_B = (2D)^p \sqrt{\frac{\log(8/\beta)}{n}}.$$

■

11 LEVEL AND POWER GUARANTEES FOR ALGORITHM 1

11.1 Proof of Theorem 4

Following the proof of Schrab et al. (2023, Proposition 1), we obtain the following chain of implications

$$\begin{aligned} \Delta(\mathcal{Y}_n, \mathcal{Z}_m, \mathbb{Z}_B, \Theta) = 1 &\Rightarrow \widehat{\text{SW}}_p^p > \widehat{c}_{1-\alpha, N}^B \\ &\Rightarrow \widehat{\text{SW}}_p^p > \widehat{\text{SW}}_p^{p, r \cdot \lceil (B+1)(1-\alpha) \rceil} \\ &\Rightarrow \sum_{b=1}^{B+1} \mathbf{1}\left(\widehat{\text{SW}}_p^{p, \pi_b} < \widehat{\text{SW}}_p^{p, \pi_{B+1}}\right) \geq \lceil (B+1)(1-\alpha) \rceil \\ &\Rightarrow B+1 - \sum_{b=1}^{B+1} \mathbf{1}\left(\widehat{\text{SW}}_p^{p, \pi_b} < \widehat{\text{SW}}_p^{p, \pi_{B+1}}\right) \leq B+1 - \lceil (B+1)(1-\alpha) \rceil \\ &\Rightarrow \sum_{b=1}^{B+1} \mathbf{1}\left(\widehat{\text{SW}}_p^{p, \pi_b} \geq \widehat{\text{SW}}_p^{p, \pi_{B+1}}\right) \leq \lfloor \alpha(B+1) \rfloor \\ &\Rightarrow \sum_{b=1}^{B+1} \mathbf{1}\left(\widehat{\text{SW}}_p^{p, \pi_b} \geq \widehat{\text{SW}}_p^{p, \pi_{B+1}}\right) \leq \alpha(B+1) \\ &\Rightarrow \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbf{1}\left(\widehat{\text{SW}}_p^{p, \pi_b} \geq \widehat{\text{SW}}_p^{p, \pi_{B+1}}\right)\right] \leq \alpha, \end{aligned}$$

where the fifth deduction follows from the fact that $B+1 - \lceil (1-\alpha)(B+1) \rceil = \lfloor \alpha(B+1) \rfloor$. Besides, the notations $\widehat{\text{SW}}_p^p$ and $\widehat{\text{SW}}_p^{p, \pi}$ refer respectively to the statistics defined in (2) and (4).

Moreover, Lemma 1 in Romano and Wolf (2005) guarantees that:

$$\mathbb{P}_{\mu \times \mu \times r \times \sigma} \left\{ \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbf{1}\left(\widehat{\text{SW}}_p^{p, \pi_b} \geq \widehat{\text{SW}}_p^{p, \pi_{B+1}}\right)\right] \leq \alpha \right\} \leq \alpha.$$

As a consequence, we have

$$\mathbb{P}_{\mu \times \mu \times r \times \sigma} (\Delta(\mathcal{Y}_n, \mathcal{Z}_m, \mathbb{Z}_B, \Theta) = 1) \leq \mathbb{P}_{\mu \times \mu \times r \times \sigma} \left\{ \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbf{1}\left(\widehat{\text{SW}}_p^{p, \pi_b} \geq \widehat{\text{SW}}_p^{p, \pi_{B+1}}\right)\right] \leq \alpha \right\} \leq \alpha.$$

Hence, the test controls the Type I error at level α .

11.2 Proof of Theorem 5

Let r denote the uniform distribution over the symmetric group S_{n+m} , i.e., the set of all permutations of $\{1, \dots, n+m\}$. Let π be a permutation drawn from r . Moreover, we defined $c_{1-\alpha, N}$ as the $(1-\alpha)$ -quantile (with respect to the randomness of π) of the test statistics $\widehat{\text{SW}}_p^{p, \pi}$ (introduced in Eq. (4)), that is,

$$c_{1-\alpha, N} = \inf \left\{ t : \mathbb{P}_\pi \left(\widehat{\text{SW}}_p^{p, \pi} \geq t \mid \mathcal{Y}_n, \mathcal{Z}_m, \Theta \right) \leq \alpha \right\}. \quad (6)$$

As explained in Remark 2, in practice we draw B independent permutations $(\pi_b)_{1 \leq b \leq B}$ from the uniform distribution r on S_{n+m} and set $\pi_{B+1} := \text{id}$ to denote the identity permutation. We then compute the corresponding statistics $\widehat{\text{SW}}_p^{p, \pi_b}$ for $b = 1, \dots, B+1$ and estimate the $(1-\alpha)$ empirical quantile of the permutation distribution by

$$\widehat{c}_{1-\alpha, N}^B := \inf \left\{ t : \frac{1}{B+1} \sum_{i=1}^{B+1} \mathbf{1}\{\widehat{\text{SW}}_p^{p, \pi_i} \leq t\} \geq 1 - \alpha \right\}.$$

In accordance with the convention introduced earlier, we use the notation $\mathbb{P}_r$ to denote the probability with respect to the B random permutations $\pi_1, \dots, \pi_B \stackrel{\text{i.i.d.}}{\sim} r$.

With these preliminaries in place, the proof of Theorem 5 proceeds as follows. Taking inspiration from Schrab et al. (2023, Lemma 4), we first provide a condition on the separation between the distributions μ and ν that guarantees a desired level of test power (see Lemma 19). This condition relates the separation to the (random) empirical quantile $\widehat{c}_{1-\alpha, N}^B$. Next, we control $\widehat{c}_{1-\alpha, N}^B$ by linking it to the (deterministic) quantile $c_{1-\alpha, N}$, conditional on the samples and the projection directions (see Lemma 20). Then, the quantile $c_{1-\alpha, N}$ is bounded using the auxiliary results prepared in Section 9 (see Proposition 21). Finally, we derive the bound for $\widehat{c}_{1-\alpha, N}^B$ stated in Proposition 22. We now develop this program in detail.

Main ingredients for the proof

We begin with the following lemma, which is a straightforward adaptation of Schrab et al. (2023, Lemma 4). It provides a sufficient condition on the separation between distributions μ and ν to ensure a desired level of test power.

Lemma 19 *Let $\beta \in (0, 1)$ and $p \geq 1$. Consider the setting of Lemma 18. Let*

$$\gamma(n, p, \beta, D, L) := (2D)^p \left(\sqrt{\frac{\log(8/\beta)}{2L}} + \sqrt{\frac{\log(8/\beta)}{n}} \right).$$

The test defined in Algorithm 1 achieves power at least $1 - \beta$ provided that

$$\mathbb{P}_{\mu \times \nu \times r \times \sigma} \left(\text{SW}_p^p(\mu, \nu) \geq \gamma(n, p, \beta, D, L) + \widehat{c}_{1-\alpha, N}^B \right) > 1 - \frac{\beta}{2}. \quad (7)$$

Proof Lemma 18 guarantees

$$\mathbb{P}_{\mu \times \nu \times r \times \sigma} \left(\left| \widehat{\text{SW}}_p^p(\widehat{\mu}_n, \widehat{\nu}_m) - \text{SW}_p^p(\mu, \nu) \right| \geq \gamma(n, p, \beta, D, L) \right) \leq \frac{\beta}{2}, \quad (8)$$

where

$$\gamma(n, p, \beta, D, L) := (2D)^p \left(\sqrt{\frac{\log(8/\beta)}{2L}} + \sqrt{\frac{\log(8/\beta)}{n}} \right).$$

Define the events

$$\mathcal{A} := \{ \widehat{\text{SW}}_p^p(\widehat{\mu}_n, \widehat{\nu}_m) \leq \widehat{c}_{1-\alpha, N}^B \} \quad \text{and} \quad \mathcal{B} := \{ \text{SW}_p^p(\mu, \nu) \geq \gamma(n, p, \beta, D, L) + \widehat{c}_{1-\alpha, N}^B \}.$$

Let us show that $\mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{A}) \leq \beta$. By definition of the events and Eq. (8), we have

$$\begin{aligned} \mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{A} \cap \mathcal{B}) &= \mathbb{P}_{\mu \times \nu \times r \times \sigma} \left(\widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m) \leq \hat{c}_{1-\alpha, N}^B, \hat{c}_{1-\alpha, N}^B \leq \text{SW}_p^p(\mu, \nu) - \gamma(n, p, \beta, D, L) \right) \\ &\leq \mathbb{P}_{\mu \times \nu \times r \times \sigma} \left(\widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m) - \text{SW}_p^p(\mu, \nu) \leq -\gamma(n, p, \beta, D, L) \right) \\ &\leq \mathbb{P}_{\mu \times \nu \times r \times \sigma} \left(\left| \widehat{\text{SW}}_p^p(\hat{\mu}_n, \hat{\nu}_m) - \text{SW}_p^p(\mu, \nu) \right| \geq \gamma(n, p, \beta, D, L) \right) \\ &\leq \frac{\beta}{2}. \end{aligned}$$

To conclude, note that whenever $\mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{B}) > 1 - \frac{\beta}{2}$, the theorem of total probability implies

$$\begin{aligned} \mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{A}) &= \mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{A} \cap \mathcal{B}) + \mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{A} \cap \mathcal{B}^c) \\ &= \mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{A} \cap \mathcal{B}) + \mathbb{P}_{p \times q \times r \times \sigma}(\mathcal{A} \mid \mathcal{B}^c) \mathbb{P}_{\mu \times \nu \times r \times \sigma}(\mathcal{B}^c) \\ &\leq \frac{\beta}{2} + \frac{\beta}{2} \cdot 1 \\ &= \beta. \end{aligned}$$

■

The condition for achieving the desired test power, given in Lemma 19, involves two main components.: the term $\gamma(n, p, \beta, D, L)$, which corresponds to the sample complexity bound (see Lemma 18), and the random empirical quantile threshold $\hat{c}_{1-\alpha, N}^B$, which constitutes the main difficulty in the theoretical analysis of permutation-based testing. To address this issue, we first rely on Domingo-Enrich et al. (2023, Lemma 6), which establishes a connection between the (random) empirical quantile threshold $\hat{c}_{1-\alpha, N}^B$ and the (deterministic) quantile $c_{1-\alpha, N}$ of the permuted test statistic, conditional on the samples and the projection directions.

Lemma 20 (Relation between population and empirical quantiles) *Let $\alpha, \beta \in (0, 1)$, and let $N \geq 1$. Denote by B the number of sampled permutations as described in Algorithm 1. Given the datasets $\mathcal{Y}_n = (y_1, \dots, y_n)$ and $\mathcal{Z}_m = (z_1, \dots, z_m)$, whose points lie in the ball centered at the origin with radius D , as well as a set of projection directions $\Theta = (\theta_1, \dots, \theta_L)$ on the unit sphere $\mathbb{S}^{d-1}$, we have*

$$\mathbb{P}_r(\hat{c}_{1-\alpha, N}^B \leq c_{1-\alpha, N} \mid \mathcal{Y}_n, \mathcal{Z}_m, \Theta) > 1 - \frac{\beta}{2},$$

$$\text{where } \alpha_1 := \left(\frac{\beta/2}{\binom{B}{\lfloor \alpha(B+1) \rfloor}} \right)^{1/\lfloor \alpha(B+1) \rfloor}.$$

Consequently, obtaining an upper bound for $c_{1-\alpha_1, N}$ immediately implies an upper bound for the empirical quantile threshold $\hat{c}_{1-\alpha, N}^B$. The next proposition establishes such an upper bound for the population quantile.

Proposition 21 (Upper bound for $c_{1-\alpha, N}$) *Let $p \geq 1$ and $\alpha \in (0, 1)$. Under the same assumptions as in Lemma 20, and for brevity, we denote $N := n + m$. Moreover, without loss of generality, we assume that $n \leq m$. Then the $(1 - \alpha)$ -quantile of the permutation distribution, $c_{1-\alpha, N}$ introduced in (5), satisfies*

$$c_{1-\alpha, N} \leq \sqrt{\frac{8(2D)^{2p} \log(1/\alpha)}{3n}} + \frac{(4D)^p \sqrt{2}}{\sqrt{n}}. \quad (9)$$

Proof For readability, we denote

$$h(n, m) := \frac{2(n + m - \frac{1}{2})}{nm \left(\frac{(2D)^p}{n} + \frac{(2D)^p}{m} \right)^2} \cdot \left(1 - \frac{1}{2 \max\{m, n\}} \right).$$

Setting $\epsilon := \sqrt{\ln \left(\frac{1}{\alpha} \right) \frac{1}{h(n, m)}}$, Lemma 12 guarantees

$$\mathbb{P}_\pi \left[\widehat{\text{SW}}_p^{p, \pi} \geq \sqrt{\ln \left(\frac{1}{\alpha} \right) \frac{1}{h(n, m)}} + \mathbb{E}_\pi(\widehat{\text{SW}}_p^{p, \pi}) \mid \mathcal{Y}_n, \mathcal{Z}_m, \Theta \right] \leq \alpha.$$

It implies, by definition of $c_{1-\alpha,N}$ (recalled in Eq. (6)), that

$$c_{1-\alpha,N} \leq \sqrt{\ln\left(\frac{1}{\alpha}\right) \cdot \frac{1}{h(n,m)}} + \mathbb{E}_\pi(\widehat{\text{SW}}_p^{p,\pi}). \quad (10)$$

We now derive a more explicit and tractable expression for $h(n,m)$. Since $1 \leq n \leq m$, we have

$$h(n,m) = \frac{n(2m-1)(n+m-\frac{1}{2})}{(2D)^{2p}(n+m)^2}.$$

Setting $x := \frac{m}{n} \geq 1$, we obtain

$$\begin{aligned} h(n,m) &= \frac{n(2x - \frac{1}{n})(1 + x - \frac{1}{2n})}{(2D)^{2p}(1+x)^2} \\ &\geq \frac{n(2x - \frac{1}{n})(1 + x - \frac{1}{2n})}{(2D)^{2p}(1+x)^2}, \end{aligned}$$

where the last inequality follows from the facts that $2x - \frac{1}{n} \geq 2x - 1$ and $1 + x - \frac{1}{2n} \geq x + \frac{1}{2}$.

Moreover, since the function $\phi(x) = \frac{(2x-1)(x+\frac{1}{2})}{(1+x)^2}$ is increasing for $x \geq 1$, it follows that

$$h(n,m) \geq \frac{3n}{8(2D)^{2p}}. \quad (11)$$

In Proposition 17, we have shown that, given the samples $\mathcal{Y}_n, \mathcal{Z}_m$ and the projection directions $\Theta = (\theta_1, \dots, \theta_L)$, the following bound holds:

$$\mathbb{E}_\pi[\widehat{\text{SW}}_p^{p,\pi}] \leq (4D)^p \frac{\sqrt{2}}{\sqrt{n}}. \quad (12)$$

Combining inequalities (10), (11), and (12), we obtain

$$c_{1-\alpha,N} \leq \sqrt{\frac{8(2D)^{2p} \log(1/\alpha)}{3n}} + \frac{(4D)^p \sqrt{2}}{\sqrt{n}}.$$

■

Having established an upper bound for $c_{1-\alpha,N}$, we now rely on Lemma 20 to derive the following proposition, which provides a bound for the empirical quantile $\widehat{c}_{1-\alpha,N}^B$.

Proposition 22 (Bound on the empirical quantile) *Let $\beta > 0$ and $\frac{1}{B+1} \leq \alpha < 1$, where B denotes the number of permutations as described in Algorithm 1. Set $\omega_\alpha := \lfloor \alpha(B+1) \rfloor$. Then, conditional on the samples $\mathcal{Y}_n$ and $\mathcal{Z}_m$ (with $n \leq m$ and $N := n+m$) whose points lie in the ball centered at the origin with radius D , as well as on a set of projection directions $\Theta = (\theta_1, \dots, \theta_L)$ on the unit sphere $\mathbb{S}^{d-1}$, we have*

$$\mathbb{P}_r \left(\widehat{c}_{1-\alpha,N}^B \leq \sqrt{\frac{8(2D)^{2p}}{3n} \log \left(\frac{2e}{\alpha \left(\frac{\beta}{2} \right)^{1/\omega_\alpha}} \right)} + \frac{(4D)^p}{\sqrt{n}} \right) > 1 - \frac{\beta}{2}.$$

Proof Proposition 21 guarantees that

$$c_{1-\alpha_1,N} \leq \sqrt{\frac{8(2D)^{2p} \log(1/\alpha_1)}{3n}} + \frac{(4D)^p \sqrt{2}}{\sqrt{n}}, \quad (13)$$

where

$$\alpha_1 = \left(\frac{\beta/2}{\binom{B}{\lfloor \alpha(B+1) \rfloor}} \right)^{1/\lfloor \alpha(B+1) \rfloor}.$$

To make this bound more explicit, we upper bound the logarithmic term. Let $\omega_\alpha := \lfloor \alpha(B+1) \rfloor \geq 1$. Then

$$\log \left(\binom{B}{\omega_\alpha}^{1/\omega_\alpha} \right) \leq \log \left(\frac{eB}{\omega_\alpha} \right) \leq \log \left(\frac{2eB}{\alpha(B+1)} \right) \leq \log \left(\frac{2e}{\alpha} \right). \quad (14)$$

The first inequality follows from the bound $\binom{n}{k} \leq \left(\frac{en}{k} \right)^k$, and the second from the fact that $\lfloor x \rfloor \geq \frac{x}{2}$ whenever $x \geq 1$.

Combining Eq. (13) and Eq. (14) yields

$$c_{1-\alpha_1, N} \leq \sqrt{\frac{8(2D)^{2p}}{3n} \log \left(\frac{2e}{\alpha \left(\frac{\beta}{2} \right)^{1/\omega_\alpha}} \right)} + \frac{(4D)^p \sqrt{2}}{\sqrt{n}}. \quad (15)$$

Hence, from Lemma 20, Eq. (15), and also conditioning on the samples $\mathcal{Y}_n, \mathcal{Z}_m$, and the projection directions $\theta_1, \dots, \theta_L$, we have

$$\mathbb{P}_r \left(\widehat{c}_{1-\alpha, N}^B \leq \sqrt{\frac{8(2D)^{2p}}{3n} \log \left(\frac{2e}{\alpha \left(\frac{\beta}{2} \right)^{1/\omega_\alpha}} \right)} + \frac{(4D)^p \sqrt{2}}{\sqrt{n}} \right) > 1 - \frac{\beta}{2}.$$

■

Finally, Theorem 5 is proved by combining Lemma 19 and Proposition 22, as outlined in the following.

Proof of Theorem 5

To ensure that the test has power at least $1 - \beta$, Lemma 19 shows it suffices to find conditions on μ and ν such that

$$\mathbb{P}_{\mu \times \nu \times r \times \sigma} \left[\text{SW}_p^p(\mu, \nu) \geq (2D)^p \left(\sqrt{\frac{\log(8/\beta)}{2L}} + \sqrt{\frac{\log(8/\beta)}{n}} \right) + \widehat{c}_{1-\alpha, N}^B \right] > 1 - \frac{\beta}{2}. \quad (16)$$

From Proposition 22 and the condition provided in Eq. (16), a sufficient condition for guaranteeing a power at least $1 - \beta$ is given by

$$\sqrt{\frac{8(2D)^{2p}}{3n} \log \left(\frac{2e}{\alpha \left(\frac{\beta}{2} \right)^{1/\omega_\alpha}} \right)} + \frac{(4D)^p \sqrt{2}}{\sqrt{n}} \leq -(2D)^p \left(\sqrt{\frac{\log(8/\beta)}{2L}} + \sqrt{\frac{\log(8/\beta)}{n}} \right) + \text{SW}_p^p(\mu, \nu)$$

or equivalently,

$$\text{SW}_p^p(\mu, \nu) \geq \sqrt{\frac{8(2D)^{2p}}{3n} \log \left(\frac{2e}{\alpha \left(\frac{\beta}{2} \right)^{1/\omega_\alpha}} \right)} + \frac{(4D)^p \sqrt{2}}{\sqrt{n}} + (2D)^p \left(\sqrt{\frac{\log(8/\beta)}{2L}} + \sqrt{\frac{\log(8/\beta)}{n}} \right).$$

The above inequality implies the claimed bound on the test's power.

12 LOWER BOUNDS

In this section, we present the proofs of the lower bounds stated in the main body, namely Proposition 6 and Proposition 7. We begin by explaining, in Section 12.1, why the problem of obtaining lower bounds for two-sample testing can be reduced to that of one-sample testing—i.e., goodness-of-fit testing. Next, we present the proofs of the two propositions.

12.1 Reduction to goodness-of-fit testing

Following Arias-Castro et al. (2018, Lemma 1), we recast the task of establishing lower bounds for the two-sample testing problem as the conceptually simpler task of establishing lower bounds for the one-sample testing problem (also called goodness-of-fit). Intuitively, goodness-of-fit can be viewed as a special case of two-sample testing in which one of the distributions is fully known—which is equivalent to having access to an infinite number of samples from it.

Fix a known reference distribution ν_0 and observe independent samples $Y_1, \dots, Y_n \sim \nu$, where ν is an unknown distribution. We consider the following goodness-of-fit testing problem:

$$\mathcal{H}_0 : \nu = \nu_0 \quad \text{vs.} \quad \mathcal{H}_1 : \text{SW}_p^p(\nu, \nu_0) \geq \epsilon_n.$$

Similarly to two-sample testing (see Section 2.3), we define the minimax separation for one-sample testing at level α and power $1 - \beta$ as

$$\epsilon_n^\dagger := \inf \left\{ \epsilon_n > 0 : \inf_{\Delta_n \in \Phi_{\alpha, n, \nu_0}} \sup_{\nu \in \mathcal{P}_1(\epsilon_n, \nu_0)} \mathbb{E}_\nu[1 - \Delta_n] \leq \beta \right\},$$

where

- Φ_{α, n, ν_0} is the set of tests Δ_n , based on n samples, of level at most α , that is, $\mathbb{E}_{\nu_0}[\Delta_n] \leq \alpha$;
- $\mathcal{P}_1(\epsilon_n, \nu_0)$ is the set of distributions ν such that $\text{SW}_p^p(\nu, \nu_0) \geq \epsilon_n$.

Because the two-sample setting involves estimating both distributions from data—unlike the one-sample case, where one distribution is known exactly—the minimax separation for two-sample testing cannot be smaller than that for one-sample testing. Consequently,

$$\epsilon_{n, m}^\dagger \geq \epsilon_n^\dagger,$$

where $\epsilon_n^\dagger$ and $\epsilon_{n, m}^\dagger$ denote the minimax separations for the one-sample and two-sample problems, respectively. In particular, a lower bound on the minimax separation in the one-sample setting directly implies a lower bound in the two-sample setting.

In what follows, we prove that $\epsilon_n^\dagger \gtrsim n^{-1/2}$ which immediately yields $\epsilon_{n, m}^\dagger \gtrsim n^{-1/2}$.

12.2 Proof of Proposition 6

Let $[d] := \{1, \dots, d\}$, and denote by $\mathcal{P}_{\text{Multi}}^{(d)}$ the class of multinomial distributions on $[d]$. In this section, we provide the proof of Proposition 6, which establishes a lower bound on the minimax separation rate for the two-sample testing problem over $\mathcal{P}_{\text{Multi}}^{(d)} \times \mathcal{P}_{\text{Multi}}^{(d)}$.

The proof of Proposition 6 is based on Ingster’s method (Ingster, 1993), a classical approach for deriving minimax lower bounds in hypothesis testing. We first provide a brief overview of this technique.

Ingster’s method for minimax lower bounds. Let ν_0 denote a fixed null distribution. The minimax Type II error at level α (for one-sample testing problem) is defined as

$$R_{n, \alpha}^\dagger := \inf_{\phi \in \Phi_{\alpha, n, \nu_0}} \sup_{\nu \in \mathcal{P}_1(\epsilon_n, \nu_0)} \mathbb{P}_\nu(\phi = 0),$$

where Φ_{α,n,ν_0} and $\mathcal{P}_1(\epsilon_n, \nu_0)$ are defined in Section 12.1.

Select $T \geq 1$ distributions $\nu_1, \dots, \nu_T \in \mathcal{P}_1(\epsilon_n, \nu_0)$, and define a mixture distribution Q such that, for any measurable set A ,

$$Q(A) := \frac{1}{T} \sum_{t=1}^T \nu_t^{\otimes n}(A),$$

where the n -fold product distribution of a distribution ν_t is denoted as $\nu_t^{\otimes n}$.

Given n i.i.d. observations $Y_1, \dots, Y_n$ drawn from an unknown distribution ν , we denote the likelihood ratio between Q and the null distribution ν_0 by

$$L_n(Y_1, \dots, Y_n) := \frac{dQ}{d\nu_0^{\otimes n}}(Y_1, \dots, Y_n) = \frac{1}{T} \sum_{t=1}^T \prod_{i=1}^n \frac{\nu_t(Y_i)}{\nu_0(Y_i)},$$

where, with a slight abuse of notation, we denote by ν_0 and $(\nu_t)_{t=1}^T$ both the probability measures and their densities with respect to the Lebesgue measure.

Given these definitions, the following lemma provides a condition on the likelihood ratio between the mixture distribution Q and the null distribution ν_0 to obtain a lower bound for the minimax Type II error $R_{n,\alpha}^\dagger$.

Lemma 23 (Lower bound via a mixture) *Fix $\alpha \in (0, 1)$ and $\beta \in (0, 1 - \alpha)$. If*

$$\mathbb{E}_{\nu_0}[L_n^2] \leq 1 + 4(1 - \alpha - \beta)^2,$$

then $R_{n,\alpha}^\dagger \geq \beta$.

The proof of this lemma can be found in Wasserman (2014, Section 11) and in Kim et al. (2022, Appendix H).

To apply Lemma 23, we need to (i) specify the distributions Q and ν_0 , (ii) compute the expectation of the squared likelihood ratio under the null. We borrow the construction of perturbed distributions around the uniform measure on $[d]$ from Kim et al. (2022, Appendix H), and adapt the corresponding computations to the sliced Wasserstein distance setting.

Let ν_0 denote the uniform distribution over the set $[d] := \{1, \dots, d\}$, that is, $\nu_0(k) = \frac{1}{d}$ for $k = 1, \dots, d$.

Define the set

$$\mathcal{M}_d := \{\eta \in \{-1, 1\}^d : \sum_{k=1}^d \eta_k = 0\}. \quad (17)$$

We may assume without loss of generality that d is even, as the proof for the odd-dimensional case follows analogously by setting the last coordinate to zero $\eta_d = 0$. Given $\eta \in \mathcal{M}_d$, we define the measure ν_η as

$$\nu_\eta(k) := \nu(k) + \frac{2\epsilon_n}{d} \eta_k, \quad \text{for } k = 1, \dots, d,$$

for some $\epsilon_n \leq \frac{1}{2}$ that we will specify later in the proof.

Note that, for any $\eta \in \mathcal{M}_d$, the distribution ν_η is a valid probability distribution on the set $[d]$. Indeed,

1. ν_η is non negative.

$$2. \sum_{k=1}^d \nu_\eta(k) = \sum_{k=1}^d \left(\nu(k) + \frac{2\epsilon_n}{d} \eta_k \right) = \sum_{k=1}^d \nu(k) + \frac{2\epsilon_n}{d} \sum_{k=1}^d \eta_k = 1.$$

Let us now evaluate the sliced Wasserstein distance between ν_0 and its perturbation ν_η . By construction, exactly $d/2$ coordinates k satisfy $\eta_k = 1$, and the remaining $d/2$ coordinates satisfy $\eta_k = -1$. For coordinates with $\eta_k = 1$, the weight in ν_η is increased by $\frac{2\epsilon_n}{d}$ compared to ν_0 , while for coordinates with $\eta_k = -1$, the weight is

decreased by the same amount. Hence, the total excess mass between the $d/2$ heavier points and the $d/2$ lighter points is

$$\frac{d}{2} \cdot \frac{2\epsilon_n}{d} = \epsilon_n.$$

Since ν_0 is uniform over $[d]$ and the distance between any two distinct points in $[d]$ is at least 1, moving any unit of excess mass incurs a cost of at least $1^2 = 1$. Therefore, the total transportation cost is at least ϵ_n . Since both ν_0 and ν_η are one-dimensional distributions, the sliced Wasserstein distance coincides with the standard Wasserstein distance and

$$\text{SW}_p^p(\nu_0, \nu_\eta) = \text{W}_p^p(\nu_0, \nu_\eta) \geq \epsilon_n.$$

Let T be the cardinality of $\mathcal{M}_d$. In order to apply Ingster's method, we introduce the uniform mixture $\nu_{\eta(1)}, \dots, \nu_{\eta(T)}$,

$$Q := \frac{1}{T} \sum_{\eta \in \mathcal{M}_d} \nu_\eta.$$

The likelihood ratio between Q and the null distribution ν_0 will be

$$L_n(Y_1, \dots, Y_n) := \frac{1}{T} \sum_{\eta \in \mathcal{M}_d} \prod_{i=1}^n \frac{\nu_\eta(Y_i)}{\nu_0(Y_i)}.$$

With those ingredients on hand, we are now ready to compute the expected value of the squared likelihood ratio. The subsequent computation follows exactly the same steps as in Kim et al. (2022, Appendix H). We present it here for completeness.

$$\begin{aligned} L_n^2(Y_1, \dots, Y_n) &= \frac{1}{T^2} \sum_{\eta, \eta' \in \mathcal{M}_d} \prod_{i=1}^n \frac{\nu_\eta(Y_i) \nu_{\eta'}(Y_i)}{\nu_0(Y_i)^2} \\ &= \frac{1}{T^2} \sum_{\eta, \eta' \in \mathcal{M}_d} \prod_{i=1}^n \frac{(\frac{1}{d} + \frac{2\epsilon_n}{d} \eta_{Y_i})(\frac{1}{d} + \frac{2\epsilon_n}{d} \eta'_{Y_i})}{1/d^2} \\ &= \frac{1}{T^2} \sum_{\eta, \eta' \in \mathcal{M}_d} \prod_{i=1}^n (1 + 2\epsilon_n \eta_{Y_i})(1 + 2\epsilon_n \eta'_{Y_i}). \end{aligned}$$

By definition of $\eta \in \mathcal{M}_d$, taking the expectation over $Y_1, \dots, Y_n \sim \nu_0$, we obtain

$$\mathbb{E}_{\nu_0}[L_n^2] = \frac{1}{T^2} \sum_{\eta, \eta' \in \mathcal{M}_d} \left(1 + \frac{4\epsilon_n^2}{d} \sum_{k=1}^d \eta_k \eta'_k \right)^n.$$

Moreover, using the inequality $1 + x \leq e^x$, which holds for any $x \in \mathbb{R}$,

$$\mathbb{E}_{\nu_0}[L_n^2] \leq \frac{1}{T^2} \sum_{\eta, \eta' \in \mathcal{M}_d} \exp\left(\frac{4n\epsilon_n^2}{d} \sum_{k=1}^d \eta_k \eta'_k\right).$$

Let η and η^* be independent random variables uniformly distributed on $\mathcal{M}_d$. We have

$$\frac{1}{T^2} \sum_{\eta, \eta' \in \mathcal{M}_d} \exp\left(\frac{4n\epsilon_n^2}{d} \sum_{k=1}^d \eta_k \eta'_k\right) = \mathbb{E}_{\eta, \eta^*} \left[\exp\left(\frac{4n\epsilon_n^2}{d} \langle \eta, \eta^* \rangle\right) \right].$$

Applying Lemma 2 in Dubhashi and Ranjan (1996), we obtain

$$\mathbb{E}_{\nu_0}[L_n^2] \leq \mathbb{E}_{\eta, \eta^*} \left[\exp\left(\frac{4n\epsilon_n^2}{d} \langle \eta, \eta^* \rangle\right) \right]$$

$$\begin{aligned}
 &\leq \prod_{k=1}^d \mathbb{E}_{\eta_k, \eta_k^*} \left[\exp \left(\frac{4n\epsilon_n^2}{d} \eta_k \eta_k^* \right) \right] \\
 &= \prod_{k=1}^d \cosh \left(\frac{4n\epsilon_n^2}{d} \right) \\
 &\leq \exp \left(\frac{d}{2} \left(\frac{4n\epsilon_n^2}{d} \right)^2 \right) = \exp \left(\frac{8n^2\epsilon_n^4}{d} \right),
 \end{aligned}$$

where the last inequality follows from $\cosh(x) \leq e^{x^2/2}$, which holds for any $x \in \mathbb{R}$.

The proof is concluded by observing that Lemma 23 guarantees that the minimax Type II error is lower bounded by β provided that

$$\epsilon_n \leq \frac{1}{\sqrt{n}} \sqrt[4]{\frac{d \log [1 + 4(1 - \alpha - \beta)^2]}{8}}.$$

12.3 Proof of Proposition 7

Proposition 6 states a lower bound on the minimax separation rate for the two-sample testing problem over the class of multinomial distributions. Turning to the more general class of distributions with bounded support, denoted by $\mathcal{P}_{\mathbb{R}^d}(D)$ for some $D > 0$, the proof of Proposition 7 establishes the corresponding lower bound on the minimax separation rate for the two-sample testing problem over this class.

We begin by presenting some preliminary ingredients required for the proof of Proposition 7.

Main ingredients for the proof

We begin this section by recalling a basic fact: the uniform distribution on the sphere is invariant under orthogonal transformations, as stated below.

Lemma 24 (Rotation invariance) *If U is uniform on $\mathbb{S}^{d-1}$ and Q is an orthogonal matrix, then QU has the same distribution as U .*

Proof Let σ denote the $(d-1)$ -dimensional surface area measure on $\mathbb{S}^{d-1}$. Orthogonal maps preserve surface area on the sphere, that is, $\sigma(QA) = \sigma(A)$ for every Borel set $A \subset \mathbb{S}^{d-1}$. Hence, for any Borel set A , we have

$$\mathbb{P}(QU \in A) = \mathbb{P}(U \in Q^{-1}A) = \frac{\sigma(Q^{-1}A)}{\sigma(\mathbb{S}^{d-1})} = \frac{\sigma A}{\sigma(\mathbb{S}^{d-1})} = \mathbb{P}(U \in A).$$

Therefore QU and U have the same law. ■

We now identify the law of a single coordinate under the uniform distribution on the sphere.

Lemma 25 (First coordinate of a spherical uniform distribution) *Let $U = (U_1, \dots, U_d) \sim \text{Unif}(\mathbb{S}^{d-1})$.*

1. For $d = 1$, $U_1^2 \equiv 1$ a.s. (a degenerate distribution at 1).

2. For $d \geq 2$, $U_1^2 \sim \text{Beta}\left(\frac{1}{2}, \frac{d-1}{2}\right)$.

Proof

Case $d = 1$. It is clear that $\mathbb{S}^0 = \{\pm 1\}$ and $U_1 = \pm 1$ with probability $1/2$ each. Therefore, $U_1^2 \equiv 1$ almost surely.

Case $d \geq 2$. Let $Z = (Z_1, \dots, Z_d)$ be standard normal in $\mathbb{R}^d$ (i.e., $Z \sim \mathcal{N}(0, I_d)$). By Equation (3.15) in Vershynin (2018), $U := Z/\|Z\| \sim \text{Unif}(\mathbb{S}^{d-1})$. Since $Z_1 \sim \mathcal{N}(0, 1)$ and the Z_i 's are independent,

$$X := Z_1^2 \sim \chi_1^2 = \Gamma\left(\frac{1}{2}; 2\right), \quad Y := \sum_{i=2}^d Z_i^2 \sim \chi_{d-1}^2 = \Gamma\left(\frac{d-1}{2}; 2\right),$$

with $X \perp Y$.

Finally, using the classical fact that if $M \sim \chi^2(m)$ and $N \sim \chi^2(n)$ are independent, then

$$\frac{M}{M+N} \sim \text{Beta}\left(\frac{m}{2}, \frac{n}{2}\right),$$

(see, e.g., Soch et al. (2025, Section 3.9)), we conclude that

$$U_1^2 \sim \text{Beta}\left(\frac{1}{2}, \frac{d-1}{2}\right).$$

As a consequence, for $p \geq 1$, and following the result of Balakrishnan and Nevzorov (2004, Page 145), we have

$$\mathbb{E}|U_1|^p = \begin{cases} \frac{\Gamma(\frac{p+1}{2})\Gamma(\frac{d}{2})}{\sqrt{\pi}\Gamma(\frac{p+d}{2})}, & d \geq 2, \\ 1, & d = 1, \end{cases}$$

where $\Gamma(\cdot)$ denotes the classical Gamma function. ■

We recall two classical inequalities that are particularly useful for bounding Gamma functions: Wendel's inequality and Stirling's inequalities.

Lemma 26 (Wendel's inequality) *We denote by $\Gamma(\cdot)$ the classical Gamma function. For $0 \leq a < 1$ and $x > 0$,*

$$\left(\frac{x}{x+a}\right)^{1-a} \leq \frac{\Gamma(x+a)}{x^a \Gamma(x)} \leq 1.$$

Lemma 27 (Stirling's inequality) *We denote by $\Gamma(\cdot)$ the classical Gamma function. For $x > 0$,*

$$\sqrt{2\pi} x^{x-\frac{1}{2}} e^{-x} \leq \Gamma(x) \leq \sqrt{2\pi} x^{x-\frac{1}{2}} e^{-x} e^{\frac{1}{12x}}.$$

Proofs of Lemmas 26 and 27 can be found in Wendel (1948) and Jameson (2015, Theorem 1), respectively.

We then close these preliminaries with a simple yet essential fact from linear algebra.

Lemma 28 *For any unit vector $u \in \mathbb{R}^d$ with $\|u\| = 1$, there exists a rotation $Q \in SO(d)$ such that $Qe_1 = u$, where $e_1 = (1, 0, \dots, 0)^\top$ and $SO(d) = \{Q \in \mathbb{R}^{d \times d} : Q^\top Q = I, \det Q = 1\}$.*

Proof Extend $\{u\}$ to a basis of $\mathbb{R}^d$ and apply the Gram–Schmidt procedure to obtain an orthonormal basis $\{u, v_2, \dots, v_d\}$. Let

$$Q := [u \mid v_2 \mid \dots \mid v_d] \in \mathbb{R}^{d \times d}.$$

Then the columns of Q are orthonormal and $Q^\top Q = I$. In particular, Q is orthogonal and satisfies $Qe_1 = u$.

If $\det(Q) = 1$, then $Q \in SO(d)$ and we are done. If $\det(Q) = -1$, replace one column v_j with $-v_j$ for some $j \geq 2$ (e.g., $j = 2$). The resulting matrix

$$\tilde{Q} = [u \mid -v_2 \mid v_3 \mid \dots \mid v_d]$$

remains orthogonal, still satisfies $\tilde{Q}e_1 = u$, and has $\det(\tilde{Q}) = -\det(Q) = 1$. Therefore, there exists $Q \in SO(d)$ such that $Qe_1 = u$. ■

With these preparatory results in place, we are now ready to prove Proposition 7.

12.3.1 Formal proof

As discussed in Section 12.1, once a lower bound for the minimax rate $\epsilon_n^\dagger$ for goodness-of-fit testing is established, a lower bound $\epsilon_{n,m}^\dagger$ for two-sample testing follows immediately. For the general case of distributions with bounded support, we apply Le Cam two-point method (Le Cam, 1973, 2012), a classical techniques for deriving minimax lower bounds.

Let ν_0 be a reference distribution. Let $\epsilon_n > 0$ and let μ_0 be an alternative distribution in $\mathcal{P}_1(\epsilon_n, \nu_0)$. Using the same notation as for the one-sample testing problem (see Section 12.1), we obtain

$$\begin{aligned}
 \inf_{\Delta_n \in \Phi_{\alpha,n,\nu_0}} \sup_{\mu \in \mathcal{P}_1(\epsilon_n, \nu_0)} \mathbb{E}_\mu[1 - \Delta_n] &\geq \inf_{\Delta_n \in \Phi_{\alpha,n,\nu_0}} \mathbb{E}_{\mu_0}[1 - \Delta_n] \\
 &= 1 - \sup_{\Delta_n \in \Phi_{\alpha,n,\nu_0}} \mathbb{E}_{\mu_0}[\Delta_n] \\
 &= 1 - \sup_{\Delta_n \in \Phi_{\alpha,n,\nu_0}} \{\mathbb{E}_{\mu_0}[\Delta_n] - \mathbb{E}_{\nu_0}[\Delta_n] + \mathbb{E}_{\nu_0}[\Delta_n]\} \\
 &\geq 1 - \alpha - \sup_{\Delta_n \in \Phi_{\alpha,n,\nu_0}} \{\mathbb{E}_{\mu_0}[\Delta_n] - \mathbb{E}_{\nu_0}[\Delta_n]\} \\
 &\geq 1 - \alpha - d_{\text{TV}}(\mu_0^{\otimes n}, \nu_0^{\otimes n}) \\
 &\stackrel{(i)}{\geq} 1 - \alpha - 1 + \frac{1}{2}e^{-D_{\text{KL}}(\mu_0^{\otimes n} \parallel \nu_0^{\otimes n})} \\
 &\stackrel{(ii)}{=} \frac{1}{2}e^{-n D_{\text{KL}}(\mu_0 \parallel \nu_0)} - \alpha,
 \end{aligned}$$

where (i) follows from the Bretagnolle–Huber inequality introduced in Canonne (2022, Lemma 3), and (ii) uses the chain rule for the KL divergence. We recall that $\mu_0^{\otimes n}$ and $\nu_0^{\otimes n}$ are the n -fold product distributions of μ_0 and ν_0 , respectively.

Consequently, the minimax Type II error is at least β , that is,

$$\inf_{\Delta_n \in \Phi_{\alpha,n,\nu_0}} \sup_{\mu \in \mathcal{P}_1(\epsilon_n, \nu_0)} \mathbb{E}_\mu[1 - \Delta_n] \geq \beta,$$

provided that

$$\alpha + \beta < 0.5 \quad \text{and} \quad D_{\text{KL}}(\mu_0 \parallel \nu_0) \leq \frac{1}{n} \log \left(\frac{1}{2(\alpha + \beta)} \right).$$

In order to obtain a lower bound, we need to choose $\mu_0, \nu_0 \in \mathcal{P}_{\mathbb{R}^d}(D)$ such that

$$D_{\text{KL}}(\mu_0 \parallel \nu_0) \leq \frac{1}{n} \log \left(\frac{1}{2(\alpha + \beta)} \right). \quad (18)$$

and

$$\text{SW}_p^p(\mu_0, \nu_0) \geq \epsilon_n. \quad (19)$$

Inspired by the works of Kim and Schrab (2023, Section E.10.1) and adapting their work to the case of the sliced Wasserstein distance, we choose

$$\mu_0 = p_0 \delta_x + (1 - p_0) \delta_v, \quad \nu_0 = q_0 \delta_x + (1 - q_0) \delta_v,$$

where $x, v \in \mathbb{R}^d$ with $\|x - v\|_2 = D$, $0 \leq p_0, q_0 \leq 1$, and δ_x denotes the Dirac measure at x . Furthermore, we also choose

$$p_0 = \frac{1}{2} + \min \left\{ \sqrt{\frac{1}{4n} \log \left(\frac{1}{2(\alpha + \beta)} \right)}, \frac{1}{2} \right\}, \quad q_0 = \frac{1}{2}.$$

With those parameters, we have

$$D_{\text{KL}}(\mu_0 \parallel \nu_0) \leq \frac{(p_0 - q_0)^2}{q_0(1 - q_0)} = \min \left\{ \frac{1}{n} \log \left(\frac{1}{2(\alpha + \beta)} \right), 1 \right\} \leq \frac{1}{n} \log \left(\frac{1}{2(\alpha + \beta)} \right),$$

where the first inequality follows directly from a classical inequality between the KL and χ^2 divergences (see, e.g., Tsybakov (2008, Lemma 2.7)).

Moreover, we have

$$\begin{aligned} \text{SW}_p^p(\mu_0, \nu_0) &= \int_{\mathbb{S}^{d-1}} W_p^p(\Pi_{\#}^{\theta} \mu_0, \Pi_{\#}^{\theta} \nu_0) \sigma(d\theta) \\ &\stackrel{(i)}{=} \int_{\mathbb{S}^{d-1}} |p_0 - q_0| |\langle x - v, \theta \rangle|^p \sigma(d\theta) \\ &\stackrel{(ii)}{=} |p_0 - q_0| \|x - v\|^p \int_{\mathbb{S}^{d-1}} |\langle u, \theta \rangle|^p \sigma(d\theta), \end{aligned}$$

where (i) follows from the computation of the Wasserstein distance between two Dirac measures in one dimension (see, e.g., Chewi et al. (2024, Example 1.4)), and (ii) from setting $u = \frac{x-v}{\|x-v\|}$.

We can further develop this expression as

$$\begin{aligned} \text{SW}_p^p(\mu_0, \nu_0) &= |p_0 - q_0| \|x - v\|^p \mathbb{E}_{\theta \sim \text{Unif}(\mathbb{S}^{d-1})} [|\langle u, \theta \rangle|^p] \\ &\stackrel{(iii)}{=} |p_0 - q_0| \|x - v\|^p \mathbb{E}_{\theta \sim \text{Unif}(\mathbb{S}^{d-1})} [|\langle Qe_1, \theta \rangle|^p] \\ &= |p_0 - q_0| \|x - v\|^p \mathbb{E}_{\theta \sim \text{Unif}(\mathbb{S}^{d-1})} [|\langle e_1, Q^\top \theta \rangle|^p] \\ &\stackrel{(iv)}{=} |p_0 - q_0| \|x - v\|^p \mathbb{E}_{\theta \sim \text{Unif}(\mathbb{S}^{d-1})} [|\langle e_1, \theta \rangle|^p] \\ &= |p_0 - q_0| \|x - v\|^p \mathbb{E}_{\theta \sim \text{Unif}(\mathbb{S}^{d-1})} [|\theta_1|^p] \\ &\stackrel{(v)}{=} |p_0 - q_0| \|x - v\|^p \begin{cases} \frac{\Gamma(\frac{p+1}{2}) \Gamma(\frac{d}{2})}{\sqrt{\pi} \Gamma(\frac{p+d}{2})}, & d \geq 2, \\ 1, & d = 1, \end{cases} \end{aligned}$$

where (iii) follows from Lemma 28, (iv) from Lemma 24, and (v) from Lemma 25.

We now derive lower bound for the term

$$c_{d,p} := \frac{\Gamma(\frac{p+1}{2}) \Gamma(\frac{d}{2})}{\sqrt{\pi} \Gamma(\frac{p+d}{2})}.$$

The term $\Gamma(\frac{p+1}{2})$ can be controled by applying Stirling's lower bound (recalled in Lemma 27). In particular, we have

$$\Gamma\left(\frac{p+1}{2}\right) \geq \sqrt{2\pi} \left(\frac{p+1}{2}\right)^{\frac{p}{2}} e^{-\frac{(p+1)}{2}}.$$

We next control the remaining term.

Write $\frac{p}{2} = m + r$ with $m = \lfloor p/2 \rfloor \in \mathbb{N}$ and $r \in [0, 1)$. We then can decompose the remaining terms as follow:

$$\frac{\Gamma(\frac{d}{2})}{\Gamma(\frac{d}{2} + m + r)} = \frac{\Gamma(\frac{d}{2})}{\Gamma(\frac{d}{2} + m)} \cdot \frac{\Gamma(\frac{d}{2} + m)}{\Gamma(\frac{d}{2} + m + r)}.$$

By the recursion formula of the gamma function, $\Gamma(p+1) = p\Gamma(p)$ for $p \geq 0$, we obtain

$$\frac{\Gamma(\frac{d}{2})}{\Gamma(\frac{d}{2} + m)} \geq (\frac{d}{2} + m)^{-m}.$$

For the second factor, since $0 \leq r < 1$, Wendel's inequality (recalled in Lemma 26) yields

$$\frac{\Gamma(\frac{d}{2} + m)}{\Gamma(\frac{d}{2} + m + r)} \geq (\frac{d}{2} + m + r)^{-r}.$$

Hence,

$$\frac{\Gamma(\frac{d}{2})}{\Gamma(\frac{d+p}{2})} = \frac{\Gamma(\frac{d}{2})}{\Gamma(\frac{d}{2} + m + r)} \geq (\frac{d}{2} + m)^{-m} (\frac{d}{2} + m + r)^{-r} \geq \left(\frac{d+p}{2}\right)^{-\frac{p}{2}}.$$

Combining the above estimates gives

$$c_{d,p} \geq \frac{\sqrt{2\pi}}{\sqrt{\pi}} \left(\frac{p+1}{2}\right)^{\frac{p}{2}} e^{-\frac{(p+1)}{2}} \cdot \left(\frac{d+p}{2}\right)^{-\frac{p}{2}} = \sqrt{\frac{2}{e}} \left[\frac{p+1}{e(d+p)}\right]^{\frac{p}{2}}.$$

As a result, we have

$$\text{SW}_p^p(\mu_0, \nu_0) \geq \frac{D^p}{\sqrt{2e}} \left[\frac{p+1}{e(d+p)}\right]^{\frac{p}{2}} \min \left\{ \sqrt{\frac{1}{n} \log \left(\frac{1}{2(\alpha + \beta)} \right)}, 1 \right\}.$$

Setting $\epsilon_n := \frac{D^p}{\sqrt{2e}} \left[\frac{p+1}{e(d+p)}\right]^{\frac{p}{2}} \min \left\{ \sqrt{\frac{1}{n} \log \left(\frac{1}{2(\alpha + \beta)} \right)}, 1 \right\}$, conditions (18) and (19) hold, and thus the minimax separation satisfies

$$\epsilon_n^\dagger \geq \epsilon_n = \frac{D^p}{\sqrt{2e}} \left[\frac{p+1}{e(d+p)}\right]^{\frac{p}{2}} \min \left\{ \sqrt{\frac{1}{n} \log \left(\frac{1}{2(\alpha + \beta)} \right)}, 1 \right\}.$$

The above bound implies the bound stated in Proposition 7.

13 PERMUTATION APPROACH

In this section, we present the motivation for employing a permutation-based strategy in the two-sample testing problem using the sliced Wasserstein distance.

In hypothesis testing, the determination of the critical value plays a central role, as it directly governs the decision to reject the null hypothesis. In general, existing approaches for determining the critical value can be grouped into two main categories: asymptotic and non-asymptotic methods.

Asymptotic method. In the asymptotic approach, the critical value is determined from the limiting distribution of the test statistic under the null hypothesis (see, e.g., Shekhar et al. (2022, 2023); Zaremba et al. (2013)). Specifically, to control the Type I error in the asymptotic regime, the critical value is chosen as the $(1 - \alpha)$ quantile of the null distribution.

However, this asymptotic approach is not without limitations. In many cases, the limiting null distribution is either intractable or lacks a convenient closed form. For instance, Gretton et al. (2012, Theorem 12) show that the null distribution of the squared empirical biased estimator of MMD is an infinite weighted sum of independent χ^2 random variables, with weights given by the eigenvalues of the kernel operator. This makes the distribution highly dependent on the kernel choice and analytically difficult to handle. Likewise, more general tests based on U-statistics face similar challenges in the multinomial setting: the shape of the null distribution is intricately tied to the probabilities defining the underlying multinomial structure (see Kim et al. (2022, Figure 1)).

The test statistic $\widehat{\text{SW}}_p^p$ defined in (2) faces the same challenge. To gain intuition about its null distribution, we approximate it empirically through simulation. Specifically, we consider the case $p = 2$ and generate samples from three pairs of identical distributions with $n = m = 8000$ and $L = n = 8000$ projection directions. Repeating this procedure 4000 times produces the histograms shown in Figure 5. The resulting distributions vary significantly across underlying data distributions, which are unknown in practice, making it difficult to compute the critical threshold analytically.




Figure 5: Histograms of the test statistic $\widehat{SW}_2^2$ computed from 8000 i.i.d. samples drawn respectively from (left) Gaussian $\mathcal{N}(0, I_2)$, (middle) Uniform on $[-1, 1]^2$, and (right) a two-component Gaussian mixture $\frac{1}{2}\mathcal{N}(0, I_2) + \frac{1}{2}\mathcal{N}(\mathbf{m}, I_2)$, where $\mathbf{m} = (2.5, 2.5)^\top$.

Recent work by Rodríguez-Vítores et al. (2025) established a Central Limit Theorem (CLT) that characterizes the asymptotic distribution of the empirical Sliced Wasserstein distance. Under suitable regularity conditions, the theorem states that: As $n, m = m(n) \rightarrow +\infty$, we obtain that

$$\sqrt{\frac{k \frac{nm}{n+m}}{k + \frac{nm}{n+m}}} \left(\widehat{SW}_p^p - SW_p^p(\mu, \nu) \right) \xrightarrow{d} \mathcal{N}(0, (1 - \tau)w_{\mu, \nu}^2 + \tau((1 - \lambda)v_{\mu, \nu}^2 + \lambda v_{\nu, \mu}^2)). \quad (20)$$

Here, $\xrightarrow{d}$ denotes convergence in distribution, n and $m = m(n)$ are the respective sample sizes, and $k = k(n)$ denotes the number of projection directions, and the parameters τ and λ are asymptotic limits defined by

$$\tau = \lim_{n \rightarrow +\infty} \frac{k}{k + \frac{nm}{n+m}} \quad \text{and} \quad \lambda = \lim_{n \rightarrow +\infty} \frac{n}{n + m}.$$

Moreover, the two components of the asymptotic variance are defined as²

$$w_{\mu, \nu}^2 := \int_{\mathbb{S}^{d-1}} W_p^{2p}(\Pi_{\#}^{\theta} \mu, \Pi_{\#}^{\theta} \nu) d\sigma(\theta) - SW_p^{2p}(\mu, \nu),$$

and

$$v_{\mu, \nu}^2 := \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} \text{Cov}_{\mu}(\phi^{\theta}, \phi^{\eta}) d\sigma(\theta) d\sigma(\eta),$$

where ϕ^{θ} is any c-concave optimal transport potential from $\Pi_{\#}^{\theta} \mu$ to $\Pi_{\#}^{\theta} \nu$.

As can be observed, the asymptotic variances $w_{\mu, \nu}^2$ and $v_{\mu, \nu}^2$ are hard to compute and depend on the underlying unknown distributions. Furthermore, those asymptotic variances are equal to zero under the null hypothesis. As a consequence, it is not possible to construct a meaningful test based on this asymptotic null distribution. More broadly, research on the limiting distribution of the sliced Wasserstein distance remains limited, which makes it difficult to address testing problems through asymptotic analysis.

Finally, the number of observations is often limited due to economical or biological constraints (see, e.g., the discussion about applications in neurosciences in Albert (2015, Section 0.4.1)). Consequently, asymptotic procedures may not be suitable in such small sample size settings. This is one of the reasons why non-asymptotic methods have been developed.

Non-asymptotic methods. A classical approach is to determine the threshold of a test using a concentration inequality or a tail bound for the test statistic under the null hypothesis (see, e.g., Gretton et al. (2012, Section 4.2), Wei and Dudley (2011)). This non-asymptotic approach avoids dependence on asymptotic approximations and is therefore less sensitive to distributional assumptions. On the downside, the resulting bounds

²We use the same notation as in the definition of the sliced Wasserstein distance presented in Subsection 2.2

are usually conservative, leading to Type I errors much smaller than α and lower power. Moreover, many of these guarantees involve constants that are not explicitly determined in theory, making the resulting thresholds difficult to compute in practice (cf. Wang et al. (2021a, Section III) and the illustration of Kim et al. (2022, Section 9)).

There exist other data-driven approaches such as the bootstrap (Tibshirani and Efron, 1993) and subsampling (Politis et al., 1999). Many tests based on these methods have been extensively studied (see, e.g., Van Der Vaart and Wellner (1996); Romano (1988); Hu and Lin (2025)). However, as noted in Albert (2015, Section 0.5.2), these methods do not guarantee Type I error control in a non-asymptotic sense, i.e., for small sample sizes — a property that is ensured by permutation-based approaches. As a result, when both methods are applicable and non-asymptotic Type I error control is required, the permutation approach should be preferred.

For the reasons discussed above, we adopt the permutation approach for the two-sample test based on the sliced Wasserstein distance.

14 EXISTING APPROACHES FOR CONTROLLING THE RANDOM CRITICAL VALUE OF PERMUTATION-BASED (TWO-SAMPLE) TESTS

In the previous section, we motivated the use of a permutation approach for the sliced Wasserstein distance two-sample test. We now review how the challenges associated with the theoretical analysis of permutation based two-sample tests have been addressed in the literature (see also, for instance, Kim et al. (2022, Section 1.2)).

A central difficulty in analyzing the non-asymptotic power of permutation tests lies in controlling the random critical value of the test, and in particular the dependence structure induced by permutation sampling. To the best of our knowledge, no prior work has addressed this issue in the context of the sliced Wasserstein distance.

In this section, we review two existing approaches for controlling the random rejection threshold in permutation-based tests. Both rely on the structure of U-statistics: the first through a coupling argument (Kim et al., 2022), and the second via a concentration inequality for permuted sums (Albert, 2019). We then explain why current techniques cannot be directly extended to the sliced Wasserstein setting, motivating the development of a new analytical framework—one of the main theoretical contributions of this work.

14.1 Coupling technique for U-statistics

Kim et al. (2022, Section 6) investigate the performance of permutation-based procedures with a focus on degenerate second-order U-statistics, a broad class of estimators which encompasses many commonly used two-sample test statistics (Serfling, 2009, Chapter 5). This line of work has inspired several recent theoretical analysis in permutation-based two-sample testing (Schrab et al., 2023; Kim and Schrab, 2023; Choi and Kim, 2024; Chatalic et al., 2025). In order to clarify this approach, we begin by formally introducing second-order U-statistics. We use the same setting and notation as in Section 4.1.

Definition 29 (Second order U-statistic) *Let X be a measurable space. Let $g : \mathsf{X} \times \mathsf{X} \rightarrow \mathbb{R}$ be a measurable symmetric bivariate function, that is, $g(x, y) = g(y, x)$ for all $x, y \in \mathsf{X}$. We introduce*

$$h(y_1, y_2; z_1, z_2) := g(y_1, y_2) + g(z_1, z_2) - g(y_1, z_2) - g(y_2, z_1).$$

Let $\mathbf{i}_n^2$ be the set of all couples drawn without replacement from the set $\{1, \dots, n\}$. Then, the corresponding U-statistic is given by

$$U_{n,m} := \frac{1}{n(n-1)m(m-1)} \sum_{(i_1, i_2) \in \mathbf{i}_n^2} \sum_{(j_1, j_2) \in \mathbf{i}_m^2} h(Y_{i_1}, Y_{i_2}; Z_{j_1}, Z_{j_2}).$$

Moreover, given a permutation $\pi \in S_N$ (with $N = n + m$), the permuted U-statistic associated with π is defined as

$$U_{n,m}^\pi := \frac{1}{n(n-1)m(m-1)} \sum_{(i_1, i_2) \in \mathbf{i}_n^2} \sum_{(j_1, j_2) \in \mathbf{i}_m^2} h(X_{\pi_{i_1}}, X_{\pi_{i_2}}; X_{\pi_{n+j_1}}, X_{\pi_{n+j_2}}).$$

Assume $n \leq m$, and let $L = (l_1, \dots, l_n)$ denote an n -tuple drawn uniformly without replacement from the set $\{1, \dots, m\}$. Given L , we define the auxiliary test statistic

$$\tilde{U}_{n,m}^{\pi,L} := \frac{1}{n(n-1)} \sum_{(k_1, k_2) \in \mathbf{i}_2^n} h(X_{\pi_{k_1}}, X_{\pi_{k_2}}; X_{\pi_{n+l_{k_1}}}, X_{\pi_{n+l_{k_2}}}).$$

Note that the statistic $U_{n,m}^{\pi}$ is the conditional expectation of $\tilde{U}_{n,m}^{\pi,L}$ with respect to L , given all other random quantities, that is,

$$U_{n,m}^{\pi} = \mathbb{E}_L \left[\tilde{U}_{n,m}^{\pi,L} \mid \mathcal{X}_N, \pi \right]. \quad (21)$$

The usefulness of this construction becomes apparent when applying the following lemma from Kim et al. (2022). It is a coupling argument based on a symmetrization trick (Dümbgen, 1998).

Lemma 30 (Coupling with i.i.d. random variables) *Let $N \geq 1$ and let π be a random permutation uniformly distributed on the symmetric group S_N , i.e., the set of all permutations of $\{1, \dots, N\}$. Let $k = \lfloor N/2 \rfloor$ and let $\Delta = (\delta_1, \dots, \delta_k)$ be a random vector of independent Bernoulli trials, taking values in $\{0, 1\}^k$ and independent of π . Define the transformation $T_\Delta : S_N \rightarrow S_N$ of any permutation $\tau = (\tau_1, \dots, \tau_N) \in S_N$ as,*

$$(T_\Delta(\tau))_{2i-1} = \delta_i \tau_{2i-1} + (1 - \delta_i) \tau_{2i}, \quad (T_\Delta(\tau))_{2i} = (1 - \delta_i) \tau_{2i-1} + \delta_i \tau_{2i},$$

for $i = 1, \dots, k$. Let $\pi' := T_\Delta(\pi)$. Then π' and π are identically distributed.

For completeness we provide a proof of this lemma.

Proof Assume that N is even. If N is odd, we simply set $(T_\Delta(\tau))_N = \tau_N$ and the proof remains unchanged.

To prove that π' and π are identically distributed, it suffices to show that, for any arbitrary permutation $\sigma \in S_N$,

$$\mathbb{P}(\pi' = \sigma) = \mathbb{P}(\pi = \sigma).$$

By applying the law of total probability, we obtain

$$\mathbb{P}(\pi' = \sigma) = \mathbb{E}[\mathbb{P}(T_\Delta(\pi) = \sigma \mid \Delta)].$$

Since T_Δ is an involution, i.e., $T_\Delta(T_\Delta(\tau)) = \tau$ for all $\tau \in S_N$, the event $\{T_\Delta(\pi) = \sigma\}$ can be equivalently expressed as $\{\pi = T_\Delta(\sigma)\}$. Moreover, since π and Δ are independent, conditioning on Δ does not affect the distribution of π . Following these observations, it holds that

$$\mathbb{P}(\pi' = \sigma) = \mathbb{E}[\mathbb{P}(\pi = T_\Delta(\sigma) \mid \Delta)] = \mathbb{E}[\mathbb{P}(\pi = T_\Delta(\sigma))] = \mathbb{E}\left(\frac{1}{N!}\right) = \frac{1}{N!} = \mathbb{P}(\pi = \sigma).$$

Since this equality holds for every $\sigma \in S_N$, the claim follows. ■

The above lemma implies that the distribution of $\tilde{U}_{n,m}^{\pi,L}$ remains unchanged if we randomly swap X_{π_k} and $X_{\pi_{n+l_k}}$ for $k \in \{1, \dots, n\}$. In other words, it allows us to connect the statistic to i.i.d. Bernoulli random variables, which are easier to handle analytically. Moreover, due to the symmetry of $g(x, y)$ and the definition of h , this argument naturally extends to i.i.d. Rademacher random variables, viewed as random sign flips.

Indeed, let $\zeta_1, \dots, \zeta_n$ be i.i.d. Rademacher random variables. Then

$$\tilde{U}_{n,m}^{\pi,L,\zeta} := \frac{1}{n(n-1)} \sum_{(k_1, k_2) \in \mathbf{i}_2^n} \zeta_{k_1} \zeta_{k_2} h(X_{\pi_{k_1}}, X_{\pi_{k_2}}; X_{\pi_{n+l_{k_1}}}, X_{\pi_{n+l_{k_2}}}) \stackrel{(d)}{=} \tilde{U}_{n,m}^{\pi,L}. \quad (22)$$

This machinery provides a way to bound the tail probability of $U_{n,m}^{\pi}$. Indeed, for any $\lambda > 0$ and $t > 0$, by applying a Chernoff bound (Vershynin, 2018, Section 2.3) and Jensen's inequality (Vershynin, 2018, Section 1.6) together with (21), we obtain

$$\mathbb{P}_\pi(U_{n,m}^{\pi} > \lambda \mid \mathcal{X}_N) \leq e^{-\lambda t} \mathbb{E}_\pi[\exp(\lambda U_{n,m}^{\pi}) \mid \mathcal{X}_N] \leq e^{-\lambda t} \mathbb{E}_{\pi,L}[\exp(\lambda \tilde{U}_{n,m}^{\pi,L}) \mid \mathcal{X}_N].$$

Then, by (22), it follows that

$$\mathbb{E}_{\pi, L} \left[\exp \left(\lambda \tilde{U}_{n, m}^{\pi, L} \right) \middle| \mathcal{X}_N \right] = \mathbb{E}_{\pi, L, \zeta} \left[\exp \left(\lambda \tilde{U}_{n, m}^{\pi, L, \zeta} \right) \middle| \mathcal{X}_N \right].$$

As a result,

$$\mathbb{P}_{\pi} \left(U_{n, m}^{\pi} > \lambda \mid \mathcal{X}_N \right) \leq \mathbb{E}_{\pi, L, \zeta} \left[\exp \left(\lambda \tilde{U}_{n, m}^{\pi, L, \zeta} \right) \mid \mathcal{X}_N \right].$$

The right-hand side, which involves Rademacher averages, is a well-studied quantity that can be controlled using standard decoupling arguments (Vershynin, 2018, Chapter 6). This, in turn, provides an upper bound on the concentration of $U_{n, m}^{\pi}$ and allows control of the random critical value in the permutation U-statistic test.

Several works on MMD-based two-sample test use this approach by expressing their test statistic as the sum of a U-statistics with respect to the positive definite kernel g associated to the MMD and a controllable remainder term. In contrast, the sliced Wasserstein distance does not relate naturally to U-statistics, preventing a direct application of this approach.

While Kim et al. (2022) rely on the symmetrization trick of Dümbgen (1998) to address the dependence introduced by permutations, Albert (2019) takes a different route and establish a concentration bound for permuted sums, offering an alternative way to control the random quantile term. We describe this second approach in the next section.

14.2 Concentration of permuted sums

Continuing our review of methods for controlling the random critical value in permutation-based tests, we now turn to the work of Albert (2019). As shown in Albert et al. (2015, Section 2.2), the test statistic considered in Albert (2019) is a rescaled version of a U-statistic. Although their analysis focuses on independence testing, its contribution extends beyond this specific setting by introducing a distinct framework for analyzing random critical values and highlighting the key role of concentration inequalities in this task.

The main contribution of this work lies in Section 2.2, where the author establishes a concentration inequality for permuted sums in a general setting by leveraging fundamental inequalities for random permutations of Talagrand (1995). In particular, they present the following result

Lemma 31 (Theorem 2.1 of Albert (2019)) *Let $\{a_{i, j}\}_{1 \leq i, j \leq N}$ be a collection of any real numbers, and π be random uniform permutation in S_N . Consider $Z_N = \sum_{i=1}^N a_{i, \pi(i)}$. Then, for all $x > 0$,*

$$\mathbb{P} \left(|Z_N - \mathbb{E}[Z_N]| \geq 2 \sqrt{2 \left(\frac{1}{n} \sum_{i, j=1}^N a_{i, j}^2 \right)} + 2 \max_{1 \leq i, j \leq N} \{ |a_{i, j}| \} x \right) \leq 16e^{1/16} \exp \left(-\frac{x}{16} \right). \quad (23)$$

In the final section, the author illustrates the use of this theorem by analyzing the non-asymptotic behavior of a permutation-based independence test introduced in Albert et al. (2015). Before going into the details of how they applied the theorem, we now recall in detail their problem of interest and the corresponding test statistic.

Let $\mathbf{X}$ be a separable space, and let $\mathcal{X}_N = (X_1, \dots, X_N)$ be i.i.d. samples from a joint distribution P on $\mathbf{X}^2$, where each $X_i = (X_i^1, X_i^2)$ has marginals P^1 and P^2 corresponding to its coordinates. The goal is to test whether $P = P^1 \otimes P^2$. To this end, they consider the following test statistic, whose motivation is detailed in Albert et al. (2015):

$$T_{\delta}(\mathcal{X}_N) = \frac{1}{N-1} \left(\sum_{i=1}^N \varphi_{\delta}(X_i^1, X_i^2) - \frac{1}{n} \sum_{i=1}^N \sum_{j=1}^N \varphi_{\delta}(X_i^1, X_j^2) \right).$$

where φ_{δ} is a measurable real-valued function on $\mathbf{X}^2$ potentially depending on some unknown parameter δ .

Then, for any random permutation π uniformly distributed over $\{1, \dots, N\}$, the corresponding permuted sample is defined as

$$\mathcal{X}_N^{\pi} = (X_1^{\pi}, \dots, X_N^{\pi}), \quad \text{where } X_i^{\pi} = (X_i^1, X_{\pi(i)}^2), \quad \forall 1 \leq i \leq N,$$

and the associated permuted test statistic is given by

$$T_\delta(\mathcal{X}_N^\pi) = \frac{1}{N-1} \left(\sum_{i=1}^N \varphi_\delta(X_i^1, X_{\pi(i)}^2) - \frac{1}{n} \sum_{i=1}^N \sum_{j=1}^N \varphi_\delta(X_i^1, X_j^2) \right). \quad (24)$$

To apply their concentration bound, they express the permuted test statistic as the difference between a random variable and its expectation. In particular, by introducing $\tilde{Z}(\mathcal{X}_N) = \sum_{i=1}^N \varphi_\delta(X_i^1, X_{\pi(i)}^2)$, the permuted test statistic defined in (24) can be rewritten as

$$T_\delta(\mathcal{X}_N^\pi) = \frac{1}{N-1} \left(\tilde{Z}(\mathcal{X}_N) - \mathbb{E} \left[\tilde{Z}(\mathcal{X}_N) \mid \mathcal{X}_N \right] \right).$$

This reformulation aligns the left-hand side of the probability inequality in (23) with the permuted test statistic, thereby allowing the authors to directly apply the concentration result to derive an upper bound on its quantile. However, this approach cannot be extended to the case of the sliced Wasserstein distance, since the latter cannot be written as such a centered difference.

From the preceding technical review concerning the control of the conditional quantile of a permuted test statistic, two main insights can be drawn:

- Concentration inequalities are a powerful tool for analyzing conditional quantiles, as they provide sharp non-asymptotic high-probability bounds—typically with exponentially small tails—that lead to the desired logarithmic dependence on the Type I error level α , as discussed in Albert (2019, Section 3.2).
- Unlike U -statistics, which are centered at zero under permutation, or the statistic in Albert et al. (2015), which can be written as a deviation from its expectation, analyzing the permuted sliced Wasserstein statistic requires explicit control of its expectation under permutation of the samples. We address this challenge in Section 9.

15 GAUSSIAN MEAN SHIFT EXPERIMENT

Due to page constraints in the main paper, we present an additional experiment here. In this setting, samples are drawn from $\mu = \mathcal{N}(0, I_{60})$ and $\nu = \mathcal{N}(\mathbf{m}, I_{60})$, where $\mathbf{m} = (0.6, 0.6, 0, \dots, 0) \in \mathbb{R}^{60}$ is a 60-dimensional vector whose first two entries are equal to 0.6.

Following the experimental setup from Section 6.1, we compare the sliced Wasserstein (SW) tests with the Projected Wasserstein (PW) test (Wang et al., 2021a) and the MMD test (Gretton et al., 2012) using linear, Gaussian and Laplace kernels. As observed, the MMD test with a linear kernel shows a significant improvement in this setting. This improvement can be explained by the fact that for the linear kernel $k(x, y) = x^\top y$, the MMD test reduces to a mean-difference test, making it particularly well-suited for detecting to mean shifts.

It is also observed that, in this scenario, the SW-based tests exhibit lower performance compared to the three MMD-based tests. Nevertheless, the statistical power of SW tests can be considerably improved by increasing the number of projections. This observation motivates the extended investigation presented in Section 6.3.




Figure 6: Power vs. number of sample size: Gaussian mean shift scenario