Integral Probability Metrics Meet Neural Networks: The Radon-Kolmogorov-Smirnov Test

Seunghoon Paik¹, Michael Celentano¹, Alden Green², Ryan J. Tibshirani¹

¹University of California, Berkeley, ²Stanford University

Abstract

Integral probability metrics (IPMs) constitute a general class of nonparametric two-sample tests that are based on maximizing the mean difference between samples from one distribution P versus another Q, over all choices of data transformations f living in some function space \mathcal{F} . Inspired by recent work that connects what are known as functions of *Radon bounded variation* (RBV) and neural networks (Parhi and Nowak, 2021, 2023), we study the IPM defined by taking \mathcal{F} to be the unit ball in the RBV space of a given smoothness degree $k \geq 0$. This test, which we refer to as the *Radon-Kolmogorov-Smirnov* (RKS) test, can be viewed as a generalization of the well-known and classical Kolmogorov-Smirnov (KS) test to multiple dimensions and higher orders of smoothness. It is also intimately connected to neural networks: we prove that the witness in the RKS test—the function f achieving the maximum mean difference—is always a ridge spline of degree k, i.e., a single neuron in a neural network. We can thus leverage the power of modern neural network optimization toolkits to (approximately) maximize the criterion that underlies the RKS test. We prove that the RKS test has asymptotically full power at distinguishing any distinct pair $P \neq Q$ of distributions, derive its asymptotic null distribution, and carry out experiments to elucidate the strengths and weaknesses of the RKS test versus the more traditional kernel MMD test.

1 Introduction

In this paper, we consider the fundamental problem of nonparametric two-sample testing, where we observe independent samples $x_i \sim P$, i = 1, ..., m and $y_i \sim Q$, i = 1, ..., n, all samples assumed to be in \mathbb{R}^d , and we use the data to test the hypothesis

$$H_0: P = Q$$
, versus $H_1: P \neq Q$.

This is not only a problem of core interest in the traditional literature on hypothesis testing in statistics, it also has a key role in more modern applications in machine learning, such as out-of-distribution detection (Hendrycks and Gimpel, 2017), transfer learning (Long et al., 2015), generative modeling (Goodfellow et al., 2014), among others. *Nonparametric* two-sample testing in particular refers to the problem of testing P = Qas above, without parametric assumptions on the forms of the distributions P, Q. Although there are many nonparametric two-sample tests, we will focus on a class of tests called *integral probability metrics* (IPMs) (Müller, 1997), which measure the discrepancy between two sets of samples by taking the maximum difference in sample means over a function class \mathcal{F} :

$$\rho(P_m, Q_n; \mathcal{F}) = \sup_{f \in \mathcal{F}} |P_m(f) - Q_n(f)|.$$
(1)

Here $P_m = \frac{1}{m} \sum_{i=1}^m \delta_{x_i}$ is the empirical distribution and $P_m(f) = \frac{1}{m} \sum_{i=1}^m f(x_i)$ the corresponding empirical expectation operator based on x_i , i = 1, ..., m; and likewise for Q_n and $Q_n(f)$.

The discrepancy in (1) is quite general and certain choices of \mathcal{F} recover well-known distances that give rise to two-sample tests. A noteworthy univariate (d = 1) example is the *Kolmogorov-Smirnov* (KS) distance (Kolmogorov, 1933; Smirnov, 1948). Though it is more typically defined in terms of cumulative distribution functions (CDFs) or ranks, the KS distance can be seen as the IPM that results when \mathcal{F} is taken to be the class of functions with at most unit *total variation* (TV). Some attempts have been made to define a multivariate KS distance based on multivariate CDFs (Bickel, 1969) or ranks (Friedman and Rafsky, 1979), but the resulting tests have a few limitations, such as being expensive to compute or exhibiting poor power in practice. In fact, even the univariate KS test is known to be insensitive to certain kinds of alternatives, such as those $P \neq Q$ which differ in the tails (Bryson, 1974). This led Wang et al. (2014) to recently propose a higher-order generalization of the KS test, which uses the IPM based on a class of functions whose *derivatives* are of bounded TV.

In this work we propose and study a new class of IPMs based on a kind of multivariate total variation known as *Radon total variation* (RTV). The definition of RTV is necessarily technical, and we delay it until Section 2. For now, we remark that if d = 1 then RTV reduces to the usual notion of total variation, which means that our RTV-based IPM—the central focus of this paper—directly generalizes the KS distance to multiple dimensions. We refer to this RTV-based IPM as the *Radon-Kolmogorov-Smirnov* (RKS) distance, and the resulting two-sample test as the RKS test. We also will consider the IPM defined using a class of functions whose derivatives are of bounded RTV, which in turn generalizes the higher-order KS distance of Wang et al. (2014) to multiple dimensions. Sadhanala et al. (2019) recently showed that the higher-order KS test can be more sensitive to departures in the tails, and we will give empirical evidence that our higher-order RKS test can be similarly sensitive to tail behavior.

Our test statistic also has a simple but revealing characterization in terms of *ridge splines*. A ridge spline of integer degree $k \ge 0$ is defined for a direction $w \in \mathbb{S}^{d-1}$ and an offset $b \in \mathbb{R}$ as the truncated polynomial $f(x) = (w^{\mathsf{T}}x - b)_{+}^{k}$. Here we use \mathbb{S}^{d-1} to denote the unit ℓ_2 sphere in \mathbb{R}^d (the set of vectors with unit ℓ_2 norm), and we use the abbreviation $t_+ = \max\{t, 0\}$. We prove (Theorem 3) that a k^{th} degree ridge spline witnesses the k^{th} degree RKS distance: there exists a k^{th} degree ridge spline f with unit Radon total variation for which $P_m(f) - Q_n(f)$ equals the RKS distance. Thus the RKS test statistic can be written as

$$T_{d,k} = \max_{(w,b)\in\mathbb{S}^{d-1}\times[0,\infty)} \left| \frac{1}{m} \sum_{i=1}^{m} (w^{\mathsf{T}}x_i - b)_+^k - \frac{1}{n} \sum_{i=1}^{n} (w^{\mathsf{T}}y_i - b)_+^k \right|.$$
(2)

Ridge splines—and in particular the first-degree ridge spline $f(x) = (w^{\mathsf{T}}x - b)_+$, also known as the rectified linear (ReLU) unit—are fundamental building blocks of modern neural networks. From this perspective, the representation (2) says that the RKS distance is achieved by a single neuron in a two-layer neural network. In fact, the connections between Radon total variation and neural networks run deeper. The RKS distance can be viewed as finding the maximum discrepancy between sample means over all two-layer neural networks (of arbitrarily large width) subject to a weight decay constraint (Section 2). Although we focus throughout on two-layer neural networks, the IPM defined over deeper networks can be viewed as the RKS distance in the representation learned by the hidden layers.

The representation (2) also hints at some interesting statistical properties. Let $X \sim P_m$ and $Y \sim Q_n$ be random variables drawn from the empirical distributions of $\{x_i\}_{i=1}^m$ and $\{y_i\}_{i=1}^n$, respectively. The 0th degree RKS distance finds the direction $w \in \mathbb{S}^{d-1}$ along which the usual Kolmogorov-Smirnov distance—maximum gap in CDFs—between the univariate random variables $w^T X$ and $w^T Y$ is as large as possible. For $k \ge 1$, the k^{th} degree statistic finds the direction $w \in \mathbb{S}^{d-1}$ along which the higher-order KS distance—maximum gap in truncated k^{th} moments—between $w^T X$ and $w^T Y$ is maximized.

Thus, speaking informally, we expect the RKS test to be particularly sensitive to the case when the laws of $w^{\mathsf{T}}X$ and $w^{\mathsf{T}}Y$ differ primarily in just a few directions w (since we are taking a maximum over all such directions in (2)). In other words, we expect the RKS test to be sensitive to *anisotropic* differences between P and Q. On the other hand, taking a larger value of k results in higher-order truncated sample moments, which can be much more sensitive to small differences in the tails. Figure 1 gives an illustrative example. We consider two normal distributions with equal means and covariances differing in just one direction (along the x-axis). The figure displays the witnesses (ridge splines) for the RKS test when k = 0, 1, 2. For larger k, the witness is more aligned with the direction in which P and Q vary, and also places more weight on the tails.

Summary of contributions. Our main contributions are as follows.

• We propose a new class of two-sample tests using a k^{th} degree Radon total variation IPM. We prove a representer theorem (Theorem 3) which shows that the IPM is witnessed by a ridge spline (establishing the representation in (2)). This connects our proposal to two-sample testing based on neural networks, and helps with optimization.



Figure 1: Illustration of RKS tests for $P = \mathcal{N}_2(0, I)$ and $Q = \mathcal{N}_2(0, D)$, where D = diag(1.4, 1).

- Under the null P = Q, we derive the asymptotic distribution of the test statistic (Theorem 5).
- Under the alternative $P \neq Q$, we give upper bounds on the rate at which the test statistic concentrates around the population-level IPM. Using the fact that the population-level IPM is a metric, we then show that the RKS test is *consistent* against any fixed $P \neq Q$: it has asymptotic error (sum of type I and type II errors) tending to zero (Theorem 6).
- We complement our theory with numerical experiments to explore the operating characteristics of the RKS test compared to other popular nonparametric two-sample tests.

Related work. When \mathcal{F} is the unit ball in a reproducing kernel Hilbert space (RKHS), the IPM in (1) is known as a maximum mean discrepancy (MMD). Recently, there has been interest in studying multivariate nonparametric two-sample testing using kernel MMDs (Gretton et al., 2012), as well as energy distances (Baringhaus and Franz, 2004; Székely and Rizzo, 2005). The latter are in many cases equivalent to a kernel MMD (Sejdinovic et al., 2013). It is worth emphasizing that the Radon bounded variation space is *not* an RKHS, and as such, our test cannot be seen as a special case of a kernel MMD. There are also many other kinds of multivariate nonparametric two-sample tests, such as graph-based tests using kNN graphs (Schilling, 1986; Henze, 1988) or spanning-trees (Friedman and Rafsky, 1979). The RKS test statistic (2) avoids the need to construct a graph over the samples. However, we note that one could also use a graph to measure (approximate) multivariate total variation of a different variety (often called *measure-theoretic* TV, which is not the same as Radon TV), and one could then form a related two-sample test, accordingly.

From a different point of view, based on (2), the RKS test can be understood as a projection test which uses the higher-order KS test statistic (Wang et al., 2014; Sadhanala et al., 2019) as its base univariate tool. In general, projections are a natural way to lift univariate tools for statistical modeling (whether in testing, estimation, or prediction) to multivariate data settings. The literature is rich with contributions studying projection tests in two-sample settings, including Mueller and Jaakkola (2015); Ghosh and Biswas (2016); Wei et al. (2016); Paty and Cuturi (2019); Lin et al. (2020); Niles-Weed and Rigollet (2022). Our paper complements this line of work by connecting projection tests to IPMs over bounded variation spaces, which, to our knowledge, has not been done before.

In machine learning, two-sample tests play a fundamental role in generative adversarial networks (GANs) (Goodfellow et al., 2014), which have proven to be an effective way to generate new (synthetic) draws that adhere to the distribution of a given high-dimensional set of samples. In short, GANs are trained to produce synthetic data which a two-sample test cannot distinguish from "real" samples. However, optimizing a GAN is notoriously unstable. It has been suggested that the stability of GANs can be improved if the underlying two-sample test is based on an IPM, with constraints on, e.g., the Lipschitz constant (Arjovsky et al., 2017), RKHS norm (Li et al., 2017), or Sobolev norm (Mrouch et al., 2018) of the witness. In practice, a neural network is used to train the discriminator in the GAN, which only approximates the witness (the function achieving the supremum) in IPMs over RKHS or Sobolev spaces. Interestingly, for RBV spaces, this is not an approximation and a two-layer neural network captures the witness exactly, as shown in (2).

At a technical level, our paper falls into a line of work exploring neural networks from the functional analytic perspective. The equivalences between weight decay of parameters in a two-layer neural network and first-degree (k = 1) Radon total variation were discovered in Savarese et al. (2019); Ongie et al. (2020), and further formalized and extended by Parhi and Nowak (2021) to the higher-degree case (k > 1). We make use of these equivalences to prove results on the Radon TV IPM. Lastly, we note that the sensitivity of neural networks to variation in only a few directions was observed even earlier, in Bach (2017).

2 The Radon-Kolmogorov-Smirnov test

1

For independent samples $x_i \sim P$, i = 1, ..., m and $y_i \sim Q$, i = 1, ..., n, and a given integer $k \geq 0$, the RKS test statistic $T_{d,k}$ is defined as in (2). This statistic searches for a marginal (defined by projection onto the direction w) with the largest discrepancy in a truncated k^{th} moment. As with any two-sample test, we can use permutations to compute a finite-sample valid p-value,

$$p = \frac{\sum_{b=1}^{B} \mathbb{1}\left\{T_{d,k}(z_{\pi_b}) \ge T_{d,k}(z)\right\} + 1}{B+1}.$$
(3)

Here we use $z = (x_1, \ldots, x_m, y_1, \ldots, y_n) \in \mathbb{R}^{m+n}$ to denote the original data sequence, and we use $T_{d,k}(z)$ to emphasize that the statistic in (2) is computed on z. Moreover, each π_b is a permutation of $\{1, \ldots, m+n\}$, drawn uniformly at random (from the set of all possible permutations), and we use $T_{d,k}(z_{\pi_b})$ to denote the test statistic computed on the permuted sequence $z_{\pi_b} = (z_{\pi_b(1)}, \ldots, z_{\pi_b(m+n)})$. Given any user-chosen level $\alpha \in [0, 1]$, we reject $H_0: P = Q$ if $p \leq \alpha$. By well-known results on permutation tests (see, e.g., Hemerik and Goeman (2018), and references therein), this gives finite-sample type I error control, $\mathbb{P}_{H_0}(p \leq \alpha) \leq \alpha$.

Finding the value of w and b achieving the supremum in (2) is a nonconvex problem. However, it bears a connection to neural network optimization, which is of course among the most familiar and widely-studied nonconvex problems in machine learning. We show in Corollary 1 that problem (2) can be equivalently cast as finding the maximum discrepancy between sample means over all two-layer neural networks (of arbitrarily large width) under a weight decay-like constraint. In particular, if we denote by $f_{(a_j,w_j,b_j)_{j=1}^N}$ the N-neuron two-layer neural network defined by

$$f_{(a_j,w_j,b_j)}(x) = a_j (w_j^\mathsf{T} x - b_j)_+^k \quad \text{and} \quad f_{(a_j,w_j,b_j)_{j=1}^N}(x) = \sum_{j=1}^N f_{(a_j,w_j,b_j)}(x), \tag{4}$$

then the RKS test statistic is equivalently

$$T_{d,k} = \max_{\substack{f = f_{(a_j, w_j, b_j)_{j=1}^N} \\ \text{subject to}}} P_m(f) - Q_n(f)$$
(5)

This equivalence holds for any number of neurons $N \ge 1$, which is why we say that the network can be of "arbitrarily large width".

In Section 4, we use a Lagrangian form of this constrained optimization, allowing us to leverage existing deep learning toolkits to compute the RKS distance. As with neural network optimization in general, we cannot prove that any practical first-order scheme—such as gradient descent or any its variants—is able to find the global optimum of (5), or its Lagrange form. However, we find empirically that the test statistic resulting from first-order optimization has favorable behavior in practice, such as stability (with respect to the choice of learning rate) and strong power (with respect to certain classes of alternatives). Further, we recall that in practice we would typically use the permutation p-value (3), and as long as the same first-order optimization algorithm is applied to the original data sequence and to every permuted sequence, this still provides finite-sample type I error control. In effect, we can view this as *redefining* the RKS test to be the output of the given first-order algorithm we have at our disposal.

2.1 Functions of Radon bounded variation

The RKS distance, like the univariate KS distance, can be viewed as the maximum mean discrepancy with respect to a function class defined by a certain notion of smoothness. Whereas the univariate KS distance is

the IPM over functions with total variation is bounded by 1, the RKS distance is the IPM over functions with Radon total variation bounded by 1, subject to a boundary condition. This will be shown a bit later; first, we must cover preliminaries needed to understand Radon total variation.

Heuristically, the k^{th} degree Radon total variation of $f : \mathbb{R}^d \to \mathbb{R}$ measures the average smoothness of the "marginals" of f. To be more concrete, consider a function $f : \mathbb{R}^d \to \mathbb{R}$ which is infinitely differentiable, and whose derivatives of any order decay super-polynomially: $f \in C^{\infty}(\mathbb{R}^d)$ and $\sup_{x \in \mathbb{R}^d} |x^{\alpha} \partial^{\beta} f(x)| < \infty$. Here $\alpha = (\alpha_1, \ldots, \alpha_d)$ and $\beta = (\beta_1, \ldots, \beta_d)$ are multi-indices (with nonnegative integer elements), and we use the abbreviations

$$x^{\alpha} = \prod_{i=1}^{d} x_{i}^{\alpha_{i}}, \text{ and } \partial^{\beta} f = \partial_{x_{1}}^{\beta_{1}} \cdots \partial_{x_{d}}^{\beta_{d}} f.$$

The set of such functions is called the *Schwartz class*, and denoted by $\mathcal{S}(\mathbb{R}^d)$. The *Radon transform* of a function $f \in \mathcal{S}(\mathbb{R}^d)$ is itself another function $\mathcal{R}\{f\}: \mathbb{S}^{d-1} \times \mathbb{R} \to \mathbb{R}$ defined as

$$\mathcal{R}{f}(w,b) = \int_{w^{\mathsf{T}}x=b} f(x) \,\mathrm{d}x,$$

where the integral is with respect to the Lebesgue surface measure on the hyperplane $\{w^{\mathsf{T}}x = b\}$. As a function of b, the Radon transform can be viewed as the marginal of the function f in the direction w. The k^{th} degree Radon total variation of f is then defined to be (Parhi and Nowak, 2021, 2023):¹

$$\|f\|_{\mathrm{RTV}^{k}} = c_{d} \int_{\mathbb{S}^{d-1} \times \mathbb{R}} \left| \partial_{b}^{k+1} \Lambda^{d-1} \mathcal{R}\{f\}(w, b) \right| \mathrm{d}\sigma(w, b),$$

where σ is the Hausdorff measure on $\mathbb{S}^{d-1} \times \mathbb{R}$, $c_d = 1/(2(2\pi)^{d-1})$, and Λ^{d-1} is the "ramp filter" defined by $\Lambda^{d-1} = (-\partial_b^2)^{\frac{d-1}{2}}$ with fractional derivatives interpreted in terms of Riesz potentials (see Parhi and Nowak (2021, 2023) for details). Equivalently, if we define an operator by $R_k\{f\}(w,b) = c_d \partial_b^{k+1}(-\partial_b^2)^{\frac{d-1}{2}} \mathcal{R}\{f\}(w,b)$, then we can write k^{th} degree Radon TV of f simply as $||f||_{\text{RTV}^k} = ||R_k\{f\}||_{L^1}$. The mapping $f \mapsto ||f||_{\text{RTV}^k}$ is a seminorm, which we call the RTV^k seminorm.

Parhi and Nowak (2021, 2023) and Parhi (2022) extend the RTV^k seminorm beyond Schwartz functions using functional analytic tools based on duality. The space of functions with bounded RTV^k seminorm is referred to as the RBV^k space (where RBV stands for "Radon bounded variation"). After this extension, one can still think of RTV^k seminorm as measuring average higher-order smoothness of the marginals of f, but it does not require that all derivatives and integrals exist in a classical sense. We refer the reader especially to Parhi (2022) for the complete functional analytic definition of the RBV^k space. For our purposes, it is only important that RBV^k functions enjoy the representation established by the following theorem. Note that in order to define the RBV^k space, Parhi and Nowak (2021); Parhi (2022) must view it as a space of equivalence classes of functions under the equivalence relation $f \sim g$ if f(x) = g(x) for Lebesgue almost every $x \in \mathbb{R}^d$. We denote elements of RBV^k as f, and the functions they contain as $f \in f$.

Proposition 1 (Adaptation of Theorem 22 in Parhi and Nowak (2021); Theorem 3.8 in Parhi (2022)). Fix any $k \ge 0$. For each $f \in RBV^k$, there exists a representative $f \in f$ which satisfies for all $x \in \mathbb{R}^d$

$$f(x) = \int_{\mathbb{S}^{d-1} \times [0,\infty)} (w^{\mathsf{T}} x - b)^k_+ \,\mathrm{d}\mu(w,b) + q(x),$$

for some finite signed Borel measure μ on $\mathbb{S}^{d-1} \times [0, \infty)$ satisfying $\|\mu\|_{\mathrm{TV}} = \|\mathbf{f}\|_{\mathrm{RTV}^k}$, and some polynomial q of degree at most k, where we use $\|\cdot\|_{\mathrm{TV}}$ for the total variation norm in the sense of measures.

Proposition 1 is the key result connecting the RBV^k space to two-layer neural networks. It states that any RBV^k functions is equal almost everywhere to the sum of a (possibly infinite-width) two-layer neural network and a polynomial. Furthermore, the ℓ_1 norm of coefficients in the last layer of this neural network (formally, the total variation norm $\|\mu\|_{\text{TV}}$ of the measure μ defined over the coefficients) equals the RBV^k

¹A remark on notation and nomenclature: we use RBV^k to refer to the space which Parhi and Nowak (2021); Parhi (2022); Parhi and Nowak (2023) instead call RBV^{k+1}. That is, we index based on the *degree* of the underlying ridge spline, whereas the aforementioned papers index based on the *order*, traditionally defined as the degree plus one, m = k + 1.

seminorm of the function in question. As stated, Proposition 1 is a slight refinement (and simplification) of results in Parhi and Nowak (2021); Parhi (2022), which we show in Appendix A.

Parhi and Nowak (2021) consider the nonparametric regression problem defined by minimizing, over all functions f, the squared loss incurred by f with respect to a given finite data set plus a penalty on $||f||_{\text{RTV}^k}$. They show this is solved by the sum of a *finite-width* two-layer neural network and a polynomial, and hence provide a representation theorem for RTV^k -penalized nonparametric regression analogous to that for kernel ridge regression (Wahba, 1990). In a coming subsection, we will establish a similar connection in the context of two-sample testing: the IPM under an RTV^k constraint is realized by a single neuron.

Before moving on, it is worth providing some intuition for the results just discussed. First, the operator R_k should be understood as transforming a function f into the Radon domain: namely, it specifies f via the higher-order derivatives of its marginals. Then, the RTV^k acts as the L^1 norm in the Radon domain. Lastly, the reason two-layer neural networks and ridge splines emerge as the solutions to nonparametric regression and maximum mean discrepancy problems under an RTV^k constraint is that ridge splines are *sparse* in the Radon domain. In particular, for $f(x) = (w^{\mathsf{T}}x - b)_{+}^k$, we have (Parhi and Nowak, 2021):

$$R_k\{f\}(w,b) = \frac{1}{2} \left(\delta_{(w,b)} + (-1)^{k+1} \delta_{(-w,-b)} \right),$$

where $\delta_{(\pm w,\pm b)}$ denote point masses at $(\pm w,\pm b)$. Thus, RTV^k-penalized optimization problems are solved by two-layer neural networks due to the tendency of the L^1 norm to encourage sparse solutions.

2.2 Pointwise evaluation of RBV functions

Recall, the RBV^k space contains equivalence classes of functions whose members differ on sets of Lebesgue measure zero. Therefore, at face value, point evaluation is meaningless for elements of the RBV^k space. This poses a problem for defining an IPM with respect to RBV functions, since the IPM depends on empirical averages, which require function evaluations over samples. Fortunately, we show that there is a natural way to resolve point evaluation over the RBV^k space: each equivalence class in RBV^k has, possibly subject to a mild boundary condition, a unique representative which satisfies a suitable notion of continuity. We can then define point evaluation with respect to this representative.

For the case k = 0, we need to introduce a new notion of continuity, which we call *radial cone continuity*. For $k \ge 1$, we rely on the standard notion of continuity. Radial cone continuity is defined as follows.

Definition 1. A function f is said to be radially cone continuous if it is continuous at the origin, and for any $x \in \mathbb{R}^d \setminus \{0\}$, we have $f(x + \epsilon v) \to f(x)$ whenever (i) $\epsilon \to 0$ and (ii) $v \in \mathbb{S}^{d-1}$ and $v \to x/||x||_2$.

Radial cone continuity can be viewed as a generalization of right-continuity to higher dimensions. Indeed a consequence of radial cone continuity is that for every $v \in \mathbb{R}^d$, the restriction of the function f to the ray emanating from the origin in the direction v is right-continuous (i.e., $t \mapsto f(tv)$, t > 0 is right-continuous). The concept of radial cone continuity, however, is stronger than right-continuity along rays—it additionally requires continuity along any path which is tangent to and approaches x from the same direction as the ray that points from the origin to x. The following theorem shows that (subject to a boundary condition for the case k = 0) every equivalence class $f \in \text{RBV}^k$ contains a unique continuous representative when $k \ge 1$, and radially cone continuous representative when k = 0.

Theorem 1. For k = 0, if $f \in RBV^0$ contains a function that is continuous at the origin, then it contains a unique representative that is radially cone continuous. Moreover, for $k \ge 1$, each $f \in RBV^k$ contains a unique representative that is continuous and this representative is in fact (k - 1)-times continuously differentiable.

Theorem 1, which as far as we can tell is a new result for RBV spaces, is proved in Appendix B. This theorem motivates the definition of the function spaces:

$$RBV_{c}^{0} = \left\{ f : f \in f \text{ for some } f \in RBV^{0}, f \text{ radially cone continuous} \right\},$$

$$RBV_{c}^{k} = \left\{ f : f \in f \text{ for some } f \in RBV^{k}, f \text{ continuous} \right\}, \quad k \ge 1.$$
(6)

Unlike RBV^k , the space RBV_c^k contains functions rather than equivalence classes of functions, so that point evaluation is well-defined. The Radon total variation of a member $f \in \mathsf{f}$ is defined in the natural way as the Radon total variation of the function class in RBV^k to which it belongs.

2.3 The RKS distance is an IPM

We are ready to show that the RKS distance is equivalent to the IPM over functions f with $||f||_{\text{RTV}^k} \leq 1$, subject to a boundary condition. In this sense, it can be viewed as a generalization of the classical univariate KS distance, which, as we have mentioned, is the IPM over the space of functions with total variation at most 1. Precisely, the RKS distance is the IPM over the function space

$$\mathcal{F}_{k} = \Big\{ f \in \mathrm{RBV}_{c}^{k} : \|f\|_{\mathrm{RTV}^{k}} \leq 1, \ \partial^{\alpha} f(0) \text{ exists and equals 0 for all } |\alpha| \leq k \Big\}.$$
(7)

The boundary condition $\partial^{\alpha} f(0) = 0$ for all multi-indices $\alpha = (\alpha_1, \ldots, \alpha_d)$ with $|\alpha| = \sum_{j=1}^d \alpha_j \leq k$ is needed to guarantee that the IPM is finite. We can interpret this condition intuitively as requiring the polynomial q to be zero in the representation in Proposition 1. Otherwise, q would be able to grow without bound under the constraint $||f||_{\text{RTV}^k} \leq 1$, which we could then use to drive the IPM criterion to ∞ , provided that P and Q had any moment differences (in their first k moments).

The following result proves an important representation for functions in (7). Its proof is in Appendix C.

Theorem 2. Fix any $k \ge 0$. For any $f \in \mathcal{F}_k$, there exists a finite, signed Borel measure μ on $\mathbb{S}^{d-1} \times [0, \infty)$ such that for all $x \in \mathbb{R}^d$,

$$f(x) = \int_{\mathbb{S}^{d-1} \times [0,\infty)} (w^{\mathsf{T}} x - b)_+^k \,\mathrm{d}\mu(w,b),$$

and $\|\mu\|_{\mathrm{TV}} = \|f\|_{\mathrm{RTV}^k}$. When k = 0, the measure μ may be taken to be supported on $\mathbb{S}^{d-1} \times (0, \infty)$.

Note that the representation provided by Theorem 2 is similar to that in Proposition 1, except that the former explicitly deals with point evaluation and sets the polynomial term to q = 0, which is a consequence of the boundary condition imposed in the definition of \mathcal{F}_k . (Theorem 2 also appears similar to results in Parhi and Nowak (2021), but differs from their results for the same reasons.)

The next theorem establishes the equivalence between $T_{d,k}$ as defined in (2) and the IPM over \mathcal{F}_k . Its proof is in Appendix D.

Theorem 3. Fix any $k \ge 0$. Define

$$\mathcal{G}_k = \left\{ (w^{\mathsf{T}} \cdot - b)_+^k : (w, b) \in \mathbb{S}^{d-1} \times [0, \infty) \right\},\$$

where by convention we take $t^0_+ = \mathbb{1}\{t \ge 0\}$. Then for any P and Q with finite k^{th} moments, we have

$$\rho(P,Q;\mathcal{F}_k) = \rho(P,Q;\mathcal{G}_k).$$

In particular, this implies that for the empirical distributions P_m and Q_n over any sets of samples $\{x_i\}_{i=1}^m$ and $\{y_i\}_{i=1}^n$, we have $\rho(P_m, Q_n; \mathcal{F}_k) = T_{d,k}$, where the latter is as defined in (2).

Lastly, as a consequence of the above result, we establish the representation claimed earlier in (5). Its proof is simple enough that we present it here.

Corollary 1. Fix any $k \ge 0$. For the empirical distributions P_m and Q_n over any sets of samples $\{x_i\}_{i=1}^m$ and $\{y_i\}_{i=1}^n$, the RKS test statistic $T_{d,k}$ defined in (2) is equivalent to that in (4), (5), for any $N \ge 1$.

Proof. For $||w||_2 = 1$ and $b \in \mathbb{R}$, Parhi and Nowak (2021) show that $||(w^{\mathsf{T}} \cdot -b)_+^k||_{\mathrm{RTV}^k} = 1$. Then, for any $w \in \mathbb{R}^d$ and $a \in \mathbb{R}$, by the k^{th} order homogeneity of the ridge splines, $||a(w^{\mathsf{T}} \cdot -b)_+^k||_{\mathrm{RTV}^k} = |a|||w||_2^k$. By the triangle inequality,

$$\|f_{(a_j,w_j,b_j)_{j=1}^N}\|_{\mathrm{RTV}^k} \le \sum_{j=1}^N |a_j| \|w_j\|_2^k.$$

Thus, defining

$$\mathcal{F}_{k,N} = \left\{ f_{(a_j,w_j,b_j)_{j=1}^N} : (a_j,w_j,b_j) \in \mathbb{R} \times \mathbb{R}^d \times (0,\infty), \ \sum_{j=1}^N |a_j| \|w_j\|_2^k \le 1 \right\},$$

we have $\mathcal{F}_{k,N} \subseteq \mathcal{F}_k$. Defining also $\mathcal{G}_k^+ = \{(w^{\mathsf{T}} \cdot -b)_+^k : (w,b) \in \mathbb{S}^{d-1} \times (0,\infty)\}$, we have $\mathcal{G}_k^+ \subseteq \mathcal{F}_{k,N}$, so

$$\sup_{f \in \mathcal{G}_k^+} |P_m(f) - Q_n(f)| \le \sup_{f \in \mathcal{F}_{k,N}} |P_m(f) - Q_n(f)| \le \sup_{f \in \mathcal{F}_k} |P_m(f) - Q_n(f)|.$$

The leftmost supremum is unchanged if we include the origin, i.e., take the supremum over $f \in \mathcal{G}_k$, which precisely defines $T_{d,k}$. By Theorem 3, this is equal to $\sup_{f \in \mathcal{F}_k} |P_m(f) - Q_n(f)|$, the rightmost supremum above, so we conclude that all inequalities are equalities, thus $T_{d,k} = \sup_{f \in \mathcal{F}_{k,N}} |P_m(f) - Q_n(f)|$. Finally, we can remove the absolute value sign in the criterion because $\mathcal{F}_{k,N}$ is symmetric $(f \in \mathcal{F}_{k,N} \iff -f \in \mathcal{F}_{k,N})$, which completes the proof of (5) as an equivalent representation.

2.4 The RKS distance identifies the null hypothesis

An important property of the RKS distance for the purposes of two-sample testing is its ability to distinguish any two distinct distributions. In particular, we have the following result, whose proof is in Appendix E.

Theorem 4. For any P, Q with finite k^{th} moments, $\rho(P, Q; \mathcal{F}_k) = 0$ if and only if P = Q, and is positive otherwise.

Theorem 4 states that the RKS distance, in the population, perfectly distinguishes the null $H_0: P = Q$ from the alternative $H_1: P \neq Q$. Thus, the function space \mathcal{F}_k is rich enough to make the IPM a *metric* at the population level. As an interesting contrast, we remark that this is not true of the kernel MMD with a polynomial kernel of degree k, as studied in Gretton et al. (2012). Indeed, if P and Q have the same moments up to order k (where "moments" is to be interpreted in the appropriate multivariate sense) then the kernel MMD in the population between P and Q is 0. In other words, we can see that the use of *truncated* moments, as given by averages of ridge splines, is the key to the discrepancy being a metric.

3 Asymptotics

Given any probability distribution P on \mathbb{R}^d with finite moments of order $2k + \Delta$ for any fixed $\Delta > 0$, let us define the corresponding Gaussian process $\mathbb{G}_P = \{G_{w,b} : (w,b) \in \mathbb{S}^{d-1} \times [0,\infty)\}$ by

$$G_{w,b} \sim \mathcal{N}\Big(0, \mathbb{E}_{x \sim P}\Big[(w^{\mathsf{T}}x - b)^{2k}_{+}\Big]\Big), \quad \operatorname{Cov}(G_{w,b}, G_{w',b'}) = \mathbb{E}_{x \sim P}\Big[(w^{\mathsf{T}}x - b)^{k}_{+}(w'^{\mathsf{T}}x - b')^{k}_{+}\Big]. \tag{8}$$

The importance of this Gaussian process is explained in the next result, which shows that its supremum provides the asymptotic null distribution of the RKS test statistic.

Theorem 5. Fix any $k \ge 0$. Assume that P has finite moments of order $2k + \Delta$, for any fixed $\Delta > 0$. If k = 0, then additionally assume P is absolutely continuous with respect to Lebesgue measure, and satisfies

$$\sup_{w\in\mathbb{S}^{d-1}} \left\| p_{w^{\mathsf{T}}x} \right\|_{\infty} < \infty,\tag{9}$$

where $p_{w^{\mathsf{T}_x}}$ denotes the density of w^{T_x} for $x \sim P$, and $||f||_{\infty} = \sup_{t \in \mathbb{R}} |f(t)|$ denotes the sup norm of f. The condition (9) holds when P has a bounded density and bounded support, but also holds for many distributions without bounded support, such as the multivariate Gaussian. Under these conditions, when P = Q, we have

$$\sqrt{\frac{mn}{m+n}} T_{d,k} \xrightarrow{d} \sup_{(w,b) \in \mathbb{S}^{d-1} \times [0,\infty)} |G_{w,b}|, \quad as \ m, n \to \infty.$$

Moreover, with a proper rejection threshold, the RKS test can have asymptotically zero type I error and full power against any fixed $P \neq Q$.

Theorem 6. Fix any $k \ge 0$. Assume that P, Q each have finite moments of order $2k + \Delta$, for any fixed $\Delta > 0$. If k = 0, then additionally assume P, Q each have densities which satisfy (9). Let $t_{m,n}$ be such that $t_{m,n} \to 0$ and $t_{m,n}\sqrt{m+n} \to \infty$ as $m, n \to \infty$. Then the test which rejects when $T_{d,k} > t_{m,n}$ rejects with asymptotic probability 0 if P = Q and asymptotic probability 1 if $P \neq Q$, as $m, n \to \infty$.

We prove Theorems 5 and 6 in Appendices F and G. Roughly, Theorem 5 follows from general uniform central limit theorem results in Dudley (2014), after controlling the bracketing number of the function class \mathcal{G}_k . Theorem 6 follows from tail bounds that we establish for the concentration of the RKS distance around its population value, combined with the fact that the population distance is a metric (Theorem 4).

Let us pause to discuss the above theorems and how they fit into our understanding of the RKS test at large. While these results answer (what in our view are) standard and yet fundamental questions about the asymptotic behavior of the RKS test, neither really guides practical usage of the test. If the asymptotic null distribution in Theorem 5 had been *pivotal*, i.e., not depending on P, then we could have used this (at least in principle, ignoring computational issues) to approximately control type I error in large samples. However, the asymptotic null appears to depend on P except in the special case when k = 0 and d = 1, which recall, coincides with classical KS test, whose asymptotic null distribution is well-known to be pivotal (defined in terms of the supremum of a Brownian bridge).

In practice, as discussed at the start of Section 2, we would typically use permutations to calibrate the RKS test, controlling type I error in finte-sample. Our next result shows that calibration via permutations can indeed also achieve asymptotically zero type I error and full power against any fixed $P \neq Q$.

Theorem 7. Fix any $k \ge 0$. Assume the conditions of Theorem 6 on P, Q, and assume further that each has bounded support. Consider the permutation p-value in (3), for B + 1 = (m + n)!, and let $\alpha_{m,n}$ be such that $\alpha_{m,n} \to 0$ and $\log(1/\alpha_{m,n})\sqrt{1/m + 1/n}$ as $m, n \to \infty$. Then the test which rejects when $p > \alpha_{m,n}$ rejects with asymptotic probability 0 if P = Q and asymptotic probability 1 if $P \neq Q$, as $m, n \to \infty$.

The proof of Theorem 7 is given in Appendix H. It follows techniques developed in Green et al. (2024).

4 Experiments

We conduct experiments to compare the power of the RKS test with other multivariate nonparametric tests in multiple settings (different pairs of P and Q), both to investigate the strengths and weaknesses of the RKS test, and to investigate its behavior as we vary the smoothness degree k.

4.1 Computation of the RKS distance

First, we discuss computation of the RKS distance, as initially defined in (2). This a difficult optimization problem because of the nonconvexity of the criterion and the unit ℓ_2 norm constraint on w. Though we will not be able to circumvent the challenge posed by nonconvexity entirely, we can ameliorate it by moving to the overparametrized formulation in (5): this now casts the same computation as an optimization over the *N*-neuron neural network f in (4), subject to a constraint on $\sum_{j=1}^{N} |a_j| ||w_j||_2^k$ (which is sometimes called the "path norm" of f). We call this problem "overparametrized" since any individual pair (w_j, b_j) would be sufficient to obtain the global optimum in (2), and yet we allow the optimization to search over N such pairs $(w_j, b_j), j = 1, \ldots, N$. Perhaps unsurprisingly, we find that overparametrization (larger N) generally makes optimization easier—more stable with respect to the initialization and choice of learning rate, discussed in more detail in Appendix K.

Problem (5) is still more difficult than we would like it to be, i.e., not within scope for typical first-order optimization techniques, due to the path norm constraint $\sum_{j=1}^{N} |a_j| ||w_j||_2^k \leq 1$. Fortunately, it is easy to see that any bound here suffices to compute the RKS distance: we can instead use $\sum_{j=1}^{N} |a_j| ||w_j||_2^k \leq C$, for any C > 0, rescale the criterion $P_m(f) - Q_n(f)$ by 1/C, and then the value of the criterion at its optimum (the IPM value) is unchanged. This motivates us to instead solve a Lagrangian version of (5):

$$\min_{\substack{f=f_{(a_j,w_j,b_j)_{j=1}^N}\\b_j\ge 0,\,j=1,\dots,N}} -\frac{1}{k}\log\left(\frac{1}{N}\Big|P_m(f) - Q_n(f)\Big|\right) + \frac{\lambda}{kN}\sum_{i=1}^N |a_i| \|w_i\|_2^k,\tag{10}$$

which is the focus in all of our experiments henceforth. The log transform for the mean difference in (10) is used because we find that it leads to greater stability with respect to the choice of learning rate, which we cover in Appendix J. Importantly, the choice of Lagrangian parameter $\lambda > 0$ in (10) can be arbitrary, by the same rescaling argument as explained above: given any (local) optimizer f^* in (10), we can simply return the mean difference over the rescaled witness $f^*/||f^*||_{\mathrm{RTV}^k}$, where $||f^*||_{\mathrm{RTV}^k} = \sum_{i=1}^N |a_i^*|| |w_i^*||_2^k$. For $k \ge 1$, we apply the torch.optim.Adam optimizer (a variation on gradient descent), as implemented

in PyTorch, to (10). We use a betas parameter (0.9, 0.99), learning rate 0.5, number of iterations T = 200. penalty parameter $\lambda = 1$, and number of neurons N = 10. To enforce the nonnegativity condition on b, we project b to $[0,\infty)$ after each gradient step. Rather than take the last iterate, we choose the maximal IPM values among the iterates (after rescaling by the RTV^k seminorm of each iterate so that it lies in the unit seminorm ball). Further, we repeat this over three random initializations, and select the best resultant IPM value to be the final output. While this optimization scheme does not come with guarantees on convergence to a minimizer of the criterion (10), we reiterate that if we calibrate this test statistic using the permutation null (i.e., we use permutations to calculate a p-value as in (3), where in each permutation we apply the same optimization scheme), then we are still guaranteed to control type I error in finite-sample.

For k = 0, such a first-order scheme is not applicable due to the fact that the gradient of the 0th degree ridge spline $(w^{\mathsf{T}}x - b)^0_+ = \mathbb{1}\{w^{\mathsf{T}}x \ge b\}$ (with respect to w and b) is almost everywhere zero. As a surrogate, we directly approximate the optimum (w^*, b^*) in (2) using logistic regression where the class labels identify samples from P versus Q, as implemented in sklearn.linear_model.LogisticRegression in Python.

For concreteness, we summarize our computational approach below in Algorithm 1.

Algorithm 1 RKS test statistic

Input: Samples $\{x_i\}_{i=1}^m, \{y_i\}_{i=1}^n \subseteq \mathbb{R}^d$; smoothness degree $k \ge 0$; number of neurons $N \ge 1$; number of iterations $T \ge 1$; learning rate $\eta > 0$; regularization parameter $\lambda > 0$

Output: RKS test statistic $T_{d,k}$

Run:

Step 1 (optional). Center the data to have sample mean zero:

 $\widehat{\mu} = \frac{1}{m+n} \left(\sum_{\ell=1}^{m} x_{\ell} + \sum_{\ell=1}^{n} y_{\ell} \right); x_i \leftarrow x_i - \widehat{\mu}; y_i \leftarrow y_i - \widehat{\mu}$

 $\frac{\text{Step 2} (\text{optional}). \text{ Scale the data to have average } \ell_2 \text{ norm of 1:}}{\hat{L}^2 = \frac{1}{m+n} (\sum_{\ell=1}^m \|x_\ell\|_2^2 + \sum_{\ell=1}^n \|y_\ell\|_2^2); x_i \leftarrow x_i/\hat{L}; y_i \leftarrow y_i/\hat{L}}$

Step 3. Apply optimization procedure:

if $k \ge 1$ then

Run Adam on the criterion in (10) for T iterations with learning rate η to compute f^*

else

Run logistic regression on data $\{z_i\}_{i=1}^{m+n} = \{x_i\}_{i=1}^m \cup \{y_i\}_{i=1}^n$, with class labels indicating whether each $z_i \in \{x_\ell\}_{\ell=1}^m$ or $z_i \in \{y_\ell\}_{\ell=1}^n$. Given the solution (w^*, b^*) , define $f^*(x) = \mathbb{1}\{w^{*\mathsf{T}}x \ge b\}$ end if

Step 4. Rescale $f^*/||f^*||_{\mathrm{RTV}^k}$, where $||f^*||_{\mathrm{RTV}^k} = \sum_{i=1}^N |a_i^*|||w_i^*||_2^k$, and compute

$$T_{d,k} = |P_m(f^*) - Q_n(f^*)|$$

return $T_{d,k}$

Experimental setup 4.2

We compare the RKS tests of various degrees across several experimental settings, both to one another and to other all-purpose nonparametric tests: the energy distance test (Székely and Rizzo, 2005) as well as kernel MMD (Gretton et al., 2012) with a Gaussian kernel. Python code to replicate all of our experimental results is available at https://github.com/100shpaik/.

We fix the sample sizes to m = n = 512 throughout, and study four choices of dimension: d = 2, 4, 8, 16. For each dimension d, we consider five settings for P, Q, which are described in Table 1. In each setting, the parameter v controls the discrepancy between P and Q, but its precise meaning depends on the setting. The settings were broadly chosen in order to study the operating characteristics of the RKS test when differences between P and Q occur in one direction (settings 1–4), and in all directions (setting 5). Among the settings in which the differences occur in one direction, we also investigate different varieties (settings 1 and 2: mean shift under different geometries, setting 3: tail difference, setting 4: variance difference). Figure 2 visualizes samples from drawn each task in d = 2 dimensions, whereas Figure 3 exaggerates the deviation between P, Q(larger values of v) to better illustrate the geometry. Finally, we note that since the RKS test is *rotationally invariant*,² the fact that the chosen differences in Table 1 are axis-aligned is just a matter of convenience, and the results would not change if these differences instead occurred along arbitrary directions in \mathbb{R}^d .

Setting	Р	Q	v	
Pancake mean shift	one axis $\sim \mathcal{N}(0, 1)$	one axis $\sim \mathcal{N}(v, 1)$	0.3	
("pancake-shift")	the others $\sim \mathcal{N}_{d-1}(0,9I)$	the others $\sim \mathcal{N}_{d-1}(0,9I)$		
Ball mean shift		one axis $\sim \mathcal{N}(v, 1)$ the others $\sim \mathcal{N}_{d-1}(0, I)$		
("ball-shift")	$\mathcal{N}_d(0, I)$			
t coordinate		one axis $\sim t(v)$ the others $\sim \mathcal{N}_{d-1}(0, I)$		
("t-coord")	$\mathcal{N}_d(0, I)$			
Variance inflation		one arise $\Lambda(0, u)$		
in one direction	$\mathcal{N}_d(0, I)$	one axis $\sim \mathcal{N}(0, v)$	1.4	
("var-one")		the others $\sim \mathcal{N}_{d-1}(0, I)$		
Variance inflation				
in all directions	$\mathcal{N}_d(0, I)$	$\mathcal{N}_d(0, vI)$	1.2	
("var-all")				

Table 1: Experimental settings. Here $\mathcal{N}_d(\mu, \Sigma)$ means the *d*-dimensional normal distribution with mean μ and covariance Σ , and t(v) means the *t* distribution with *v* degrees of freedom.



Figure 2: Samples drawn from the settings in Table 1, with d = 2.



Figure 3: Samples drawn from the settings in Table 1, with d = 2, and exeggerated values of v.

²Meaning, $\rho(P_m, Q_n; \mathcal{F}_k)$ is invariant to rotations of the underlying samples $\{x_i\}_{i=1}^m$ and $\{y_i\}_{i=1}^n$ by any arbitrary orthogonal transformation $U \in \mathbb{R}^{d \times d}$. This is true because the RTV^k seminorm is itself invariant to rotations of the underlying domain, as can be seen directly from the representation in Proposition 1.

For the RKS tests, we examine smoothness degrees k = 0, 1, 2, 3, and we center the input data to have the sample mean zero jointly across both samples. That is, each x_i is replaced by $x_i - \frac{1}{m+n} \left(\sum_{\ell=1}^m x_\ell + \sum_{\ell=1}^n y_\ell \right)$, and likewise for y_i . This makes the RKS test translationally-invariant. For kernel MMD, we set the Gaussian kernel bandwidth according to a standard heuristic, based on the median of all pairwise squared distances between the input samples (Caputo et al., 2002). (The energy distance test has no tuning parameters.) For each setting, we compute these test statistics under the null where each x_i and y_i are sampled i.i.d. from the mixture $\frac{m}{m+n}P + \frac{n}{m+n}Q$, and under the alternative where x_i are i.i.d. from P and y_i from Q. We then repeat this 100 times (draws of samples, and computation of test statistics), and trace out ROC curves—true positive versus false positive rates—as we vary the rejection threshold for each test.

In Appendix L, we also compare the RKS test to kernel MMD with a polynomial kernel. An important remark is that the latter is not actually a nonparametric test, for any finite polynomial degree, because the population IPM does not define a metric (it cannot distinguish distributions P, Q that agree up to some number of moments, interpreted in the multivariate sense, but differ otherwise). We include comparisons to polynomial kernels for completeness nonetheless.

4.3 Experimental results

Figure 4 displays the ROC curves from all tests across all settings (rows) and all dimensions (columns). We see that in the first four settings (first four rows), the RKS tests exhibit generally strong performance: in each case, one of the k^{th} degree RKS tests performs either the best or near-best among the methods considered. This supports the intuition discussed earlier, because in each of these settings, P and Q differ in only a small number (in fact, one) direction. Conversely, in the last setting, "var-all", the distributions P and Q differ by a small amount in all directions, and the Gaussian kernel MMD and energy distance tests outperform the RKS tests, especially for larger d.

Recalling that larger choices of k correspond to comparing higher-order truncated moments, we can look further into the first four tasks (and their underlying data models) in order to try to learn something about why and when different degree RKS tests perform better or worse. In the "pancake-shift" setting, the mean shift occurs in a low-variance direction, which leads to nearly separable samples. The smallest choice k = 0performs best (and also better than kernel-based tests). In the "var-one" and "t-coord" settings, P, Q have the same mean, but different variance or tails, respectively, in one direction. We see that larger choices of k perform better: k = 1 for "t-coord", and k = 2, 3 for "var-one". In the "ball-shift" setting, which is a mean shift between isotropic Gaussians, all degrees k = 0, 1, 2 are reasonably competitive.

As the dimension d increases (across the columns), the performance of each test generally gets worse, except for "var-all", which will be explained shortly. However, when performance decreases, it does so at a rate which depends on the task and the test. For example, in the "pancake-shift" task, the performance of the RKS test with k = 0 decreases quite slowly as d increases, whereas all other methods quickly lose power. (In experiments not shown, RKS with k = 0 continues to have nontrivial power in this setting even when dis in the hundreds.) On the other hand, in the "var-one" task, all tests—including the leading RKS tests with k = 2, 3—noticeably lose power by the last column, d = 16. Finally, in "var-all", the problem actually gets *easier* as d increases, recalling that v parametrizes a variance inflation in each of the d total directions. Indeed, we can see that the kernel MMD and energy distance tests improve as d increases.

4.4 Aggregating RKS tests

Now we discuss how to aggregate the RKS test statistics, across various degrees k. In an ideal case, such an aggregate RKS test could leverage the strengths of different degrees for different settings, and have more robust power than any individual degree itself.

Suppose we consider RKS tests for degrees k = 0, ..., K, for some choice $K \ge 0$ of maximum degree; in most practical scenarios, it seems reasonable to choose to K = 3. A classical approach for aggregating these RKS tests would be Bonferroni correction, which takes p-values p_k , k = 0, ..., K, and returns

$$p_{\min} = (K+1) \cdot \min_{k=0,\dots,K} p_k$$

as an aggregate p-value. The controls type I error by the union bound: $\mathbb{P}_{H_0}((K+1) \cdot \min_{k=0,\ldots,K} p_k \leq \alpha) \leq \sum_{k=0}^{K} \mathbb{P}_{H_0}(p_k \leq \alpha/(K+1)) \leq \alpha.$



Figure 4: ROC curves across the experimental settings described in Table 1. Each row represents a different setting, and each column a different dimension. The RKS, kernel MMD ("KMMD") with a Gaussian kernel, and energy distance tests are compared. Each is coded by a combination of color and line type.

A variety of other aggregation strategies are available; for example,

$$p_{\mathrm{avg}} = \frac{2}{K+1} \sum_{k=0}^{K} p_k$$

is also a valid p-value. Table 1 of Vovk and Wang (2020) lists a number of other examples. Interestingly, if the p-values for individual RKS tests are computed via permutations (as have advocated for generally in this paper), then we can take any aggregation function G, and calibrate the corresponding p-value

$$G(p_0,\ldots,p_K)$$

itself via permutations at no additional computational cost, provided the same permutations π_b , $b = 1, \ldots, B$ were used to compute each p-value p_k , and these permutations form a group. This can be explained as follows (similar ideas appear in Solmi and Onghena (2014)). Abbreviating $T(b, k) = T_{d,k}(z_{\pi_b})$ for the test statistic for each b, k, and $T(0, k) = T_{d,k}(z)$ for the test statistic on the original sequence, it helps to visualize the matrix of available test statistics:

T(0,0)	T(0,1)	 T(0,K)
T(1,0)	T(1, 1)	 T(1, K)
•		
T(B,0)	T(B,1)	 T(B,K)

The permutation p-value, per column (per degree), is given by the following computation:

$$p_k = \frac{\sum_{b=0}^{B} \mathbb{1}\{T(b,k) \ge T(0,k)\}}{B+1}, \quad k = 0, \dots, K$$

The aggregator, $G(p_0, \ldots, p_K)$, combines this information across the columns. To calibrate it, we can treat this as yet another test statistic of the data, writing $\tilde{T}(0) = G(p_0(z), \ldots, p_K(z))$, and writing

$$T(b) = G(p_0(z_{\pi_b}), \dots, p_K(z_{\pi_b}))$$

for the result when the original sequence z is replaced by z_{π_b} . Now comes the key realization: each $p_k(z_{\pi_b})$ can be computed from the same matrix of available test statistics as above, just given by swapping the b^{th} row with the first row:

$$p_k(z_{\pi_b}) = \frac{\sum_{b'=0}^B \mathbb{1}\{T(b',k) \ge T(b,k)\}}{B+1}, \quad k = 0, \dots, K.$$

This is true because the set of permutations generated by applying $\pi_{b'}$, $b' = 0, \ldots, B$ to z_{π_b} is the same as that generated by applying $\pi_{b'}$, $b' = 0, \ldots, B$ to z (due to the group structure). Therefore each $\tilde{T}(b)$ can be computed from the same matrix of available test statistics, as can the permutation p-value corresponding to the aggregate statistic:

$$p_G = \frac{\sum_{b=0}^B \mathbb{1}\{\tilde{T}(b) \ge \tilde{T}(0)\}}{B+1}$$

This is a valid p-value in finite-sample, for the same reason that permutation p-values are valid whenever the set of permutations used forms a group (see, e.g., Hemerik and Goeman (2018)).

Figure 5 examines two aggregation methods over the same setup as that in Figure 4. As before, we use a mixture null as a proxy for the permutation null (for the sake of computational efficiency; our simulations use 100 repetitions). The aggregation rules considered are the min-rule: $G(p_0, \ldots, p_K) = \min_{k=0,\ldots,K} p_k$, and the so-called Fisher-rule: $G(p_0, \ldots, p_K) = -\sum_{k=0}^{K} \log p_k$. For each task (row), Figure 5 displays the ROC curves corresponding to these rules as dotted colored lines, and the ROC curve from the best RKS (single k) test as a solid colored line; ROC curves from the other RKS tests (other values of k) are grayed-out as a visual aid. We can see that either aggregation method frequently competes with the best RKS test, with the min-rule performing slightly better overall.

Figure 5: ROC curves across the same experimental settings as in Figure 4, along with ROC curves from the min aggregation rule ("agg-min") and Fisher aggregation rule ("agg-Fisher"). In each row, only the ROC curve from the best-performing RKS test is drawn in color, and the rest are drawn in gray.

5 Discussion

We proposed and analyzed a new multivariate nonparametric two-sample test, which is an integral probability metric (IPM) test with respect to a class of functions having unit Radon total variation of a given smoothness degree $k \ge 0$. We call this the Radon-Kolmogorov-Smirnov (RKS) test, and it can be seen as a generalization of the Kolmogorov-Smirnov (KS) test to multiple dimensions and higher orders of smoothness.

The RKS test, like any statistical test, is not expected to be most powerful in all scenarios. Roughly, we expect the RKS test to be favorable in settings in which the given distributions P and Q differ in only a few directions in the ambient space; and higher orders of smoothness k can be more sensitive to tail differences along these few directions. When P and Q differ in many or all directions, the more traditional kernel MMD test with (say) Gaussian kernel will probably perform better.

It is worth revisiting computation of the RKS distance and discussing some limitations of our work and possible future directions. In this paper, we considered *full batch* gradient descent applied to (10) (we actually used Adam, but the differences for this discussion are unimportant), i.e., in each update the entire data set $\{x_i\}_{i=1}^m \cup \{y_i\}_{i=1}^n$ is used to calculate the gradient. This results in the following complexity: T iterations of gradient descent costs O((n+m)dNT) operations, where recall N is the number of neurons. However, *mini batch* gradient descent should be also investigated, and is likely to be preferred in large problem sizes. That said, implementing mini batch gradients for (10) is nontrivial due to the nonseparable nature of the criterion across observations (due to the log and the absolute value in the first term), and pursuing this carefully is an important direction for future work. For now, we note that by just subsetting P_m, Q_n to a mini batch of size $b \leq n + m$, this results in a T-iteration complexity of O(bdNT) operations. To compare kernel MMD: this costs $O((n+m)^2d)$ operations, and thus we can see that for large enough problem sizes computation of the RKS statistic via full batch and especially mini batch gradient descent can be more efficient, provided that small or moderate values of N, T still suffice to obtain good performance.

Finally, another direction worth investigating is the use of deeper networks, which could be interpreted as learning the feature representation in which there is the largest discrepancy in projected truncated moments. This would help further distinguish the RKS test from kernel MMD, because the latter is "stuck with" the representation provided by the choice of kernel, and it cannot learn the representation based on the data. A deep RKS test could dovetail nicely with existing architectures (and even pre-trained networks) for related tasks in areas such as generative modeling, and out-of-distribution detection.

Acknowledgments

We thank Rahul Parhi and Jonathan Siegel and helpful discussions about Radon bounded variation spaces. This work was supported by NSF SCALE MoDL grant 2134133, ONR MURI grant N00014-20-1-2787, and the Miller Institute for Basic Research in Science at the University of California, Berkeley.

References

- Mélisande Albert. Concentration inequalities for randomly permuted sums. In *High Dimensional Probability* VIII: The Oaxaca Volume, 2019.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, 2017.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. Journal of Machine Learning Research, 18(19):1–53, 2017.
- Ludwig Baringhaus and Carsten Franz. On a new multivariate two-sample test. Journal of Multivariate Analysis, 88(1):190–206, 2004.
- Peter J. Bickel. A distribution free version of the Smirnov two sample test in the p-variate case. Annals of Mathematical Statistics, 40(1):1–23, 1969.
- Maurice C. Bryson. Heavy-tailed distributions: Properties and tests. Technometrics, 16(1):61-68, 1974.

- Barbara Caputo, K. Sim, Fredrik Furesjö, and Alexander J. Smola. Appearance-based object recognition using SVMs: Which kernel should I use? In Advances in Neural Information Processing Systems Workshop on Statistical Methods for Computational Experiments in Visual Processing and Computer Vision, 2002.
- Richard M. Dudley. Uniform Central Limit Theorems. Cambridge University Press, 2014. Second edition.
- Ilya Dumer. Covering spheres with spheres. Discrete & Computational Geometry, 38:665–679, 2007.
- Jerome H. Friedman and Lawrence C. Rafsky. Multivariate generalizations of the Wald-Wolfowitz and Smirnov two-sample tests. *Annals of Statistics*, 7(4):697–717, 1979.
- Anil K. Ghosh and Munmun Biswas. Distribution-free high-dimensional two-sample tests based on discriminating hyperplanes. Test, 25:525–547, 2016.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Advances in Neural Information Processing Systems, 2014.
- Alden Green, Sivaraman Balakrishnan, and Ryan J. Tibshirani. Two-sample testing with a graph-based total variation integral probability metric. arXiv preprint arXiv:2409.15628, 2024.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander J. Smola. A kernel two-sample test. Journal of Machine Learning Research, 13:723–773, 2012.
- Jesse Hemerik and Jelle Goeman. Exact testing with random permutations. Test, 27:811–825, 2018.
- Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017.
- Norbert Henze. A multivariate two-sample test based on the number of nearest neighbor type coincidences. Annals of Statistics, 16(2):772–783, 1988.
- Andrey Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *Giornale dell'Istituto Italiano degli Attuari*, 4:83–91, 1933.
- Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. MMD GAN: Towards deeper understanding of moment matching network. In Advances in Neural Information Processing Systems, 2017.
- Tianyi Lin, Chenyou Fan, Nhat Ho, Marco Cuturi, and Michael Jordan. Projection robust Wasserstein distance and Riemannian optimization. In Advances in Neural Information Processing Systems, 2020.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning*, 2015.
- Youssef Mroueh, Chun Liang Li, Tom Sercu, Anant Raj, and Yu Cheng. Sobolev GAN. In International Conference on Learning Representations, 2018.
- Jonas W. Mueller and Tommi Jaakkola. Principal differences analysis: Interpretable characterization of differences between distributions. In Advances in Neural Information Processing Systems, 2015.
- Alfred Müller. Integral probability metrics and their generating classes of functions. Advances in Applied Probability, 29(2):429–443, 1997.
- Jonathan Niles-Weed and Philippe Rigollet. Estimation of Wasserstein distances in the spiked transport model. Bernoulli, 28(4):2663–2688, 2022.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width ReLU nets: The multivariate case. In *International Conference on Learning Representations*, 2020.

- Rahul Parhi. On Ridge Splines, Neural Networks, and Variational Problems in Radon-Domain BV Spaces. PhD thesis, Department of Electrical and Computer Engineering, University of Wisconsin-Madison, 2022.
- Rahul Parhi and Robert D. Nowak. Banach space representer theorems for neural networks and ridge splines. Journal of Machine Learning Research, 22(43):1–40, 2021.
- Rahul Parhi and Robert D. Nowak. Near-minimax optimal estimation with shallow ReLU neural networks. Transactions on Information Theory, 69(2):1125–1140, 2023.
- François-Pierre Paty and Marco Cuturi. Subspace robust Wasserstein distances. In International Conference on Machine Learning, 2019.
- Veeranjaneyulu Sadhanala, Yu-Xiang Wang, Aaditya Ramdas, and Ryan J. Tibshirani. A higher-order Kolmogorov-Smirnov test. In International Conference on Artificial Intelligence and Statistics, 2019.
- Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In Annual Conference on Learning Theory, 2019.
- Mark F. Schilling. Multivariate two-sample tests based on nearest neighbors. Journal of the American Statistical Association, 81(395):799–806, 1986.
- Dino Sejdinovic, Bharath Sriperumbudur, Arthur Gretton, and Kenji Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Annals of Statistics*, 41(5):2263–2291, 2013.
- Nikolai Smirnov. Table for estimating the goodness of fit of empirical distributions. Annals of Mathematical Statistics, 19(2):279–281, 1948.
- Francesca Solmi and Patrick Onghena. Combining p-values in replicated single-case experiments with multivariate outcome. *Neuropsychological Rehabilitation*, 24(3–4):607–633, 2014.
- Gábor J. Székely and Maria L. Rizzo. A new test for multivariate normality. Journal of Multivariate Analysis, 93(1):58–80, 2005.
- Aad van der Vaart and Jon Wellner. Weak Convergence and Empirical Processes: With Applications to Statistics. Springer, 1996.
- Vladimir Vovk and Ruodu Wang. Combining p-values via averaging. Biometrika, 107(4):791–808, 2020.
- Grace Wahba. Spline Models for Observational Data. Society for Industrial and Applied Mathematics, 1990.
- Yu-Xiang Wang, Alexander J. Smola, and Ryan J. Tibshirani. The falling factorial basis and its statistical applications. In *International Conference on Machine Learning*, 2014.
- Susan Wei, Chihoon Lee, Lindsay Wichers, and J. S. Marron. Direction-projection-permutation for highdimensional hypothesis tests. Journal of Computational and Graphical Statistics, 25(2):549–569, 2016.

Proof of Proposition 1 Α

By Theorem 22 in Parhi and Nowak (2021), every $f \in RBV^k$ has a representative $f \in f$ of the form

$$f(x) = \int_{\mathbb{S}^{d-1} \times \mathbb{R}} \left[(w^{\mathsf{T}} x - b)_{+}^{k} - \sum_{|\alpha| \le k} c_{\alpha}(w, b) x^{\alpha} \right] \mathrm{d}\mu(w, b) + q_{0}(x),$$
(11)

for a finite Radon measure μ on $\mathbb{S}^{d-1} \times \mathbb{R}$ such that: (i) μ is an odd (respectively even) measure if k is even (respectively odd); (ii) the integrand is integrable in (w, b) with respect to μ for any x; (iii) $\|\mu\|_{\text{TV}} = \|f\|_{\text{RTV}^k}$; (iv) $q_0(x)$ is a polynomial of degree at most k; (v) $\sum_{|\alpha| \le k} (c_\alpha(w, b)x^\alpha + (-1)^k c_\alpha(-w, -b)x^\alpha) = (w^{\mathsf{T}}x - b)^k$. Using the conditions (i) and (v), we can show the following with a manual calculation and the change of

variables $(w, b) \rightarrow (-w, -b)$:

$$\int_{\mathbb{S}^{d-1} \times (-\infty,0)} \left[(w^{\mathsf{T}}x - b)_{+}^{k} - \sum_{|\alpha| \le k} c_{\alpha}(w, b) x^{\alpha} \right] d\mu(w, b) \\
= \int_{\mathbb{S}^{d-1} \times (0,\infty)} \left[(-w^{\mathsf{T}}x + b)_{+}^{k} - \sum_{|\alpha| \le k} c_{\alpha}(-w, -b) x^{\alpha} \right] d\mu(-w, -b) \\
= \int_{\mathbb{S}^{d-1} \times (0,\infty)} (-1)^{2k+1} \left[(w^{\mathsf{T}}x - b)_{-}^{k} - \sum_{|\alpha| \le k} (-1)^{k} c_{\alpha}(-w, -b) x^{\alpha} \right] d\mu(w, b) \\
= \int_{\mathbb{S}^{d-1} \times (0,\infty)} (-1)^{2k+1} \left[(w^{\mathsf{T}}x - b)_{-}^{k} - (w^{\mathsf{T}}x - b)^{k} + \sum_{|\alpha| \le k} c_{\alpha}(w, b) x^{\alpha} \right] d\mu(w, b) \\
= \int_{\mathbb{S}^{d-1} \times (0,\infty)} \left[(w^{\mathsf{T}}x - b)_{+}^{k} - \sum_{|\alpha| \le k} c_{\alpha}(w, b) x^{\alpha} \right] d\mu(w, b).$$
(12)

Define a measure $\tilde{\mu}$ by $\tilde{\mu}(E) = \mu(E_0) + 2\mu(E_+)$ where for an event E we denote $E_0 = (\mathbb{S}^{d-1} \times \{0\}) \cap E$ and $E_+ = (\mathbb{S}^{d-1} \times (0, \infty)) \cap E$. Then (12) rewrites (11) as

$$f(x) = \int_{\mathbb{S}^{d-1} \times [0,\infty)} \left[(w^{\mathsf{T}} x - b)_{+}^{k} - \sum_{|\alpha| \le k} c_{\alpha}(w,b) x^{\alpha} \right] d\tilde{\mu}(w,b) + q_{0}(x)$$
(13)

where $\|\tilde{\mu}\|_{\mathrm{TV}} = \|\mu\|_{\mathrm{TV}} = \|f\|_{\mathrm{RTV}^k}$. Note that $(w^{\mathsf{T}}x - b)^k_+$ is integrable in (w, b) with respect to $\tilde{\mu}$ because $\|w\|_2 = 1$ and $b \ge 0$. Thus, $\int_{\mathbb{S}^{d-1} \times [0,\infty)} (w^{\mathsf{T}}x - b)^k_+ \mathrm{d}\tilde{\mu} < \|x\|_2^k \|\tilde{\mu}\|_{\mathrm{TV}}$, hence $\sum_{|\alpha| \le k} c_{\alpha}(w, b) x^{\alpha}$ is integrable with respect to (w, b) for every x as well. Thus,

$$q_1(x) = \int_{\mathbb{S}^{d-1} \times [0,\infty)} \left(\sum_{|\alpha| \le k} c_\alpha(w,b) x^\alpha \right) \mathrm{d}\tilde{\mu}(w,b)$$

is well-defined due to the condition (ii) and (13). Therefore (13) can be rewritten as

$$f(x) = \int_{\mathbb{S}^{d-1} \times [0,\infty)} (w^{\mathsf{T}} x - b)_+^k \, \mathrm{d}\tilde{\mu}(w,b) + \tilde{q}_0(x)$$

where $\tilde{q}_0(x) = q_0(x) + q_1(x)$ is a polynomial of degree at most k. Also $\tilde{\mu}|_{\mathbb{S}^{d-1}\times\{0\}} = \mu|_{\mathbb{S}^{d-1}\times\{0\}}$ by definition of $\tilde{\mu}$, and hence, $\tilde{\mu}|_{\mathbb{S}^{d-1}\times\{0\}}$ is odd when k is even, and even when k is odd by the condition (i). Thus, we have arrived at the desired representation.

Proof of Theorem 1 Β

We rely on the representation established by Proposition 1 and treat the cases k = 0 and $k \ge 1$ separately.

Case k = 0. For this case, the boundary condition (continuity at the origin) implies that the integral may be restricted to b > 0.

Lemma 1. Consider $f \in \text{RBV}^0$. There exists a representative $f \in f$ which is continuous at $0 \in \mathbb{R}^d$ if and only if f admits a representation

$$f(x) = \int_{\mathbb{S}^{d-1} \times (0,\infty)} \mathbb{1}\{w^{\mathsf{T}} x \ge b\} \mathrm{d}\mu(w,b) + q,$$
(14)

where q is a constant and μ is a signed measure with $\|\mu\|_{TV} < \infty$.

Proof of Lemma 1. By Proposition 1, there exists $f \in f$ such that for all x,

$$f(x) = \int_{\mathbb{S}^{d-1} \times [0,\infty)} \mathbb{1}\{w^{\mathsf{T}} x \ge b\} \mathrm{d}\mu(w,b) + q = f^{=0}(x) + f^{>0}(x),$$

where

$$f^{>0}(x) = \int_{\mathbb{S}^{d-1} \times (0,\infty)} \mathbb{1}\{w^{\mathsf{T}} x \ge b\} \mathrm{d}\mu(w,b) + q,$$
$$f^{=0}(x) = \int_{\mathbb{S}^{d-1} \times \{0\}} \mathbb{1}\{w^{\mathsf{T}} x \ge 0\} \mathrm{d}\mu(w,b).$$

If b > 0, then $\lim_{x\to 0} \mathbb{1}\{w^{\mathsf{T}}x \ge b\} = 0$. Thus, by dominated convergence,

$$\lim_{x \to 0} f^{>0}(x) = q(x) = f^{>0}(0).$$

Thus, $f^{>0}$ is continuous at 0. We conclude that f is almost-everywhere equal to a function which is continuous at 0 if and only if $f^{=0}$ is almost-everywhere equal to a function which is continuous at 0.

For $x \neq 0$, $w^{\mathsf{T}} x \geq 0$ if and only if $w^{\mathsf{T}} x/||x|| \geq 0$. Thus, $f^{=0}(x) = f^{=0}(x/||x||)$. Thus, $f^{=0}$ is almost everywhere equal to a function which is continuous at 0 if and only if there exists some $c \in \mathbb{R}$ such that for almost every $x \in \mathbb{S}^{d-1}$ (almost every with respect to Lebesgue surface measure), we have $f^{=0}(x) = c$. But this implies that $f^{=0}(x) = c$ for almost every $x \in \mathbb{R}^d$. Thus, we get that for almost every $x \in \mathbb{R}^d$,

$$f(x) = f^{>0}(x) + c = \int_{\mathbb{S}^{d-1} \times (0,\infty)} \mathbb{1}\{w^{\mathsf{T}} x \ge b\} \mathrm{d}\mu(w,b) + q + c.$$

The right-hand side is the desired representative.

Theorem 1 in the case k = 0 is now a consequence of the following lemma, which additionally gives an explicit form for the radially cone continuous representative of f.

Lemma 2. Assume $f \in RBV^0$ contains a representative which is continuous at 0. Then it contains a unique representative which is radially cone continuous. This representative can be written in the form (14) where μ is a signed measure with $\|\mu\|_{TV} < \infty$. Conversely, any function of the form (14) where μ is a signed measure with $\|\mu\|_{TV} < \infty$ is the unique radial cone continuous representative of some $f \in RBV^0$.

Proof of Lemma 2. Consider $f \in \text{RBV}^0$ which contains a representative which is continuous at the origin. By Lemma 1, there exists $f \in f$ which satisfies (14) for all $x \in \mathbb{R}^d$, for some constant q and signed measure μ with $\|\mu\|_{\text{TV}} < \infty$.

Fix $x \neq 0$. Consider any sequence $\epsilon_j \in \mathbb{R}$ such that $\epsilon_j \downarrow 0$, and $v_j \in \mathbb{S}^{d-1}$ such that $v_j \to x/||x||$. Consider any $(w,b) \in \mathbb{S}^{d-1} \times (0,\infty)$. If $w^{\mathsf{T}}x < b$, then $w^{\mathsf{T}}(x+\epsilon_j v_j) < b$ eventually because $x + \epsilon_j v_j \to x$. Alternatively, $w^{\mathsf{T}}x \geq b$ gives $w^{\mathsf{T}}x/||x|| > 0$, whence $w^{\mathsf{T}}v_j > 0$ eventually. In particular, $w^{\mathsf{T}}(x+\epsilon_j v_j) > w^{\mathsf{T}}x > 0$ eventually. Thus, in both cases with have $\mathbb{1}\{w^{\mathsf{T}}(x+\epsilon_j v_j) \geq b\} \to \mathbb{1}\{w^{\mathsf{T}}x \geq b\}$. By dominated convergence, we get that

$$f(x+\epsilon_j v_j) = \int_{\mathbb{S}^{d-1} \times (0,\infty)} \mathbb{1}\{w^\mathsf{T}(x+\epsilon_j v_j) \ge b\} \mathrm{d}\mu(w,b) + q$$
$$\to \int_{\mathbb{S}^{d-1} \times (0,\infty)} \mathbb{1}\{w^\mathsf{T}x \ge b\} \mathrm{d}\mu(w,b) + q = f(x).$$

Thus, f is radially cone continuous and can be written in the form of (14).

By Theorem 22 in Parhi and Nowak (2021) and reversing the construction of $\tilde{\mu}$ from μ in the proof of Proposition 1, any function of the form (14) where μ is a signed measure with $\|\mu\|_{TV} < \infty$ is a representative of an equivalence class $f \in RBV^0$. The argument above also shows that it is radially cone continuous. Thus, we have shown the converse statement as well.

Case k = 1. We begin with a lemma.

Lemma 3. For $k \ge 1$ and every $\mathbf{f} \in \text{RBV}^k$, there exists a unique representative $f \in \mathbf{f}$ which is continuous. This representative is in fact (k-1)-times continuously differentiable and can be written in the form in Proposition 1. Moreover, for any multi-index α with $|\alpha| \le k-1$, we have for all $x \in \mathbb{R}^d$ that

$$D^{\alpha}f(x) = \frac{k!}{(k-|\alpha|)!} \int_{\mathbb{S}^{d-1}\times[0,\infty)} w^{\alpha} (w^{\mathsf{T}}x - b)^{k-|\alpha|}_{+} \mathrm{d}\mu(w,b) + D^{\alpha}q(x).$$
(15)

where $w^{\alpha} = \prod_{i=1}^{d} w_i^{\alpha_i}$.

Proof of Lemma 3. For any $f \in \text{RBV}^k$, let $f \in f$ be a representative of the form in Proposition 1. We will show that this representative f is (k-1)-times continuously differentiable by inducting on the order $m = |\alpha|$ of the multi-index α .

The base case is m = 0. In this case, the representation of the 0th derivative is equivalent to the representation in the proposition. The continuity of f follows by Lemma 12 applied to this representation.

Now consider $m \ge 1$, and assume we have established the result for m-1. Consider α with $|\alpha| = m$, and without loss of generality, assume $\alpha_1 \ge 1$ and let $\tilde{\alpha} = (\alpha_1 - 1, \alpha_2, \dots, \alpha_d)$. By the inductive hypothesis,

$$D^{\alpha}f(x) = \frac{k!}{(k-|\alpha|+1)!} \partial_{x_1} \int_{\mathbb{S}^{d-1} \times [0,\infty)} w^{\tilde{\alpha}} (w^{\mathsf{T}}x - b)_+^{k-|\alpha|+1} \mathrm{d}\mu(w,b) + D^{\alpha}q(x).$$

Note that for $x \in \mathbb{R}^d$, we have $\partial_{x_1} (w^{\tilde{\alpha}}(w^{\mathsf{T}}x-b)_+^{k-|\alpha|+1}) = (k-|\alpha|+1)w^{\alpha}(w^{\mathsf{T}}x-b)^{k-|\alpha|}$. This function is integrable in (w, b) with respect to μ for every $x \in \mathbb{R}^d$, since $||w||_2 = 1$ and $b \ge 0$. Moreover, the partial derivative is uniformly bounded on $\{x' : ||x'|| \le 2||x||\}$ by $2^{k-|\alpha|+1}(k-|\alpha|+1)(\prod_{i=1}^d |w_i|^{\alpha_i})(2||x||)^{k-|\alpha|}$, which is integrable with respect to μ . These conditions allow us to apply the Leibniz rule to conclude that we can exchange differentiation and integration, which yields (15). The continuity of $D^{\alpha}f$ then follows from Lemma 12, and hence, the induction is complete.

Thus we have shown that f is (k-1)-times continuously differentiable. All that is left to show is that f is the unique element of f which is (k-1)-times continuously differentiable. This holds because any two continuous functions which are equal almost everywhere are in fact equal everywhere.

Theorem 1 in the case $k \ge 1$ is now a consequence of the following lemma, which additionally gives an explicit form for the radially cone continuous representative of f.

Lemma 4. For $k \ge 1$, let $f \in \text{RBV}^k$ be such that the unique continuous representative $f \in f$ is k-times classically differentiable at $0 \in \mathbb{R}^d$ with $D^{\alpha}f(0) = 0$ for all $|\alpha| \le k$. Then f has representation of the form in Proposition 1 with $q(x) \equiv 0$.

Proof of Lemma 4. Because $w^{\alpha}(w^{\mathsf{T}}0-b)^{k-|\alpha|}_{+}=0$ for $|\alpha| \leq k-1$ and any $(w,b) \in \mathbb{S}^{d-1} \times [0,\infty)$, (15) implies that $D^{\alpha}f(0) = D^{\alpha}q(0)$ for all $|\alpha| \leq k-1$.

Next, we evaluate $D^{\alpha}f(0)$ for $|\alpha| = k$, which, by assumption, exists. Without loss of generality, assume $\alpha_1 \ge 1$, and let $\tilde{\alpha} = (\alpha_1 - 1, \alpha_2, \dots, \alpha_d)$. Note that for b > 0, $\partial_{x_i} w^{\tilde{\alpha}} (w^{\mathsf{T}} x - b)_+ |_{x=0} = 0$. By (15), we can write

$$D^{\alpha}f(0) = k!\partial_{x_1} \int_{\mathbb{S}^{d-1} \times [0,\infty)} w^{\tilde{\alpha}}(w^{\mathsf{T}}x - b)_+ d\mu(w,b)\Big|_{x=0} + D^{\alpha}q(0).$$

We compute the partial derivative in the first term manually:

$$\partial_{x_1} \int_{\mathbb{S}^{d-1} \times [0,\infty)} w^{\tilde{\alpha}} (w^{\mathsf{T}} x - b)_+ \, \mathrm{d}\mu(w,b) \Big|_{x=0} = \lim_{h \to 0} \frac{1}{h} \int_{\mathbb{S}^{d-1} \times [0,\infty)} w^{\tilde{\alpha}} (w_1 h - b)_+ \, \mathrm{d}\mu(w,b).$$

For b > 0 and $w \in \mathbb{S}^{d-1}$, we have $(w_1h - b)_+ \leq |h| \cdot \mathbb{1}\{w_1h > b\}$ because $|w_1| \leq 1$. Also $|w^{\alpha}| \leq 1$ because $|w_j| \leq 1$ for all j. Thus,

$$\begin{aligned} \left|\frac{1}{h} \int_{\mathbb{S}^{d-1} \times (0,\infty)} w^{\tilde{\alpha}}(w_1 h - b)_+ \, \mathrm{d}\mu(w, b) \right| &\leq \frac{1}{|h|} \int_{\mathbb{S}^{d-1} \times (0,\infty)} w^{\tilde{\alpha}} |h| \mathbb{1}\{w_j h - b\} \, \mathrm{d}\mu(w, b) \\ &\leq |\mu| \Big((w, b) \in \mathbb{S}^{d-1} \times (0,\infty) : w_j h > b\Big). \end{aligned}$$

Because $|\mu|$ is a finite positive measure and $\bigcap_{h>0} \{(w,b) \in \mathbb{S}^{d-1} \times (0,\infty) : w_j h > b\} = \emptyset$, we conclude $|\mu| ((w,b) \in \mathbb{S}^{d-1} \times (0,\infty) : w_j h > b) \to 0$ as $h \to 0$. Thus, we have that

$$D^{\alpha}f(0) = k! \lim_{h \to 0} \left[\frac{1}{h} \int_{\mathbb{S}^{d-1} \times \{0\}} w^{\tilde{\alpha}}(w_1h - b)_+ \,\mathrm{d}\mu(w, b)\right] + D^{\alpha}q(0), \tag{16}$$

and in particular, the limit in the first term exists. For b = 0, $(w_1h - b)_+ = w_1h \cdot \mathbb{1}\{w_1 > 0\}$ when h > 0, and $(w_1h - b)_+ = w_1h \cdot \mathbb{1}\{w_1 < 0\}$ when h < 0. Thus, for h > 0,

$$\frac{1}{h} \int_{\mathbb{S}^{d-1} \times \{0\}} w^{\tilde{\alpha}}(w_1 h - b)_+ \, \mathrm{d}\mu(w, b) = \int_{\substack{\mathbb{S}^{d-1} \times \{0\}\\w_1 > 0}} w^{\alpha} \, \mathrm{d}\mu(w, b),$$

and for h < 0,

$$\frac{1}{h} \int_{\mathbb{S}^{d-1} \times \{0\}} w^{\tilde{\alpha}} (w_1 h - b)_+ \, \mathrm{d}\mu(w, b) = \int_{\mathbb{S}^{d-1} \times \{0\}} w^{\alpha} \, \mathrm{d}\mu(w, b)$$
$$= (-1)^{|\alpha|} (-1)^{k+1} \int_{\substack{\mathrm{S}^{d-1} \times \{0\}\\w_1 > 0}} w^{\alpha} \, \mathrm{d}\mu(w, b),$$

where the second equality uses the change of variables $w \mapsto -w$ and that $\mu|_{\mathbb{S}^{d-1} \times \{0\}}$ is odd when k is even and even when k is odd (see Proposition 1). By differentiability, the quantities in the two previous displays must be equal. Because $(-1)^{|\alpha|}(-1)^{k+1} = -1$, this implies that the quantities in the two previous displays must be 0. That is, the limit in (16) must be 0, and we have $0 = D^{\alpha}f(0) = D^{\alpha}q(0)$.

We have thus shown that $D^{\alpha}q(0) = 0$ for all $|\alpha| \leq k$, which, because q(x) is a polynomial of degree at most k, implies $q(x) \equiv 0$.

C Proof of Theorem 2

Theorem 2 is an immediate consequence of Lemmas 2 and 4.

D Proof of Theorem 3

We again treat the cases k = 0 and $k \ge 1$ separately.

Case k = 0. Let

$$\mathcal{G}_0^+ = \Big\{ \mathbb{1}\{w^{\mathsf{T}} \ge b\} : \ (w, b) \in \mathbb{S}^{d-1} \times (0, \infty) \Big\}.$$

Because $\mathcal{G}_0^+ \subseteq \mathcal{G}_0$, for any signed measure μ on $\mathbb{S}^{d-1} \times (0, \infty)$ with $\|\mu\|_{\mathrm{TV}} = 1$ and $q \in \mathbb{R}$,

$$\sup_{f \in \mathcal{G}_0} P(f) - Q(f) \ge \sup_{f \in \mathcal{G}_0^+} P(f) - Q(f) \ge \int_{\mathbb{S}^{d-1} \times (0,\infty)} \int_{\mathbb{R}^d} \mathbb{1}\{w^\mathsf{T} x \ge b\} \, \mathrm{d}(P - Q)(x) \, \mathrm{d}\mu(w, b)$$
$$= \int_{\mathbb{R}^d} \left(\int_{\mathbb{S}^{d-1} \times (0,\infty)} \mathbb{1}\{w^\mathsf{T} x \ge b\} \, \mathrm{d}\mu(w, b) + q \right) \, \mathrm{d}(P - Q)(x).$$

Taking the supremum on the right-hand side over measure μ , we get that

$$\sup_{f \in \mathcal{G}_0} P(f) - Q(f) \ge \sup_{f \in \mathcal{F}_0} P(f) - Q(f).$$

Because $\lim_{b\to\infty} P(w^{\mathsf{T}}x \ge b) - Q(w^{\mathsf{T}}x \ge b)$, \mathbb{S}^{d-1} is compact, and $P(w^{\mathsf{T}}x \ge b) - Q(w^{\mathsf{T}}x \ge b)$ is continuous, the supremum on the left-hand side is achieved. Then, if the supremum is achieved at some (w, b) with b > 0, since $x \mapsto \mathbb{1}\{w^{\mathsf{T}}x \ge b\} \in \mathcal{F}_0$, we also have

$$\sup_{f \in \mathcal{G}_0} P(f) - Q(f) \le \sup_{f \in \mathcal{F}_0} P(f) - Q(f).$$

On the other hand, consider that the supremum is achieved at some (w, 0). Note that

$$\lim_{b \uparrow 0} Q(w^{\mathsf{T}} x \le b) - P(w^{\mathsf{T}} x \le b) = P(w^{\mathsf{T}} x \ge 0) - Q(w^{\mathsf{T}} x \ge 0).$$

But $Q(w^{\mathsf{T}}x \leq b) - P(w^{\mathsf{T}}x \leq b) = P(f)_b - Q(f)_b$ for $f_b(x) = \mathbb{1}\{-w^{\mathsf{T}}x \geq -b\}$. Note that $f_b \in \mathcal{G}_0^+$, and therefore, $f_b \in \mathcal{F}_0$. Thus, we have in this case also that $\sup_{f \in \mathcal{G}_0} P(f) - Q(f) \leq \sup_{f \in \mathcal{F}_0} P(f) - Q(f)$. Having shown both directions of the inequality, we conclude $\sup_{f \in \mathcal{G}_0} P(f) - Q(f) = \sup_{f \in \mathcal{F}_0} P(f) - Q(f)$.

Having shown both directions of the inequality, we conclude $\sup_{f \in \mathcal{G}_0} P(f) - Q(f) = \sup_{f \in \mathcal{F}_0} P(f) - Q(f)$ The same result holds with the roles of P and Q reversed, which establishes the result with the supremum over |P(f) - P(Q)|.

Case $k \geq 1$. Because P and Q have finite k^{th} moments, by dominated convergence

$$(w,b) \mapsto \mathbb{E}_P[(w^\mathsf{T} x - b)^k_+] - \mathbb{E}_Q[(w^\mathsf{T} x - b)^k_+]$$

is continuous on $\mathbb{S}^{d-1} \times [0, \infty)$. Thus,

$$\sup_{f \in \mathcal{G}_k} P(f) - Q(f) = \sup_{(w,b) \in \mathbb{S}^{d-1} \times (0,\infty)} \mathbb{E}_P[(w^{\mathsf{T}}x - b)_+^k] - \mathbb{E}_Q[(w^{\mathsf{T}}x - b)_+^k] \le \sup_{f \in \mathcal{F}_k} P(f) - Q(f),$$

where the last inequality holds because $x \mapsto (w^{\mathsf{T}}x - b)_{+}^{k} \in \mathcal{F}_{k}$ for all $w \in \mathbb{S}^{d-1}$ and b > 0 (Parhi and Nowak, 2021). On the other hand, for any $f \in \mathcal{F}_{k}$,

$$P(f) - Q(f) = \int_{\mathbb{R}^d} \int_{\mathbb{S}^{d-1} \times [0,\infty)} (w^\mathsf{T} x - b)_+^k \, \mathrm{d}\mu(w,b) \, \mathrm{d}(P - Q)(x)$$

=
$$\int_{\mathbb{S}^{d-1} \times [0,\infty)} \int_{\mathbb{R}^d} (w^\mathsf{T} x - b)_+^k \, \mathrm{d}(P - Q)(x) \, \mathrm{d}\mu(w,b) \le \sup_{f' \in \mathcal{G}_k} |P(f') - Q(f')|,$$

where the second equality uses Fubini's theorem and the final inequality uses $\|\mu\|_{\text{TV}} = \|f\|_{\text{RTV}^k} \leq 1$. Because the previous display holds for all $f \in \mathcal{F}_k$, we have $\sup_{f \in \mathcal{F}_k} P(f) - Q(f) = \sup_{f \in \mathcal{G}_k} |P(f) - Q(f)|$. Since $f \in \mathcal{G}_k$ implies $-f \in \mathcal{G}_k$, we have $\sup_{f \in \mathcal{G}_k} |P(f) - Q(f)| = \sup_{f \in \mathcal{G}_k} P(f) - Q(f)$.

Thus, we have shown both that $\sup_{f \in \mathcal{G}_k} |P(f) - Q(f)| \leq \sup_{f \in \mathcal{F}_k} |P(f) - Q(f)|$ and the reverse inequality, so that these are in fact equal.

E Proof of Theorem 4

If P = Q, then for any $f \in \mathcal{F}_k$, P(f) - Q(f) = 0, whence $\rho(P, Q; \mathcal{F}_k) = 0$. Alternatively, consider $P \neq Q$. Then, by the uniqueness of the characteristic function, there exists $w \in \mathbb{R}^d$ such that the distribution of $w^{\mathsf{T}}X$ and $w^{\mathsf{T}}Y$, where $X \sim P$ and $Y \sim Q$, are different. This implies that there exists some $t \in \mathbb{R}$ such that $P(w^{\mathsf{T}}X \geq t) \neq Q(w^{\mathsf{T}}Y \geq t)$.

In what follows, we treat the cases k = 0, k = 1, and $k \ge 2$ separately.

Case k = 0. If $P(w^{\mathsf{T}}X \ge t) > Q(w^{\mathsf{T}}Y \ge t)$ for some $t \ge 0$, then taking $f(x) = \mathbb{1}\{w^{\mathsf{T}}x \ge t\}$ gives

$$\rho(P,Q;\mathcal{F}_0) = \rho(P,Q;\mathcal{G}_0) \ge P(f) - Q(f) > 0,$$

where we have used Theorem 3. If $P(w^{\mathsf{T}}X \ge t) > Q(w^{\mathsf{T}}Y \ge t)$ for some $t \le 0$, then because $t \mapsto P(w^{\mathsf{T}}X \ge t)$ is left-continuous, for s < t sufficiently close to t, we will have $P(w^{\mathsf{T}}X > s) > Q(w^{\mathsf{T}}Y > s)$. Then, taking $f(x) = -\mathbb{1}\{(-w)^{\mathsf{T}}x \ge -s\}$ gives

$$\rho(P,Q;\mathcal{F}_0) = \rho(P,Q;\mathcal{G}_0) \ge P(f) - Q(f) > 0,$$

where we have used Theorem 3.

Similarly, if $P(w^{\mathsf{T}}X \ge t) < Q(w^{\mathsf{T}}X \ge t)$, then we can take $f(x) = -\mathbb{1}\{w^{\mathsf{T}}x \ge t\}$ in the case $t \ge 0$ and $f(x) = \mathbb{1}\{(-w)^{\mathsf{T}}x \ge -s\}$ for some s < t sufficiently close to t to conclude $\rho(P, Q; \mathcal{F}_0) > 0$.

Case k = 1. Because $\mathbb{E}_P[\mathbb{1}\{w^{\mathsf{T}}X \ge s\}]$ and $\mathbb{E}_Q[\mathbb{1}\{w^{\mathsf{T}}Y \ge s\}]$ are left-continuous in s, if $w^{\mathsf{T}}X$ and $w^{\mathsf{T}}y$ have a different distribution, there is some open interval $(\alpha, \beta) \subseteq \mathbb{R}$ on which $\mathbb{E}_P[\mathbb{1}\{w^{\mathsf{T}}X \ge s\}] > \mathbb{E}_Q[\mathbb{1}\{w^{\mathsf{T}}Y \ge s\}]$ or $\mathbb{E}_P[\mathbb{1}\{w^{\mathsf{T}}X \ge s\}] < \mathbb{E}_Q[\mathbb{1}\{w^{\mathsf{T}}Y \ge s\}]$, uniformly over all $s \in (\alpha, \beta)$. Because either $\beta > 0$ or $\alpha < 0$, this interval contains an open interval entirely contained in either $\mathbb{R}_{>0}$ or $\mathbb{R}_{<0}$. In the former case, we have that for b > 0, $\mathbb{E}_P[(w^{\mathsf{T}}X - b)_+] \neq \mathbb{E}_Q(w^{\mathsf{T}}Y - b)_+]$ because

$$\mathbb{E}_P\left[(w^\mathsf{T} X - b)_+\right] = \int_b^\infty \mathbb{E}_P\left[\mathbbm{1}\{w^\mathsf{T} X \ge s\}\right] \mathrm{d}s, \quad \mathbb{E}_Q\left[(w^\mathsf{T} Y - b)_+\right] = \int_b^\infty \mathbb{E}_Q\left[\mathbbm{1}\{w^\mathsf{T} Y \ge s\}\right] \mathrm{d}s.$$

We then get $\rho(P,Q;\mathcal{F}_1) = \rho(P,Q;\mathcal{G}_1) > 0$ by taking $f(x) = (w^{\mathsf{T}}x - b)_+$. In the latter case, we have that for b > 0, $\mathbb{E}_P[((-w)^{\mathsf{T}}X - b)_+] \neq \mathbb{E}_Q[((-w)^{\mathsf{T}}Y - b)_+]$ because

$$\mathbb{E}_{P}\left[((-w)^{\mathsf{T}}X - b)_{+}\right] = \int_{-\infty}^{-b} \mathbb{E}_{P}\left[(1 - \mathbb{1}\{w^{\mathsf{T}}X \ge s\})\right] \mathrm{d}s$$

and likewise for Q. We then get $\rho(P,Q;\mathcal{F}_1) = \rho(P,Q;\mathcal{G}_1) > 0$ by taking $f(x) = ((-w)^{\mathsf{T}}x - b)_+$.

Case $k \ge 2$. Note that the derivatives of the mappings $b \mapsto \mathbb{E}_P[(w^{\mathsf{T}}X - b)_+^k]$ and $b \mapsto \mathbb{E}_Q[(w^{\mathsf{T}}Y - b)_+^k]$ are $b \mapsto k\mathbb{E}_P[(w^{\mathsf{T}}X - b)_+^{k-1}]$ and $b \mapsto k\mathbb{E}_Q[(w^{\mathsf{T}}Y - b)_+^{k-1}]$, respectively. If there is some (w, b) with $b \ge 0$ such that $\mathbb{E}_P[(w^{\mathsf{T}}X - b)_+^{k-1}] \neq \mathbb{E}_P[(w^{\mathsf{T}}X - b)_+^{k-1}]$, by continuity, $\mathbb{E}_P[(w^{\mathsf{T}}X - s)_+^{k-1}] \neq \mathbb{E}_P[(w^{\mathsf{T}}X - s)_+^{k-1}]$ for all s in an open interval around b, whence, by integration, $\mathbb{E}_P[(w^{\mathsf{T}}X - s)_+^k] \neq \mathbb{E}_P[(w^{\mathsf{T}}X - s)_+^k]$ for some s in this interval. We can then conclude the result by induction, using k = 1 as a base case.

F Proof of Theorem 5

For a function class \mathcal{F} , a probability measure \mathbb{P} , and $\epsilon > 0$, we say that $\{[l_i, u_i] : i \in [N]\}$ is an (ϵ, \mathbb{P}) -bracket of \mathcal{F} with cardinality N if for all $i \in [N]$ we have $l_i(x) \leq u_i(x)$ for all x and

$$||u_i(X) - l_i(X)||^2_{L^2(\mathbb{P})} = \mathbb{E}_{X \sim \mathbb{P}}[(u_i(X) - l_i(X))^2] \le \epsilon^2,$$

and for all $f \in \mathcal{F}$ there exists $i \in [N]$ such that $l_i(x) \leq f(x) \leq u_i(x)$ for all x. The (ϵ, \mathbb{P}) -bracketing number, which we denote by $N_{[]}(\epsilon; \mathcal{F}, \|\cdot\|_{L^2(\mathbb{P})})$, is the minimum cardinality of an (ϵ, \mathbb{P}) -bracket of \mathcal{F} . Using standard techniques from empirical process theory, we will be able to prove Theorem 5 by bounding the bracketing number $N_{[]}(\epsilon; \mathcal{G}_k, \|\cdot\|_{L^2(\mathbb{P})})$.

F.1 Bracketing number for $\mathcal{G}_k, k \geq 1$

Theorem 8. Let $k \ge 1$ and let \mathbb{P} be a probability measure with finite moments of order $2k + \Delta$, for any fixed $\Delta > 0$, i.e., $\mathbb{E}_{X \sim \mathbb{P}} \|X\|_2^{2k+\Delta} = M < \infty$. Then, for the class \mathcal{G}_k , there is a constant $C_{[]} > 0$ depending only on M, d, Δ , and k such that

$$\log N_{[]}(\epsilon; \mathcal{G}_k, \|\cdot\|_2) \le C_{[]} \log \left(1 + \frac{1}{\epsilon}\right)$$

Proof of Theorem 8. We denote by $\mathsf{B}^d_r(x)$ the ball of radius r centered at $x \in \mathbb{R}^d$:

$$\mathsf{B}^{d}_{r}(x) = \{ z \in \mathbb{R}^{d} \mid ||z - x||_{2}^{2} \le r^{2} \}.$$

We say a set $S \subseteq \mathbb{R}^d$ is a *r*-covering of a set $A \subseteq \mathbb{R}^d$ if A is contained in the union of the balls of radius r centered at points $x \in I$:

$$A \subseteq \bigcup_{x \in I} \mathsf{B}^d_r(x).$$

For a compact set $I \subseteq \mathbb{R}^d$ and real number $b \in \mathbb{R}$, define functions

$$\ell_{I,b}(x) = \inf \left\{ (w^{\mathsf{T}} x - b)_{+}^{k} : w \in I \right\}, \text{ and } u_{I,b}(x) = \sup \left\{ (w^{\mathsf{T}} x - b)_{+}^{k} : w \in I \right\},$$

where we adopt the convention that $\ell_{I,\infty}(x) = u_{I,\infty}(x) \equiv 0$.

Lemma 5. Consider S is an r-covering of \mathbb{S}^{d-1} . Let \mathcal{I} be the collection of sets $\{I = \mathbb{S}^{d-1} \cap \mathsf{B}_r^d(x) : x \in S\}$, and let $b_j = j\delta$ for $j = -1, \ldots, N$, and $b_{N+1} = \infty$ for some $\delta > 0$ and integer N > 0. Then, the following is a bracket of \mathcal{G}_k

$$\{[l_{I,b_{j+1}}, u_{I,b_j}] : I \in \mathcal{I}, j = -1, \dots, N\}.$$
(17)

Proof of Lemma 5. Consider any $(w, b) \in \mathbb{S}^{d-1} \times [0, \infty)$ and its corresponding function $x \mapsto (w^{\mathsf{T}}x - b)_+^k \in \mathcal{G}_k$. Because S is an r-covering, there exists $I \in \mathcal{I}$ such that $w \in I$. Also clearly there exists $j \in \{-1, 0, 1, \dots, N\}$ such that $b_j < b \leq b_{j+1}$. Then $\ell_{I,b_{j+1}} \leq (w^{\mathsf{T}}x - b)_+^k$ from the following observation:

$$\ell_{I,b_{j+1}}(x) = \inf\left\{ (\tilde{w}^{\mathsf{T}}x - b_{j+1})_{+}^{k} : \tilde{w} \in I \right\} \le (w^{\mathsf{T}}x - b_{j+1})_{+}^{k} \le (w^{\mathsf{T}}x - b)_{+}^{k}.$$

Similarly, we can show $(w^{\mathsf{T}}x - b)_+^k \leq (w^{\mathsf{T}}x - b_j)_+^k \leq u_{I,b_j}(x).$

Lemma 6. In the setting of Lemma 5, if we set $r = \delta = \frac{\epsilon/3}{k\sqrt{1+\max(1,M)}}$ and $N = \lceil (\frac{M}{\epsilon})^{1/\Delta} \frac{1}{\delta} \rceil$, then we have $\|u_{I,b_j} - l_{I,b_{j+1}}\|_{L^2(P)} \leq \epsilon$ for all $j = -1, \ldots, N$ and $I \in \mathcal{I}$.

Proof of Lemma 6. We bound

$$|u_{I,b_{j}}(x) - l_{I,b_{j+1}}(x)| \leq \max_{w,v \in I} |(w^{\mathsf{T}}x - b_{j})_{+}^{k} - (v^{\mathsf{T}}x - b_{j+1})_{+}^{k}|$$

$$\leq \max_{w,v \in I} |(w^{\mathsf{T}}x - b_{j})_{+}^{k} - (v^{\mathsf{T}}x - b_{j})_{+}^{k}| + \max_{v \in I} |(v^{\mathsf{T}}x - b_{j})_{+}^{k} - (v^{\mathsf{T}}x - b_{j+1})_{+}^{k}|.$$

Using a telescoping sum, we may bound for $j \leq N$, using Lemma 11:

$$\begin{aligned} |(w^{\mathsf{T}}x - b_j)_+^k - (v^{\mathsf{T}}x - b_j)_+^k| &= |(w^{\mathsf{T}}x - b_j)_+ - (v^{\mathsf{T}}x - b_j)_+| \times \sum_{i=0}^{k-1} (w^{\mathsf{T}}x - b_j)_+^i (v^{\mathsf{T}}x - b_j)_+^{k-1-i} \\ &\leq ||w - v||_2 ||x||_2 \times k \cdot \max\{(w^{\mathsf{T}}x)_+, (v^{\mathsf{T}}x)_+\}^{k-1} \\ &\leq k ||w - v||_2 ||x||_2^k \\ &\leq 2kr ||x||_2^k. \end{aligned}$$

Likewise, we may bound for $j \leq N - 1$

$$\begin{aligned} |(v^{\mathsf{T}}x - b_j)_+^k - (v^{\mathsf{T}}x - b_{j+1})_+^k| &= |(v^{\mathsf{T}}x - b_j)_+ - (v^{\mathsf{T}}x - b_{j+1})_+| \times \sum_{i=0}^{k-1} (v^{\mathsf{T}}x - b_j)_+^i (v^{\mathsf{T}}x - b_{j+1})_+^{k-1-i} \\ &\leq |b_j - b_{j+1}| \times k \cdot (v^{\mathsf{T}}x - b_j)_+^{k-1} \\ &\leq k\delta ||x||_2^{k-1}. \end{aligned}$$

Therefore, for $j \leq N - 1$,

$$|u_{I,b_j}(x) - l_{I,b_{j+1}}(x)| \le k\delta ||x||_2^{k-1} + 2kr ||x||_2^k \le \begin{cases} k(\delta + 2r) & : ||x||_2 \le 1\\ k(\delta + 2r) ||x||_2^k \le k(\delta + 2r) ||x||_2^{k+\Delta/2} & : ||x||_2 > 1 \end{cases},$$

which can be simplified to

$$|u_{I,b_j}(x) - l_{I,b_{j+1}}(x)| \le k(\delta + 2r) \max\{1, \|x\|_2^{k+\Delta/2}\}$$

This gives for $j \leq N - 1$,

$$\|u_{I,b_j} - l_{I,b_{j+1}}\|_{L^2(\mathbb{P})}^2 \le k^2 (\delta + 2r)^2 (1 + \mathbb{E}_{X \sim P} \|X\|_2^{2k+\Delta}) \le k^2 (\delta + 2r)^2 (1+M) \le \epsilon^2,$$

where the final inequality is obtained by plugging in the values of r and δ specified in the statement of the lemma. Meanwhile, for j = N, we have $b_{j+1} = \infty$, so that $l_{I,b_{j+1}} = 0$. Also, $b_j = N\delta \ge (M/\epsilon)^{1/\Delta}$. Thus,

$$||u_{I,b_N} - l_{I,b_{N+1}}||^2_{L^2(\mathbb{P})} = \int_{\mathbb{R}^d} u_{I,b_N}(x)^2 \mathrm{d}\mathbb{P}(x)$$

$$\leq \int_{\mathbb{R}^d} (\|x\|_2 - b_N)_+^{2k} d\mathbb{P}(x) = \int_{\mathbb{R}^d} (\|x\|_2 - N\delta)_+^{2k} d\mathbb{P}(x)$$

$$\leq \left(\int_{\mathbb{R}^d} (\|x\|_2 - N\delta)^{2k+\Delta} d\mathbb{P}(x)\right)^{\frac{2k}{2k+\Delta}} \left(\int_{\mathbb{R}^d} 1(\|x\|_2 \ge N\delta) d\mathbb{P}(x)\right)^{\frac{\Delta}{2k+\Delta}}$$

$$= M^{\frac{2k}{2k+\Delta}} \times \mathbb{P}\Big(\|X\|_2 \ge C\delta\Big)^{\frac{\Delta}{2k+\Delta}}$$

$$\leq M^{\frac{2k}{2k+\Delta}} \times \Big(\frac{M}{(N\delta)^{2k+\Delta}}\Big)^{\frac{\Delta}{2k+\Delta}} = \frac{M}{(N\delta)^{\Delta}},$$

where the second inequality uses Hölder's inequality, and the last inequality used Markov inequality. With the value of δ and N specified by the lemma, we conclude $\|u_{I,b_N} - l_{I,b_{N+1}}\|_{L^2(\mathbb{P})}^2 \leq \epsilon^2$.

We can now complete the proof of Theorem 8. By Dumer (2007), there exists an *r*-covering of \mathbb{S}^{d-1} with cardinality at most $(1+2/r)^d$. Thus, using the choices of r, δ , and N in Lemma 6 gives an (ϵ, \mathbb{P}) -bracket of \mathcal{G}_k that has cardinality at most

$$\left(1 + \frac{2k\sqrt{1 + \max(1, M)}}{\epsilon/3}\right)^d \left(\left(\frac{M}{\epsilon}\right)^{1/\Delta} \frac{k\sqrt{1 + \max(1, M)}}{\epsilon/3} + 1\right) \le C(1 + 1/\epsilon)^{d+1+1/\Delta}.$$

with some constant C. The proof is completed.

F.2 Bracketing number for \mathcal{G}_0

Theorem 9. Let \mathbb{P} be a probability measure with finite Δ^{th} moments, i.e., $\mathbb{E}_{X \sim \mathbb{P}} \|X\|_2^{\Delta} = M < \infty$, for any fixed $\Delta > 0$. Suppose additionally that for the random vector $X \sim P$ and any $w \in \mathbb{S}^{d-1}$, the random variable $w^{\mathsf{T}}X \in \mathbb{R}$ has the density upper bounded by M_{∞} , i.e.,

$$\sup_{w \in \mathbb{S}^{d-1}} \left\| p_{w^{\mathsf{T}_{x}}} \right\|_{\infty} < M_{\infty},\tag{18}$$

where $p_{w^{\mathsf{T}_x}}$ denotes the density of w^{T_x} for $x \sim P$, and $||f||_{\infty} = \sup_{t \in \mathbb{R}} |f(t)|$ denotes the supremum norm of f. Then, for the class \mathcal{G}_0 , there is a constant $C_{[1]} > 0$ depending only on M, M_{∞} , d, and Δ such that

$$\log N_{[]}(\epsilon; \mathcal{G}_0, \|\cdot\|_2) \le C_{[]} \log \left(1 + \frac{1}{\epsilon}\right).$$

Proof of Theorem 9. We adopt the same notation used in the proof of Theorem 8. The set (17) is a bracket of \mathcal{G}_0 by Lemma 5. Now we show that it is an ϵ -bracket for appropriate choices of r, δ , and N.

Select $I \in \mathcal{I}, j = -1, ..., N-1$. There exists a $w_0 \in I$ such that $||w - w_0||_2 \leq r$ for all $w \in I$, since S is an r-cover of \mathbb{S}^{d-1} . Consequently, for any $x \in \mathbb{R}^d$

$$\max_{w \in I} 1(w^{\mathsf{T}} x \ge b_j) \le 1(w_0^{\mathsf{T}} x \ge b_j - r \|x\|_2),$$

$$\min_{v \in I} 1(v^{\mathsf{T}} x \ge b_{j+1}) \ge 1(w_0^{\mathsf{T}} x \ge r \|x\|_2 + b_{j+1})$$

Note that the random variable $w_0^{\mathsf{T}} X \in \mathbb{R}$ has a density that is upper bounded by M_{∞} as (18). Therefore, for any R > 0,

$$\begin{aligned} \|u_{I,b_j} - l_{I,b_j}\|_{L^2(\mathbb{P})}^2 &= \mathbb{P}\left(\max_{w,v} \left| 1(w^\mathsf{T} x \ge b_j) - \min_{v \in I} 1(v^\mathsf{T} x \ge b_{j+1}) \right| \right) \\ &\leq \mathbb{P}\left(w_0^\mathsf{T} X \in [b_j - r \|X\|_2, b_{j+1} + r \|X\|_2]\right) \\ &\leq \mathbb{P}\left(w_0^\mathsf{T} X \in [b_j - r R, b_{j+1} + r R]\right) + \mathbb{P}\left(\|X\|_2 \ge R\right) \\ &\leq M_\infty \left(\delta + 2r R\right) + \mathbb{P}\left(\|X\|_2 \ge R\right) \\ &\leq M_\infty \left(\delta + 2r R\right) + \frac{M}{R^\Delta}, \end{aligned}$$

with the last line following from Markov's inequality. Taking $R = (3M/\epsilon^2)^{1/\Delta}$, $r = \frac{\epsilon^2}{6RM_{\infty}} = \frac{\epsilon^{2(1+1/\Delta)}}{6(3M)^{1/\Delta}M_{\infty}}$ and $\delta = \frac{\epsilon^2}{3M_{\infty}}$ makes the above upper bound equal to ϵ^2 . On the other hand for $I \in \mathcal{I}, j = N$, we have $l_{w,b_N+1} = 0$, and

$$\|u_{I,b_N} - l_{I,b_{N+1}}\|_P^2 = \mathbb{P}\left(\max_{w \in I} w^{\mathsf{T}} X \ge b_N\right)$$
$$\leq \mathbb{P}\left(\|X\|_2 \ge b_N\right)$$
$$\leq \frac{M}{(N\delta)^{\Delta}}.$$

Taking $N = \frac{(M/\epsilon^2)^{1/\Delta}}{\delta}$ makes this upper bound equal to ϵ^2 . So for these choices of r, δ and N the bracket is an $\epsilon\text{-bracket}.$

Finally, arguing as in the proof of Theorem 7, we conclude that there exists an ϵ -bracket of \mathcal{G}_0 with cardinality at most

$$N \times \left(1 + \frac{2}{r}\right)^{d} \le C(1 + \frac{1}{\epsilon})^{2d(1+1/\Delta) + 2(1+1/\Delta)}.$$

This completes the proof.

F.3 Proof of Theorem 5

Theorem 5 is a consequence of the (ϵ, \mathbb{P}) -bracketing number bound of Theorem 8 (or Theorem 9, for k=0) and Theorems 7.6 and 9.1 in Dudley (2014), which we copy below for convenience.

Theorem 10. Let (Ω, Σ, P) be a probability space and let $\mathcal{F} \subseteq \mathcal{L}^2(\Omega, \Sigma, P)$ be such that

$$J_{[]}(\mathcal{F}, \|\cdot\|_2) = \int_0^1 \sqrt{\log N_{[]}(\epsilon; \mathcal{F}, \|\cdot\|_{L^2(P)})} \,\mathrm{d}\epsilon < \infty.$$

Then, under the P-null hypothesis, as $m, n \to \infty$,

$$\sqrt{\frac{1}{m} + \frac{1}{n}} \left(P_m - Q_n \right) \xrightarrow{d} \mathbb{G}_P, \text{ and hence } \sqrt{\frac{1}{m} + \frac{1}{n}} \sup_{f \in \mathcal{F}} \left| P_m f - Q_n f \right| \xrightarrow{d} \sup_{f \in \mathcal{F}} |\mathbb{G}_P(f)|.$$

Theorem 10 implies Theorem 5 because $\int_0^1 \sqrt{\log(1+1/\epsilon)} \, d\epsilon < \infty$.

G Proof of Theorem 6

J

Our proof follows closely the proof of the analogous result for the univariate higher-order KS test in Sadhanala et al. (2019). We begin by recalling the following lemma, which appears as Theorem 2.14.2 and 2.14.5 in van der Vaart and Wellner (1996), the statement of which we transcribe from Sadhanala et al. (2019).

Lemma 7. Let \mathcal{F} be a class of functions with an envelope function F; i.e., $f \leq F$ for all $f \in \mathcal{F}$. Define

$$W = \sqrt{n} \sup_{f \in \mathcal{F}} |P_n f - P f|,$$

and let $J = \int_0^1 \sqrt{\log N_{[]}(\epsilon; \mathcal{F}, \|\cdot\|_{L^2(P)})} d\epsilon$. If for $p \ge 2$, $\|F\|_{L^p(P)} < \infty$, then for a universal constant $c_1 > 0$, $\mathbb{E}[|W|^p]^{1/p} \le c_1(||F||_{L^2(P)}J + n^{-1/2+1/p}||F||_{L^p(P)}).$

If, for
$$0 , the exponential Orlicz norm of F , i.e., $||F||_{\Psi_p} = \inf\{t > 0 : \mathbb{E}[\exp(|F|^p/t^p)] \le 2\}$, is finite, then$$

$$\|W\|_{\Psi_p} \le c \left(\|F\|_{L^2(P)} J + n^{-1/2} (1 + \log n)^{1/2} \|F\|_{\Psi_p}\right)$$

Using this lemma, we have the following theorem on the concentration of the the RKS test statistic on its population value, from which the asymptotic result in Theorem 6 can be deduced.

Theorem 11. We have the following tail bounds.

(i) Consider k = 0. Assume P and Q have Δ^{th} moment bounded by M for $M, \Delta > 0$, are absolutely continuous with respect to Lebesuge measure, and satisfies the condition in (18). Then, for some constant c_0 depending only on M, Δ, M_{∞} , and d, the following holds with probability at least $1 - \alpha$:

$$|T_{d,k} - \rho(P,Q;\mathcal{F}_k)| \le c_0(\log(1/\alpha)) \Big(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}}\Big).$$

(ii) Consider $k \ge 1$. Assume P and Q have finite p^{th} moments bounded by M for $p = 2k + \Delta$ for $M, \Delta > 0$. Then, for some constant c_0 depending only on M, Δ, k , and d, the following holds with probability at least $1 - \alpha$:

$$|T_{d,k} - \rho(P,Q;\mathcal{F}_k)| \le \frac{c_0}{\alpha^{1/p}} \Big(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}}\Big).$$

Assume, on the other hand, that X and Y have Orlicz norm of order $0 bounded by M when <math>X \sim P$ and $Y \sim Q$. Then, for some constant c_0 depending only on M, p, k, and d, the following holds with probability at least $1 - \alpha$:

$$|T_{d,k} - \rho(P,Q;\mathcal{F}_k)| \le c_0 (\log(1/\alpha))^{1/p} \Big(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}}\Big).$$

Proof of Theorem 11. First consider $k \ge 1$. Note that a bound on any Orlicz norm of order 0 implies $a bound on the moments of order <math>2k + \Delta$ of P, Q for $\Delta = 1$. Thus, assuming either the moment bound or the Orlicz norm bound, Theorem 8 implies $J = \int_0^1 \sqrt{\log N_{[]}(\epsilon; \mathcal{G}_k, \|\cdot\|_{L^2(P)})} d\epsilon$ has a finite upper bound depending only on the moments of order $2k + \Delta$ of P, Q, and on d, k, Δ . Take $F(x) = \|x\|_2^k$. If we only assume a bound on the moments of order $2k + \Delta$ of P, Q, then take $p = 2k + \Delta \ge 2$. In this case, $n^{-1/2+1/p} \le 1$, whence using Lemma 7 and Markov's inequality, we see that $\sup_{f \in \mathcal{F}} |P_n f - Pf| \le c_0/(\alpha^{1/p}\sqrt{n})$ and $\sup_{f \in \mathcal{F}} |Q_n f - Qf| \le c_0/(\alpha^{1/p}\sqrt{m})$ with probability at least $1 - \alpha$. The asserted bound then holds by the triangle inequality. On the other hand, if we assume a bound on the Orlicz norm of order p, Lemma 7 implies a constant upper bound on the Orlicz norm of order p of $\sqrt{n} \sup_{f \in \mathcal{F}} |P_n f - Pf|$ and $\sqrt{m} \sup_{f \in \mathcal{F}} |Q_n f - Qf|$ because $(1 + \log n)^{1/2}/n^{1/2} \le 1$. This gives us the desired tail bound by taking $t = c_0(\log(1/\alpha))^{1/p}$ in the supremum defining the Orlicz norm and Markov's inequality.

Now consider k = 0. We take F as a constant 1 function, i.e., F = 1, and using Theorem 9 to bound the integral J in this case. Because F = 1 has a finite Orlicz norm of order p = 1, the asserted tail bound follows from Lemma 7.

We show how this result implies Theorem 6. By Theorem 5, if P = Q, then $T_{d,k} = O_p(\sqrt{(n+m)/(nm)})$. Because $1/t_{m,n} = o(\sqrt{n+m}) = o(\sqrt{(nm)/(n+m)})$, we see that the test rejects with asymptotic probability 0. On the other hand, if $P \neq Q$, then by Theorem 4, we have $\rho(P,Q;\mathcal{F}_k) > 0$. By Theorem 11, we have $T_{d,k} \xrightarrow{P} \rho(P,Q;\mathcal{F}_k)$. Because $t_{m,n} \to 0$, this implies that the test rejects with asymptotic probability 1.

H Proof of Theorem 7

We will develop the following nonasymptotic statement about type II error control using permutations, from which the asymptotic statement in Theorem 7 can be deduced.

Theorem 12. Assume P and Q have the bounded domains included in the L-radius Euclidean ball whose center is the origin in \mathbb{R}^d . If k = 0, additionally assume P and Q each has density satisfying (18). Then, there exists a constant C which only depends on (L, d, k) when $k \ge 1$ or (L, d, k, M_{∞}) when k = 0, such that if $m, n \ge 2$ and $\beta \in (0, 1]$ satisfies

$$\rho(P,Q;\mathcal{F}_k) \ge C\sqrt{\frac{1}{m} + \frac{1}{n}} \times \left(\log(m+n) + \log\left(\frac{1}{\alpha}\right) + \frac{1}{\beta^{1/(2k+1)}}\right),\tag{19}$$

then the permutation test with the p-value defined in (3) with the full permutations controls the type II error, i.e., under the alternative hypothesis $H_1: P \neq Q$, the p-value p in (3) satisfies $\mathbb{P}_{H_1}(p \leq \alpha) \geq 1 - \beta$.

We show how this can be used to prove Theorem 7. Choose any $\beta \in (0, 1]$. Due to the condition on $\alpha_{m,n}$, we know that for large enough m and n, (19) is satisfied. Therefore,

$$\lim_{m,n\to\infty} \mathbb{P}_{H_1}(p \le \alpha_{m,n}) \ge 1 - \beta.$$

This implies $\lim_{m,n\to\infty} \mathbb{P}_{H_1}(p \leq \alpha_{m,n}) = 1$ since β is arbitrary in (0, 1]. Combined with the finite-sample type I error property of permutation tests, and the condition $\alpha_{m,n} \to 0$, this proves Theorem 7. The rest of this section is devoted to proving Theorem 12.

H.1 Additional notation

Define S_{m+n} as a set of all possible permutations for $(1, \ldots, m+n)$. Recall the notations in (3): π denotes a permutation such that $\pi \in S_{m+n}$; $z = (z_{\ell})_{\ell=1}^{m+n} = (x_1, \ldots, x_m, y_1, \ldots, y_n)$; applying permutation π to zoutputs z^{π} . To emphasize the dependency on the data, we may use the notation $T_{d,k}(z)$ and $T_{d,k}(z^{\pi})$. Define a vector

$$a = (a_{\ell})_{\ell=1}^{m+n} = (\underbrace{1/m, \dots, 1/m}_{m}, \underbrace{-1/n, \dots, -1/n}_{n}).$$

Note that $\sum_{\ell=1}^{m+n} a_{\ell} = 0$. With the permutation π , we have $a^{\pi} = (a_{\pi(\ell)})_{\ell=1}^{m+n}$. Define $\theta(w,b) = (\theta(w,b)_{\ell})_{\ell=1}^{m+n}$ and $\theta(w,b)_{\ell} = (w^{\mathsf{T}} z_{\ell} - b)_{+}^{k}$. For brevity, we write $\theta(w,b)$ instead of $\theta(w,b;z)$ although $\theta(w,b)$ is a function of z. We can rewrite $T_{d,k}(z)$ with $\theta(w,b)$:

$$T_{d,k}(z) = \max_{w,b \in \mathbb{S}^{d-1} \times [0,\infty)} \left| \frac{1}{m} \sum_{i=1}^{m} (w^{\mathsf{T}} x_i - b)_+^k - \frac{1}{n} \sum_{i=1}^{n} (w^{\mathsf{T}} y_i - b)_+^k \right| = \max_{(w,b) \in \mathbb{S}^{d-1} \times [0,\infty)} \left| \langle \theta(w,b), a \rangle \right|.$$

Moreover, we can rewrite the permuted test statistic $T_{d,k}(z^{\pi})$ in a similar format:

$$T_{d,k}^{\pi}(z) = T_{d,k}(z^{\pi}) = \max_{(w,b)\in\mathbb{S}^{d-1}\times[0,\infty)} \left| \langle \theta(w,b), a^{\pi} \rangle \right|.$$
 (20)

Note that $T_{d,k}^{\pi}(z)$ is a function of π when z is fixed. Define the $(1 - \alpha)$ -quantile of all permuted test statistics $(T_{d,k}^{\pi}(z))_{\pi \in S_{m+n}}$, given z, as the following:

$$t_{1-\alpha}(z) = \inf \left\{ t : \mathbb{P}_{\pi \sim \mathrm{Unif}(S_{m+n})} \left(T_{d,k}^{\pi}(z) \le t \, \middle| \, z \right) \ge 1 - \alpha \right\}.$$

$$(21)$$

Define a random variable $X(w,b) = \langle \theta(w,b), a^{\pi} \rangle$, which is a function of $z \sim P^m \times Q^n$ and $\pi \sim \text{Unif}(S_{m+n})$. However, we use the notation X(w,b) instead of $X(w,b;z^{\pi})$, again for brevity. Then, $T_{d,k}^{\pi}(z)$ can be also written as

$$T_{d,k}^{\pi}(z) = \max_{(w,b) \in \mathbb{S}^{d-1} \times [0,\infty)} |X(w,b)|.$$
 (22)

H.2 Proof of Theorem 12

We need the following result, which will be proved later.

Lemma 8. Assume the same as Theorem 7. For $k \ge 0$, define a constant

$$t_{m,n,\alpha} = C_{L,d,k} \sqrt{\frac{1}{m} + \frac{1}{n}} \left(\log\left(\frac{1}{\alpha}\right) + \log\left(m+n\right) \right)$$
(23)

where $C_{L,d,k(,M_{\infty})}$ is a constant which only depends on (L,d,k) for $k \ge 1$ and (L,d,k,M_{∞}) for k = 0. Recall the definition of $t_{1-\alpha}(z)$ in (21). Then, for any given samples z, we have $t_{1-\alpha}(z) \le t_{m,n,\alpha}$.

Assuming Lemma 8 holds, we can prove Theorem 12. Since P and Q have bounded domain, they have finite moments of order 2k + 1. Also, for k = 0, we have additionally assumed the boundedness of the density $p_w^{\intercal}x$. Thus, by Theorem 11, there exists a constant $c_{L,d,k(,M_{\infty})}$ which only depends on (L,d,k) for $k \ge 1$ or (L,d,k,M_{∞}) for k = 0, such that

$$\left|T_{d,k}-\rho(P,Q;\mathcal{F}_k)\right| \leq \frac{c_{L,d,k(,M_{\infty})}}{\beta^{1/(2k+1)}} \left(\frac{1}{\sqrt{m}}+\frac{1}{\sqrt{n}}\right)$$
 with probability at least $1-\beta$.

Note that the randomness from the above equation comes from $x_i \stackrel{\text{i.i.d.}}{\sim} P$ and $y_j \stackrel{\text{i.i.d.}}{\sim} Q$. By choosing the constant C in Theorem 7 as $C = \max\{C_{L,d,k}, c_{L,d,k}(M_{\infty})\}$, we can easily observe that

$$t_{m,n,\alpha} \le \rho(P,Q;\mathcal{F}_k) - \frac{c_{L,d,k(,M_\infty)}}{\beta^{1/(2k+1)}} \Big(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}}\Big).$$

Therefore, under the alternative hypothesis H_1 , we have

$$\mathbb{P}_{z}\left(T_{d,k}(z) < t_{m,n,\alpha}\right) \leq \mathbb{P}_{z}\left(T_{d,k}(z) \leq \rho(P,Q;\mathcal{F}_{k}) - \frac{c_{L,d,k(,M_{\infty})}}{\beta^{1/(2k+1)}} \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}}\right)\right) \leq \beta.$$

$$(24)$$

Finally combining (24) with Lemma 8, we have

$$\mathbb{P}_{z}\Big(t_{1-\alpha}(z) > T_{d,k}(z)\Big) \leq \mathbb{P}\Big(t_{1-\alpha}(z) > t_{m,n,\alpha}\Big) + \mathbb{P}\Big(T_{d,k}(z) < t_{m,n,\alpha}\Big) \leq 0 + \beta = \beta,$$

which concludes the proof of Theorem 12. Now to prove Lemma 8 is the only remaining part.

H.3 Proof sketch of Lemma 8

The proof Lemma 8 follows techniques developed in Green et al. (2024). Since the proof consists of multiple steps, we introduce the outline of the proof here before we dive in.

1. Fix z. Find a "good" finite subset $\mathcal{C}_{\text{good}}$ of $\mathbb{S}^{d-1} \times [0, \infty)$, of cardinality $N_{\text{good}} = |\mathcal{C}_{\text{good}}| < \infty$. For any $\mathcal{C} \subseteq \mathbb{S}^{d-1} \times [0, \infty)$, define

$$T_{d,k,\mathcal{C}}^{\pi}(z) = \max_{(w,b)\in\mathcal{C}} \left| \langle \theta(w,b), a^{\pi} \rangle \right| = \max_{(w,b)\in\mathcal{C}} \left| X(w,b) \right|.$$
(25)

We control N_{good} and the approximation error $|T_{d,k,\mathcal{C}_{\text{good}}}^{\pi}(z) - T_{d,k}^{\pi}(z)|$.

- 2. Find a high-probability upper bound for X(w, b) for each $(w, b) \in \mathcal{C}_{good}$. Here z is fixed, and hence the only randomness is rooted from $\pi \sim \text{Unif}(S_{m+n})$.
- 3. Find a high-probability upper bound for $T^{\pi}_{d,k,\mathcal{C}_{\text{good}}}(z)$. Again, π is the only source of randomness.
- 4. Lastly, we get a high-probability upper bound for $T_{d,k}^{\pi}(z)$, with the results from the first and third steps.

H.4 Proof of Lemma 8

The proof follows the steps explained in Appendix H.3. Throughout the whole proof, we consider z is fixed.

H.4.1 C_{good} and the approximation error $|T_{d,k,\mathcal{C}_{\text{good}}}^{\pi}(z) - T_{d,k}^{\pi}(z)|$

Here we will prove the following Lemma.

Lemma 9. We have the following finite subsets of C_{\circ} and C_{ϵ} of $\mathbb{S}^{d-1} \times [0, \infty)$, whose sizes are $N_{\circ} = |C_{\circ}|$ and $N_{\epsilon} = |C_{\epsilon}|$.

• Consider k = 0. There exists C_{\circ} such that

$$\log N_{\circ} \le (d+1)\log(1+m+n) \quad and \quad \left| T_{d,k}^{\pi}(z) - T_{d,k,\mathcal{C}_{\circ}}^{\pi}(z) \right| = 0.$$
(26)

• Consider $k \geq 1$. For any $\epsilon > 0$, there exists C_{ϵ} such that

$$\log N_{\epsilon} \le C_{L,d,k}' \log \left(1 + \frac{m+n}{\epsilon} \right) \quad and \quad \left| T_{d,k}^{\pi}(z) - T_{d,k,\mathcal{C}_{\epsilon}}^{\pi}(z) \right| \le \left(\frac{1}{m} + \frac{1}{n} \right) \epsilon, \tag{27}$$

where $C'_{L,d,k}$ is a constant only depends on (L,d,k).

We will use C_{good} to denote C_{\circ} when k = 0 or C_{ϵ} when $k \geq 1$.

Case k = 0. Define $C_{\circ} = \{(w_i, b_i) : i = 1, \dots, N_{\circ}\}$ to satisfy the following condition:

$$\{\theta(w_i, b_i) : (w_i, b_i) \in \mathcal{C}_{\circ}\} = \{\theta(w, b) : (w, b) \in \mathbb{S}^{d-1} \times [0, \infty)\}.$$

Note that such \mathcal{C}_{\circ} exists since $\theta(w, b)_{\ell} \in \{0, 1\}$ for any $(w, b) \in \mathbb{S}^{d-1} \times [0, \infty)$ and $1 \leq \ell \leq m+n$.

An upper bound for N_{\circ} can be obtained from a simple observation about the Vapnik–Chervonenkis (VC) dimension and the shattering number³ of $\mathcal{H} = \{1(w^{\mathsf{T}}x + b \geq 0) : (w, b) \in \mathbb{S}^{d-1} \times [0, \infty)\}$. It is well known that the VC dimension of $\{1(w^{\mathsf{T}}x + b \geq 0) : (w, b) \in \mathbb{S}^{d-1} \times \mathbb{R}\}$ is at most d + 1. Thus the VC dimension of \mathcal{H} , saying $V_{\mathcal{H}}$, is at most d + 1 as well. Therfore the shattering number

$$S_{\mathcal{H}}(m+n) = \max_{v_1, \dots, v_{m+n} \in \mathbb{R}^d} \left| \left\{ \left(h(v_1), \dots, h(v_{m+n}) \right)^{\mathsf{T}} \in \mathbb{R}^{m+n} : h \in H \right\} \right| \\ = \max_{v_1, \dots, v_{m+n} \in \mathbb{R}^d} \left| \left\{ \left((w^{\mathsf{T}}v_1 - b)_+^0, \dots, (w^{\mathsf{T}}v_{m+n} - b)^0 \right)^{\mathsf{T}} \in \mathbb{R}^{m+n} : (w, b) \in \mathbb{S}^{d-1} \times [0, \infty) \right\} \right|$$

satisfies $S_{\mathcal{H}}(m+n) \leq (1+m+n)^{V_{\mathcal{H}}} \leq (1+m+n)^{d+1}$. This implies that we can find \mathcal{C}_{\circ} such that

$$\log N_{\circ} = \log S_{\mathcal{H}}(m+n) \le (d+1)\log(1+m+n).$$

Also, by the definition of C_{\circ} it is clear that the approximation error with C_{\circ} is zero:

$$\left|T_{d,k}^{\pi}(z) - T_{d,k,\mathcal{C}_{\circ}}^{\pi}(z)\right| = 0.$$

Case $k \ge 1$. Recall that in Appendix F, we introduced the (ϵ, \mathbb{P}) -bracketing number for a class \mathcal{F} and a measure \mathbb{P} . Define a discrete measure $R_{m+n} = \frac{1}{m+n} \sum_{\ell=1}^{m+n} \delta(z_\ell)$. Since $||z_i||_2 \le L$ for any $i \in \{1, \ldots, m+n\}$, the measure R_{m+n} has finite moments of order 2k+1 for any $k \ge 1$, upper bounded by L^{2k+1} .

Thus we can apply Lemma 5 and Theorem 8 for (ϵ, R_{m+n}) -bracketing number. Following Lemma 5, we can construct (ϵ, P_m) -bracket of \mathcal{G}_k of the form $[\ell_{I,b_{j+1}}, u_{I,b_j}]$ as (17). Denote such bracket K_{ϵ} . Theorem 8 implies the following upper bound on $|K_{\epsilon}|$, where $C'_{L,d,k}$ is a constant only depending on (L, d, k):

$$\log |K_{\epsilon}| \le C'_{L,d,k} \log \left(1 + \frac{1}{\epsilon}\right).$$

Using $K_{\epsilon/(m+n)}$, we can construct C_{ϵ} , a subset of $\mathbb{S}^{d-1} \times [0, \infty)$ which is going to be used instead of $\mathbb{S}^{d-1} \times [0, \infty)$ with certain approximation error (which will be controlled later). For each $\kappa = [\ell_{I,b_{j+1}}, u_{I,b_j}] \in K_{\epsilon/(m+n)}$, choose any $w_{\kappa} \in I$ and $b_{\kappa} \in (b_j, b_{j+1}]$. Then define $C_{\epsilon} = \{(w_{\kappa}, b_{\kappa}) : \kappa \in K_{\epsilon/m+n}\}$. Also, $N_{\epsilon} = |C_{\epsilon}|$ satisfies the following upper bound:

$$\log N_{\epsilon} = \log |K_{\epsilon/(m+n)}| \le C'_{L,d,k} \log \left(1 + \frac{m+n}{\epsilon}\right).$$

Now we will control the approximation error of using C_{ϵ} instead of $\mathbb{S}^{d-1} \times [0, \infty)$. Recall that $T_{d,k,C_{\epsilon}}^{\pi}$ is defined as (25). Choose any $(w,b) \in \mathbb{S}^{d-1} \times [0,\infty)$. By the definition of the bracket, there exists $\kappa = [\ell_{I,b_{j+1}}, u_{I,b_j}] \in K_{\epsilon/(m+n)}$ such that $\ell_{I,b_{j+1}}(x) \leq (w^{\mathsf{T}}x - b)_+^k \leq u_{I,b_j}(x)$ for any $x \in \mathbb{R}^d$. Meanwhile, by the construction of (w_{κ}, b_{κ}) , it is clear that $\ell_{I,b_{j+1}}(x) \leq (w^{\mathsf{T}}x - b_{\kappa})_+^k \leq u_{I,b_j}(x)$. Therefore, by the definition of the bracket,

$$\left\| (w^{\mathsf{T}} \cdot -b)_{+}^{k} - (w_{\kappa}^{\mathsf{T}} \cdot -b_{\kappa})_{+}^{k} \right\|_{L^{2}(R_{m+n})}^{2} \leq \left\| \ell_{I,b_{j+1}} - u_{I,b_{j}} \right\|_{L^{2}(R_{m+n})}^{2} \leq \frac{\epsilon^{2}}{(m+n)^{2}}$$

and hence, due to Cauchy-Schwarz inequality,

$$\left(\left\|\theta(w,b) - \theta(w_{\kappa},b_{\kappa})\right\|_{1}\right)^{2} = \left(\sum_{\ell=1}^{m+n} \left|\theta(w,b)_{\ell} - \theta(w_{\kappa},b_{\kappa})_{\ell}\right|\right)^{2} \le (m+n)\sum_{\ell=1}^{m+n} \left|\theta(w,b)_{\ell} - \theta(w_{\kappa},b_{\kappa})_{\ell}\right|^{2}$$

³It is also called the growth function or the shatter coefficient.

$$= (m+n) \sum_{\ell=1}^{m+n} \left| (w^{\mathsf{T}} z_{\ell} - b)_{+}^{k} - (w_{\kappa}^{\mathsf{T}} z_{\ell} - b_{\kappa})_{+}^{k} \right|^{2}$$
$$= (m+n)^{2} \left\| (w^{\mathsf{T}} \cdot -b)_{+}^{k} - (w_{\kappa}^{\mathsf{T}} \cdot -b_{\kappa})_{+}^{k} \right\|_{L^{2}(R_{m+n})}^{2} \leq \epsilon^{2}$$

Finally, we can check the upper bound for the approximation error of using (w_{κ}, b_{κ}) instead of (w, b):

$$\begin{aligned} \left| |X(w,b)| - |X(w_{\kappa},b_{\kappa})| \right| &\leq |X(w,b) - X(w_{\kappa},b_{\kappa})| = \left| \sum_{\ell=1}^{m+n} \left(\theta(w,b)_{\ell} - \theta(w_{\kappa},b_{\kappa})_{\ell} \right) a_{\ell}^{\pi} \right| \\ &\leq \sum_{\ell=1}^{m+n} \left| \left(\theta(w,b)_{\ell} - \theta(w_{\kappa},b_{\kappa})_{\ell} \right) a_{\ell}^{\pi} \right| \\ &\leq \max\left(\frac{1}{m},\frac{1}{n}\right) \left\| \theta(w,b) - \theta(w_{\kappa},b_{\kappa}) \right\|_{1} \leq \left(\frac{1}{m} + \frac{1}{n}\right) \epsilon. \end{aligned}$$

This directly implies

$$\left|T_{d,k}^{\pi}(z) - T_{d,k,\mathcal{C}_{\epsilon}}^{\pi}(z)\right| \leq \left(\frac{1}{m} + \frac{1}{n}\right)\epsilon.$$

H.4.2 High-probability upper bound for $|X_{w,b}|$ for each $(w,b) \in \mathcal{C}_{good}$

We need a following lemma from Corollary 2.2 in Albert (2019).

Lemma 10. Let $\{b_{ij}\}_{i,j=1}^{N}$ be an $N \times N$ array of numbers, and π be a permutation distributed $\text{Unif}(S_N)$. Define a random variable $U_i = b_{i,\pi(i)}$ for $i \in \{1,\ldots,N\}$. Note that π is the only randomness in U_i . Assume that (i) $\mathbb{E}[U_i] = 0$, (ii) $\text{Var}[\sum_{i=1}^{N} U_i] \leq \sigma^2$, and (iii) $|U_i| \leq b$ with probability 1. Then for any $t \geq 0$,

$$\mathbb{P}\bigg(\Big|\sum_{i=1}^{n} U_i\Big| \ge t\bigg) \le 16e^{1/16} \exp\bigg(-\frac{t^2}{256(\sigma^2 + bt)}\bigg).$$

This implies that

$$\left|\sum_{i=1}^{n} U_{i}\right| \leq \sigma \sqrt{256 \cdot \log\left(\frac{16e^{1/16}}{\delta}\right)} + 256 \cdot b \log\left(\frac{16e^{1/16}}{\delta}\right) \quad with \ probability \ at \ least \ 1-\delta.$$

In this subsection, we only consider X(w, b) with $w, b \in \mathcal{C}_{good}$. Recall the definition $X(w, b) = \sum_{\ell=1}^{m+n} \theta(w, b)_{\ell} a_{\ell}^{\pi}$, and note that $\mathbb{E}_{\pi \sim \text{Unif}(S_{m+n})}[a_{\ell}^{\pi} \mid z] = m \cdot (1/m) + n \cdot (-1/n) = 0$. Write $X(w, b)_{\ell} = \theta(w, b)_{\ell} a_{\ell}^{\pi}$. To use the lemma we will check the following three statements, where z is fixed:

- 1. $\mathbb{E}_{\pi \sim \text{Unif}(S_{m+n})}[X(w,b)_{\ell} | z] = 0$ for any ℓ .
- 2. $\operatorname{Var}(X(w,b) | z) \leq \sigma^2$ where $\sigma = L^k \sqrt{2(\frac{1}{m} + \frac{1}{n})}$.
- 3. $|X(w,b)_{\ell}| \leq b$ where $b = L^k(\frac{1}{m} + \frac{1}{n})$ for any ℓ .

First statement. Note that z is fixed. Easily observe that

$$\mathbb{E}_{\pi \sim \mathrm{Unif}(S_{m+n})}[X(w,b)_{\ell} \mid z] = \theta(w,b)_{\ell} \cdot \mathbb{E}_{\pi \sim \mathrm{Unif}(S_{m+n})}[a_{\ell}^{\pi} \mid z] = 0.$$

Second statement. We need to first calculate $\operatorname{Var}(a_{\ell}^{\pi} | z)$ and $\operatorname{Cov}(a_{\ell}^{\pi}, a_{h}^{\pi} | z)$ for $\ell \neq h$.

$$\operatorname{Var}(a_{\ell}^{\pi} \mid z) = \mathbb{E}_{\pi \sim \operatorname{Unif}(S_{m+n})}[(a_{\ell}^{\pi})^{2}] = \frac{m}{m+n} \cdot \frac{1}{m^{2}} + \frac{n}{m+n} \cdot \frac{1}{n^{2}} = \frac{1}{mn}$$
$$\operatorname{Cov}(a_{\ell}^{\pi}, a_{h}^{\pi} \mid z) = \mathbb{E}_{\pi \sim \operatorname{Unif}(S_{m+n})}[a_{\ell}^{\pi}a_{h}^{\pi} \mid z] = \frac{\binom{m}{2}}{\binom{m+n}{2}} \cdot \frac{1}{m^{2}} + \frac{mn}{\binom{m+n}{2}} \cdot \frac{-1}{mn} + \frac{\binom{n}{2}}{\binom{m+n}{2}} \cdot \frac{1}{n^{2}} \qquad (28)$$
$$= -\frac{1/(m+n-1)}{mn}$$

We are now ready to observe $\operatorname{Var}(X(w, b) | z)$.

$$\begin{aligned} \operatorname{Var}(X(w,b) \mid z) &= \operatorname{Var}\left(\sum_{\ell=1}^{m+n} \theta(w,b)_{\ell} a_{\ell}^{\pi} \mid z\right) = \sum_{\ell=1}^{m+n} \operatorname{Var}\left(\theta(w,b)_{\ell} a_{\ell}^{\pi}\right) + \sum_{\ell \neq h} \operatorname{Cov}\left(\theta(w,b)_{\ell} a_{\ell}^{\pi}, \theta(w,b)_{h} a_{h}^{\pi}\right) \\ &= \sum_{\ell=1}^{m+n} (\theta(w,b)_{\ell})^{2} \frac{1}{mn} - \sum_{\ell \neq h} \theta(w,b)_{\ell} \theta(w,b)_{h} \frac{1/(m+n-1)}{mn} \quad (by \ (28)) \\ &= \frac{1}{mn(m+n-1)} \left((m+n) \left(\sum_{\ell=1}^{m+n} (\theta(w,b)^{\ell})^{2}\right) - \left(\sum_{\ell=1}^{m+n} \theta(w,b)_{\ell}\right)^{2}\right) \\ &\leq \frac{m+n}{mn(m+n-1)} \left(\sum_{\ell=1}^{m+n} (\theta(w,b)_{\ell})^{2}\right) \\ &= \frac{m+n}{mn(m+n-1)} (m+n) L^{2k} \quad (\text{since } \theta(w,b)_{\ell} = (w^{\mathsf{T}} z_{\ell} - b)_{+}^{k} \leq ||z_{\ell}||_{2}^{k}, \text{ as } b \geq 0) \\ &\leq 2L^{2k} \left(\frac{1}{m} + \frac{1}{n}\right) \quad (\text{since } m, n \geq 1). \end{aligned}$$

Third statement. Easily check that $|X(w,b)_{\ell}| = |\theta(w,b)_{\ell}a_{\ell}^{\pi}| \le (w^{\mathsf{T}}z_{\ell}-b)_{+}^{k} \cdot \max(\frac{1}{m},\frac{1}{n}) \le L^{k}(\frac{1}{m}+\frac{1}{n}).$

Conclusion. With the above three statements and Lemma 10, we obtain the following high-probability upper bound for |X(w,b)| for each $(w,b) \in \mathcal{C}_{good}$, where z is fixed: with probability at least $1 - \delta$,

$$\left|X(w,b)\right| = \left|\sum_{\ell=1}^{m+n} X(w,b)_{\ell}\right| \le L^k \sqrt{\frac{1}{m} + \frac{1}{n}} \sqrt{512 \cdot \log\left(\frac{16e^{1/16}}{\delta}\right)} + 256L^k \left(\frac{1}{m} + \frac{1}{n}\right) \log\left(\frac{16e^{1/16}}{\delta}\right).$$
(29)

As a reminder, the randomness in the above statements only comes from $\pi \sim \text{Unif}(S_{m+n})$.

H.4.3 High-probability upper bound for $T_{d,k,C_{\epsilon}}^{\pi}$

First, for each $(w, b) \in C_{\text{good}}$, apply (29) with $\delta \leftarrow \delta/N_{\text{good}}$. Then, apply union bound over all $(w, b) \in C_{\text{good}}$ to get an upper bound for $T_{d,k,\mathcal{C}_{\text{good}}}^{\pi}(z) = \max_{(w,b)\in\mathcal{C}_{\text{good}}} |X(w,b)|$. This gives us the following, where z is fixed: with probability at least $1 - \delta$,

$$T_{d,k,\mathcal{C}_{\text{good}}}^{\pi}(z) \le L^k \sqrt{\frac{1}{m} + \frac{1}{n}} \sqrt{512 \cdot \left\{ \log\left(\frac{16e^{1/16}}{\delta}\right) + \log N_{\text{good}} \right\}} + 256L^k \left(\frac{1}{m} + \frac{1}{n}\right) \left\{ \log\left(\frac{16e^{1/16}}{\delta}\right) + \log N_{\text{good}} \right\}$$

Recall that we have already got the upper bound for N_{good} in Lemma 9. Plugging in this upper bound, since $\frac{1}{m} + \frac{1}{n} \leq 1$, we have the following statement: there exists a constant $C''_{L,d,k}$ only depending on (L, d, k) such that with probability $1 - \delta$,

• for k = 0,

$$T_{d,k,\mathcal{C}_{\circ}}^{\pi}(z) \leq C_{L,d,k}^{\prime\prime} \sqrt{\frac{1}{m} + \frac{1}{n}} \Big\{ \log\left(\frac{1}{\delta}\right) + \log\left(1 + m + n\right) \Big\},\tag{30}$$

• and for $k \ge 1$,

$$T_{d,k,\mathcal{C}_{\epsilon}}^{\pi}(z) \le C_{L,d,k}^{\prime\prime} \sqrt{\frac{1}{m} + \frac{1}{n}} \Big\{ \log\left(\frac{1}{\delta}\right) + \log\left(1 + \frac{m+n}{\epsilon}\right) \Big\}.$$
(31)

H.4.4 High-probability upper bound for $T_{d,k}^{\pi}$

Still z is fixed. Combining the approximation error in Lemma 9 and (31) gives us that, with probability at least $1 - \delta$,

• for k = 0,

$$T_{d,k}^{\pi}(z) \le 0 + C_{L,d,k}^{\prime\prime} \sqrt{\frac{1}{m} + \frac{1}{n}} \Big\{ \log\left(\frac{1}{\delta}\right) + \log\left(1 + \frac{m+n}{\epsilon}\right) \Big\}.$$

• for k = 1,

$$T_{d,k}^{\pi}(z) \leq \left(\frac{1}{m} + \frac{1}{n}\right)\epsilon + C_{L,d,k}''\sqrt{\frac{1}{m} + \frac{1}{n}}\Big\{\log\left(\frac{1}{\delta}\right) + \log\left(1 + \frac{m+n}{\epsilon}\right)\Big\}.$$

When we choose $\epsilon = (\frac{1}{m} + \frac{1}{n})^{-1/2}$, the above inequality is rewritten as

$$T_{d,k}^{\pi}(z) \le \sqrt{\frac{1}{m} + \frac{1}{n} + C_{L,d,k}''} \sqrt{\frac{1}{m} + \frac{1}{n} \Big(\log\Big(\frac{1}{\delta}\Big) + \log(1+m+n) \Big)}.$$

Then, by plugging in $\delta = \alpha$, we can rewrite the above two bullet points in one statement as the following: there exists a constant $C_{L,d,k}$ only depending on (L,d,k) such that

$$T_{d,k}^{\pi}(z) \le C_{L,d,k} \sqrt{\frac{1}{m} + \frac{1}{n}} \left(\log\left(\frac{1}{\alpha}\right) + \log(m+n) \right) = t_{m,n,\alpha} \quad \text{with probability at least } 1 - \alpha,$$

where we recall the definition of $t_{m,n,\alpha}$ in (23). In other words, $t_{1-\alpha}(z) \leq t_{m,n,\alpha}$, which concludes the proof of Lemma 8. Therefore, this completes the proof of Theorem 12, presented in Appendix H.2.

I Auxilliary technical lemmas

Lemma 11. For $a, b \in \mathbb{R}$ and integer $k \geq 1$,

$$|a_{+}^{k} - b_{+}^{k}| \le k(a_{+}^{k-1} + b_{+}^{k-1})|a - b|,$$

with the convention that when k - 1 = 0, we set $0^0 = 1$.

Proof of Lemma 11. Without loss of generality, assume $a \leq b$. The function $x \mapsto x_+^k$ has derivative kx_+^{k-1} almost everywhere (even under our convention for when k-1=0). For every point $x \in [a,b]$, we have $|kx_+^{k-1}| \leq k(a_+^{k-1} + b_+^{k-1})$, because $|x_+| \leq \max\{a_+, b_+\}$. Thus, integrating the derivative along the path between a and b gives the result.

Lemma 12. Assume $\mu \in \mathcal{M}(\mathbb{S}^{d-1} \times [0, \infty))$. Then, for any $0 \leq m \leq k-1$ and any function $c(w, b) : \mathbb{S}^{d-1} \times [0, \infty) \to \mathbb{R}$ which is bounded by $|c(w, b)| \leq C$ for some constant C, we have

$$x \mapsto \int_{\mathbb{S}^{d-1} \times [0,\infty)} \left(c(w,b) (w^{\mathsf{T}} x - b)_+^{k-m} \right) d\mu(w,b)$$
(32)

is continuous at every $x \in \mathbb{R}^d$.

Proof of Lemma 12. Denote $f(x) = \int_{\mathbb{S}^{d-1} \times [0,\infty)} c(w,b) (w^{\mathsf{T}}x - b)_+^{k-m} d\mu(w,b)$. Let $|\mu|$ be the total variation measure of the signed measure μ . Note that $\|\mu\|_{\mathrm{TV}} = \||\mu|\|_{\mathrm{TV}}$. For $x, y \in \mathbb{R}^d$,

$$\begin{split} \left| f(x) - f(y) \right| &\leq \int_{\mathbb{S}^{d-1} \times [0,\infty)} |c(w,b)| \left| (w^{\mathsf{T}}x - b)_{+}^{k-m} - (w^{\mathsf{T}}y - b)_{+}^{k-m} \right| \,\mathrm{d}|\mu|(w,b) \\ &\leq (k-m) \int_{\mathbb{S}^{d-1} \times [0,\infty)} |c(w,b)| \left((w^{\mathsf{T}}x - b)_{+}^{k-m-1} + (w^{\mathsf{T}}y - b)_{+}^{k-m-1} \right) |w^{\mathsf{T}}(x-y)| \,\mathrm{d}|\mu|(w,b) \\ &\leq (k-m) \|x-y\| \int_{\mathbb{S}^{d-1} \times [0,\infty)} |c(w,b)| \left((\|w\|_{2}\|x\|_{2})^{k-m-1} + (\|w\|_{2}\|y\|_{2})^{k-m-1} \right) \,\mathrm{d}\mu(w,b) \\ &\leq (k-m) \|x-y\| \, C \, \|\mu\|_{\mathrm{TV}} \left(\|x\|_{2}^{k-m-1} + \|y\|_{2}^{k-m-1} \right), \end{split}$$

where in the second inequality we have used Lemma 11. Since the right-hand side goes to 0 as $y \to x$, the proof is complete.

J Sensitivity analysis: log transform

In Equation (10), which is the basis for the experiments in the main paper, we use a log transform of the IPM term in the criterion. The current section compares the performance of (local) optimizers of this problem, which we call the "log" problem, for short, to those of

$$\min_{\substack{f=f_{(a_j,w_j,b_j)_{j=1}^N}\\b_j\ge 0, j=1,\dots,N}} -\frac{1}{kN} \left| P_m(f) - Q_n(f) \right| + \frac{\lambda}{k} \left(\frac{1}{N} \sum_{i=1}^N |a_i| \|w_i\|_2^k \right)^2, \tag{33}$$

which we call the "no-log" problem, for short. In the above display, we can see that there is no log transform on the IPM term, but in addition, the penalty on the RTV^k seminorm has been squared. This is done because without such a transformation, one can show that the problem (IPM term plus RTV^k seminorm) does not attain its infimum, unless λ is chosen very carefully so that the infimum in zero. By squaring the penalty, the criterion in (33) is coercive, which guarantees (since it is also continuous) that it will attain its infimum. In both (10) and (33), we fix $\lambda = 1$, and use torch.optim.Adam with betas parameter (0.9, 0.99). We explore the behavior for different learning rates and numbers of iterations. Throughout, we focus on the "var-one" setting described in Table 1, and use N = 10 neurons. The summary of our findings is as follows.

- Optimizing the "log" problem typically results in a larger IPM value compared to the "no-log" problem, especially for larger k, d, and especially under the null.
- Optimizing the "log" and "no-log" problems generally results in similar ROC curves, however, for larger k, d, the "no-log" ROC curves can actually be slightly better. This actually appears to be driven by the fact that "no-log" optimization is *relatively worse* in terms of the IPM values obtained under the null, which gives it a slightly greater separation and slightly higher power.
- The "log" problem is more robust to the choice of learning rate, both in terms of the IPM values and ROC curves obtained, especially for larger k, d.
- The "log" problem typically shows faster convergence (number of iterations required to approach a local optimum), especially for larger k, d.

Figures 6–12 display the results from our sensitivity analyses. In each one, the figure caption explains the salient points about the setup and takeaways. We remark that results for k = 4 (not included for brevity) show qualitatively similar but even more extreme behavior.

K Sensitivity analysis: number of neurons

In this section, we now turn to investigate the role of the number of neurons N. We stick with the "log" problem (10) as in the experiments in the main text (with the previous section providing evidence that the log transform leads to greater robustness in many regards). The setup is essentially the same as that used in the last section, except that we vary the number of neurons from N = 1 to N = 20. The summary of our findings is as follows.

- Using a larger number of neurons often improves the achieved IPM value, especially for larger k, d; however, going past N = 10 does not seem to make a big difference in our experiments.
- Using a larger number of neurons generally results in a better ROC curve, especially for larger k, d; again, going past N = 10 does not seem to make a big difference in our experiments.
- Using a larger number of neurons can result in faster convergence, though not dramatically.
- Using a larger number of neurons is typically more robust to the choice of learning rate, both in terms of the IPM values and ROC curves obtained, especially for larger k, d.

Figure 13–17 display the results from our sensitivity analyses. As before, each figure caption presents the salient points about the setup and takeaways, and we remark that results for k = 4 (not included for brevity) show qualitatively similar but even more extreme behavior.

Figure 6: IPM values obtained by optimizing the "log" (x-axis) and "no-log" (y-axis) problems. Each point represents a set of samples drawn from P, Q, and the result of running T = 1200 iterations with learning rate 0.01, for each criterion. (This learning rate was chosen to be favorable to the "no-log" problem.) Points below the diagonal mean that the "log" criterion results in a larger IPM value, which we see is especially prominent for larger k, d, and more prominent under the null.

Figure 7: ROC curves from the "log" and "no-log" problems, for the same data as in Figure 6. They are very similar except for the largest k, d pair, where the "no-log" ROC curve is slightly better. Inspecting Figure 6b, this is likely due to the fact that the null IPM values here are relatively smaller (poorer optimization).

Figure 8: IPM values corresponding to learning rates 0.01 (x-axis) and 10 (y-axis), for the same data as in Figure 6. There is little to no difference for the "log" problem (points near the diagonal), and a much bigger difference for the "no-log" problem, where in most cases the larger learning rate is worse (points below the diagonal).

Figure 9: IPM values corresponding to learning rates 0.01 (x-axis) and 10 (y-axis), for the same data as in Figure 6. There is little to moderate for the "log" problem (points around the diagonal), and a much bigger difference for the "no-log" problem, where in nearly all cases the larger learning rate is clearly much worse (points below the diagonal).

Figure 10: IPM values corresponding to learning rates from 0.01 to 10, for the same data as in Figure 6. In all cases, the "log" problem results in much more stable ROC curves with respect to the choice of learning rate, especially for larger k, d.

Figure 11: IPM value as a function of iteration for the "log" and "no-log" problems. (cf. the y-axis says the MMD instead of the IPM.) This is done over 20 repetitions (draws of samples from P, Q); solid lines represent the median and shaded areas the interquartile range (computed over repetitions), at each iteration. For smaller k, d, the "log" and "no-log" convergence speeds are fairly similar, but for larger k, d, the "no-log" convergence speeds are fairly similar, but for larger k, d, the "no-log" convergence speed is clearly slower. For the larger learning rate, we can also see that in many instances the "no-log" iterations fail to converge, and bounce up and down in IPM value, without staying at a local optimum.

Figure 12: ROC curves for the "log" and "no-log" problems, at different iteration numbers, over the same data in Figure 11. The "log" curves are quite robust to the choice of the number of iterations, whereas the "no-log" ones are not, especially for larger k, d. In other words, the "log" problem typically shows faster convergence (number of iterations required to approach a local optimum), especially for larger k, d.

Figure 13: IPM values when N = 10 (x-axis) and N = 2 (y-axis). Each point represents a set of samples drawn from P, Q, and the result of running T = 1200 iterations with learning rate 0.5, when k = 3. Points below the diagonal mean that N = 10 results in a larger IPM value, which we see is especially prominent for larger k, d, and more prominent under the null.

Figure 14: IPM values when N = 10 (x-axis) and N = 20 (y-axis). The setup is as above. Here we see no clear improvement in moving from N = 10 to N = 20 neurons.

Figure 15: ROC curves as the number of neurons varies from N = 1 to N = 20, for the same data as in Figure 13. The ROC curves look generally better for a larger number of neurons, especially for larger k, d.

Figure 16: IPM value as a function of iteration for N = 10 and N = 2 neurons. This is carried out over 20 repetitions (draws of samples from P, Q); solid lines represent the median and shaded areas the interquartile range (computed over repetitions), at each iteration. The convergence speed is sometimes slightly faster with N = 10 neurons but the differences are not large.

Figure 17: ROC curves as the number of neurons varies from N = 1 to N = 20, and for varying learning rates, for the same data as in Figure 13. A larger number of neurons is typically more robust to the choice of learning rate, especially for larger d.

L Comparison to polynomial kernel MMD

Figure 18 compares the RKS tests to the polynomial kernel MMD tests using the data as in Figure 4 in the main paper. We have chosen linear, quadratic, and cubic degree polynomials by (rough) analogy to the use of k = 1, 2, 3 in RKS. The colors of the RKS test of degree k and kernel MMD with a polynomial kernel of degree k are chosen to match. However, we can see that these two frameworks produce tests with generally quite different behaviors: for example, compare the behavior of RKS with k = 1 to kernel MMD with linear kernel, in the "t-coord" setting.

Altogether, the behavior of kernel MMD with polynomial kernel is as expected: it performs quite well when $P \neq Q$ differ according to the alternative (moment difference) that is anticipated by the test, and is otherwise essentially powerless.

Figure 18: ROC curves across the same experimental settings as in Figure 4, but now with kernel MMD using a polynomial kernel of linear, quadratic, and cubic degrees. The ROC curve from the likelihood ratio test is also shown, indicating the best possible performance in each case (it depends on oracle knowledge of P, Q).