
A Higher-Order Kolmogorov-Smirnov Test

Veeranjaneyulu Sadhanala¹ Yu-Xiang Wang² Aaditya Ramdas¹ Ryan J. Tibshirani¹
¹Carnegie Mellon University ²University of California at Santa Barbara

Abstract

We present an extension of the Kolmogorov-Smirnov (KS) two-sample test, which can be more sensitive to differences in the tails. Our test statistic is an integral probability metric (IPM) defined over a higher-order total variation ball, recovering the original KS test as its simplest case. We give an exact representation result for our IPM, which generalizes the fact that the original KS test statistic can be expressed in equivalent variational and CDF forms. For small enough orders ($k \leq 5$), we develop a linear-time algorithm for computing our higher-order KS test statistic; for all others ($k \geq 6$), we give a nearly linear-time approximation. We derive the asymptotic null distribution for our test, and show that our nearly linear-time approximation shares the same asymptotic null. Lastly, we complement our theory with numerical studies.

1 INTRODUCTION

The Kolmogorov-Smirnov (KS) test (Kolmogorov, 1933; Smirnov, 1948) is a classical and celebrated tool for nonparametric hypothesis testing. Let $x_1, \dots, x_m \sim P$ and $y_1, \dots, y_n \sim Q$ be independent samples. Let $X_{(m)}$ and $Y_{(n)}$ denote the two sets of samples, and also let $Z_{(N)} = X_{(m)} \cup Y_{(n)} = \{z_1, \dots, z_N\}$, where $N = m + n$. The two-sample KS test statistic is defined as

$$\max_{z \in Z_{(m+n)}} \left| \frac{1}{m} \sum_{i=1}^m 1\{x_i \leq z\} - \frac{1}{n} \sum_{i=1}^n 1\{y_i \leq z\} \right|. \quad (1)$$

In words, this measures the maximum absolute difference between the empirical cumulative distribution functions (CDFs) of $X_{(m)}$ and $Y_{(n)}$, across all points in the joint sample $Z_{(m+n)}$. Naturally, the two-sample

KS test rejects the null hypothesis of $P = Q$ for large values of the statistic. The statistic (1) can also be written in the following variational form:

$$\sup_{f: \text{TV}(f) \leq 1} |\mathbb{P}_m f - \mathbb{Q}_n f|, \quad (2)$$

where $\text{TV}(\cdot)$ denotes total variation, and we define the empirical expectation operators $\mathbb{P}_m, \mathbb{Q}_n$ via

$$\mathbb{P}_m f = \frac{1}{m} \sum_{i=1}^m f(x_i) \quad \text{and} \quad \mathbb{Q}_n f = \frac{1}{n} \sum_{i=1}^n f(y_i).$$

Later, we will give a general representation result that implies the equivalence of (1) and (2) as a special case.

The KS test is a fast, general-purpose two-sample non-parametric test. But being a general-purpose test also means that it is systematically less sensitive to some types of differences, such as tail differences (Bryson, 1974). Intuitively, this is because the empirical CDFs of $X_{(m)}$ and $Y_{(n)}$ must both tend to 0 as $z \rightarrow -\infty$ and to 1 as $z \rightarrow \infty$, so the gap in the tails will not be large.

The insensitivity of the KS test to tail differences is well-known. Several authors have proposed modifications to the KS test to improve its tail sensitivity, based on variance-reweighting (Anderson and Darling, 1952), or Renyi-type statistics (Mason and Schuenemeyer, 1983; Calitz, 1987), to name a few ideas. In a different vein, Wang et al. (2014) recently proposed a higher-order extension of the KS two-sample test, which replaces the total variation constraint on f in (2) with a total variation constraint on a derivative of f . These authors show empirically that, in some cases, this modification can lead to better tail sensitivity. In the current work, we refine the proposal of Wang et al. (2014), and give theoretical backing for this new test.

A Higher-Order KS Test. Our test statistic has the form of an integral probability metric (IPM). For a function class \mathcal{F} , the IPM between distributions P and Q , with respect to \mathcal{F} , is defined as (Muller, 1997)

$$\rho(P, Q; \mathcal{F}) = \sup_{f \in \mathcal{F}} |\mathbb{P}f - \mathbb{Q}f| \quad (3)$$

where we define the expectation operators \mathbb{P}, \mathbb{Q} by

$$\mathbb{P}f = \mathbb{E}_{X \sim P}[f(X)] \quad \text{and} \quad \mathbb{Q}f = \mathbb{E}_{Y \sim Q}[f(Y)].$$

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

For a given function class \mathcal{F} , the IPM $\rho(\cdot, \cdot; \mathcal{F})$ is a pseudometric on the space of distributions. Note that the KS test in (2) is precisely $\rho(P_m, Q_n; \mathcal{F}_0)$, where P_m, Q_n are the empirical distributions of $X_{(m)}, Y_{(n)}$, respectively, and $\mathcal{F}_0 = \{f : \text{TV}(f) \leq 1\}$.

Consider an IPM given by replacing \mathcal{F}_0 with $\mathcal{F}_k = \{f : \text{TV}(f^{(k)}) \leq 1\}$, for an integer $k \geq 1$ (where we write $f^{(k)}$ for the k th weak derivative of f). Some motivation is as follows. In the case $k = 0$, we know that the *witness functions* in the KS test (2), i.e., the functions in \mathcal{F}_0 that achieve the supremum, are piecewise constant step functions (cf. the equivalent representation (1)). These functions can only have so much action in the tails. By moving to \mathcal{F}_k , which is essentially comprised of the k th order antiderivative of functions in \mathcal{F}_0 , we should expect that the witness functions over \mathcal{F}_k are k th order antiderivatives of piecewise constant functions, i.e., k th degree piecewise polynomial functions, which can have much more sensitivity in the tails.

But simply replacing \mathcal{F}_0 by \mathcal{F}_k and proposing to compute $\rho(P_m, Q_n; \mathcal{F}_k)$ leads to an ill-defined test. This is due to the fact that \mathcal{F}_k contains all polynomials of degree k . Hence, if the i th moments of P_m, Q_n differ, for any $i \in [k]$ (where we abbreviate $[a] = \{1, \dots, a\}$ for an integer $a \geq 1$), then $\rho(P_m, Q_n; \mathcal{F}_k) = \infty$.

As such, we must modify \mathcal{F}_k to control the growth of its elements. While there are different ways to do this, not all result in computable IPMs. The approach we take yields an exact representer theorem (generalizing the equivalence between (1) and (2)). Define

$$\begin{aligned} \mathcal{F}_k = \{f : \text{TV}(f^{(k)}) \leq 1, \\ f^{(j)}(0) = 0, j \in \{0\} \cup [k-1], \\ f^{(k)}(0+) = 0 \text{ or } f^{(k)}(0-) = 0\}. \end{aligned} \quad (4)$$

Here $f^{(k)}(0+)$ and $f^{(k)}(0-)$ denote one-sided limits at 0 from above and below, respectively. Informally, the functions in \mathcal{F}_k are pinned down at 0, with all lower-order derivatives (and the limiting k th derivative from the right or left) equal to 0, which limits their growth. Now we define the k th-order KS test statistic as

$$\rho(P_m, Q_n; \mathcal{F}_k) = \sup_{f \in \mathcal{F}_k} |\mathbb{P}_m f - \mathbb{Q}_n f|. \quad (5)$$

An important remark is that for $k = 0$, this recovers the original KS test statistic (2), because \mathcal{F}_0 contains all step functions of the form $g_t(x) = 1\{x \leq t\}$, $t \geq 0$.

Another important remark is that for any $k \geq 0$, the function class \mathcal{F}_k in (4) is “rich enough” to make the IPM in (5) a metric. We state this formally next; its proof, as with all other proofs, is in the supplement.

Proposition 1. *For any $k \geq 0$, and any P, Q with k moments, $\rho(P, Q; \mathcal{F}_k) = 0$ if and only if $P = Q$.*

Motivating Example. Figure 1 shows the results of a simple simulation comparing the proposed higher-order tests (5), of orders $k = 1$ through 5, against the usual KS test (corresponding to $k = 0$). For the simulation setup, we used $P = N(0, 1)$ and $Q = N(0, 1.44)$. For 500 repetitions, “alternative repetitions”, we drew $m = 100$ samples from P , drew $n = 100$ samples from Q , and computed test statistics; for another 500 repetitions, called “null repetitions”, we drew both sets of samples from P , and again computed test statistics. Then for each test, we varied the rejection threshold, calculated its true positive rate as the fraction of rejections made on the alternative repetitions, and calculated its false positive rate similarly using the null repetitions. The oracle ROC curve corresponds to the likelihood ratio test (which knows the exact distributions P, Q). We can see that power of the higher-order KS test improves as we increase the order from $k = 0$ up to $k = 3$, then stops improving by $k = 4, 5$.

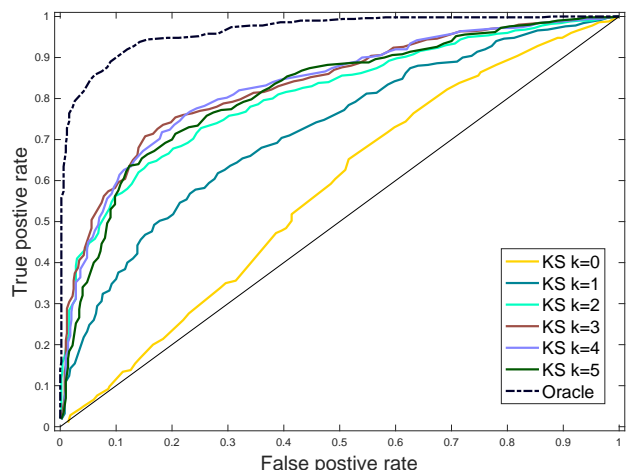


Figure 1: ROC curves from an experiment comparing the proposed higher-order KS tests in (5) (for various k) to the usual KS test, when $P = N(0, 1)$ and $Q = N(0, 1.44)$.

Figure 2 displays the witness function (which achieves the supremum in (5)) for a large-sample version of the higher-order KS test, across orders $k = 0$ through 5. We used the same distributions as in Figure 1, but now $n = m = 10^4$. We will prove in Section 2 that, for the k th order test, the witness function is always a k th degree piecewise polynomial (in fact, a rather simple one, of the form $g_t(x) = (x - t)_+^k$ or $g_t(x) = (t - x)_+^k$ for a knot t). Recall the underlying distributions P, Q here have different variances, and we can see from their witness functions that all higher-order KS tests choose to put weight on tail differences. Of course, the power of any test of is determined by the size of the statistic under the alternative, relative to typical fluctuations under the null. As we place greater and greater weight on tails, it turns out in this particular setting that we

see diminishing returns at $k = 4, 5$, which means the null fluctuations must be too great.

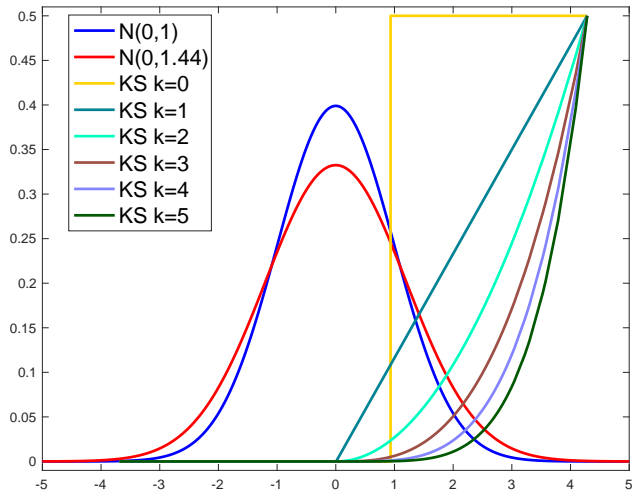


Figure 2: Witness functions (normalized for plotting purposes) for the higher-order KS tests, when $P = N(0, 1)$ and $Q = N(0, 1.44)$. They are always of piecewise polynomial form; and here they all place weight on tail differences.

Summary of Contributions. Our contributions in this work are as follows.

- We develop an exact representer theorem for the higher-order KS test statistic (5). This enables us to compute the test statistic in linear-time, for all $k \leq 5$. For $k \geq 6$, we develop a nearly linear-time approximation to the test statistic.
- We derive the asymptotic null distribution of the our higher-order KS test statistic, based on empirical process theory. For $k \geq 6$, our approximation to the test statistic has the same asymptotic null.
- We provide concentration tail bounds for the test statistic. Combined with the metric property from Proposition 1, this shows that our higher-order KS test is asymptotically powerful against any pair of fixed, distinct distributions P, Q .
- We perform extensive numerical studies to compare the newly proposed tests with several others.

Other Related Work. Recently, IPMs have been gaining in popularity due in large part to energy distance tests (Szekely and Rizzo, 2004; Baringhaus and Franz, 2004) and kernel maximum mean discrepancy (MMD) tests (Gretton et al., 2012), and in fact, there is an equivalence between the two classes (Sejdinovic et al., 2013). An IPM with a judicious choice of \mathcal{F} gives rise to a number of common distances between distributions, such as Wasserstein distance or total variation

(TV) distance. While IPMs look at differences $dP - dQ$, tests based on ϕ -divergences (such as Kullback-Leibler, or Hellinger) look at ratios dP/dQ , but can be hard to efficiently estimate in practice (Sriperumbudur et al., 2009). The TV distance is the only IPM that is also a ϕ -divergence, but it is impossible to estimate.

There is also a rich class of nonparametric tests based on graphs. Using minimum spanning trees, Friedman and Rafsky (1979) generalized both the Wald-Wolfowitz runs test and the KS test. Other tests are based on k -nearest neighbors graphs (Schilling, 1986; Henze, 1988) or matchings (Rosenbaum, 2005). The Mann-Whitney-Wilcoxon test has a multivariate generalization using the concept of data depth (Liu and Singh, 1993). Bhattacharya (2016) established that many computationally efficient graph-based tests have suboptimal statistical power, but some inefficient ones have optimal scalings.

Different computational-statistical tradeoffs were also discovered for IPMs (Ramdas et al., 2015b). Further, as noted by Janssen (2000) (in the context of one-sample testing), every nonparametric test is essentially powerless in an infinity of directions, and has nontrivial power only against a finite subspace of alternatives. In particular, this implies that no single nonparametric test can uniformly dominate all others; improved power in some directions generally implies weaker power in others. This problem only gets worse in high-dimensional settings (Ramdas et al., 2015a; Arias-Castro et al., 2018). Therefore, the question of which test to use for a given problem must be guided by a combination of simulations, computational considerations, a theoretical understanding of the pros/cons of each test, and a practical understanding of the data at hand.

Outline. In Section 2, we give computational details for the higher-order KS test statistic (5). We derive its asymptotic null in Section 3, and give concentration bounds (for the statistic around the population-level IPM) in Section 4. We give numerical experiments in Section 5, and conclude in Section 6 with a discussion.

2 COMPUTATION

Write $T = \rho(P_m, Q_n; \mathcal{F}_k)$ for the test statistic in (5). In this section, we derive a representer theorem for T , develop a linear-time algorithm for $k \leq 5$, and a nearly linear-time approximation for $k \geq 6$.

2.1 Representer Theorem

The higher-order KS test statistic in (5) is defined by an infinite-dimensional maximization over \mathcal{F}_k in (4). Fortunately, we can restrict our attention to a simpler function class, as we show next.

Theorem 1. Fix $k \geq 0$. Let $g_t^+(x) = (x - t)_+^k/k!$ and $g_t^-(x) = (t - x)_+^k/k!$ for $t \in \mathbb{R}$, where we write $(a)_+ = \max\{a, 0\}$. For the statistic T defined by (5),

$$T = \max \left\{ \sup_{t \geq 0} |(\mathbb{P}_m - \mathbb{Q}_n)g_t^+|, \sup_{t \leq 0} |(\mathbb{P}_m - \mathbb{Q}_n)g_t^-| \right\}. \quad (6)$$

The proof of this theorem uses a key result from [Mammen \(1991\)](#), where it is shown that we can construct a spline interpolant to a given function at given points, such that its higher-order total variation is no larger than that of the original function.

Remark 1. When $k = 0$, note that for $t \geq 0$,

$$\begin{aligned} |(\mathbb{P}_m - \mathbb{Q}_n)g_t^+| &= \left| \frac{1}{m} \sum_{i=1}^m \mathbf{1}\{x_i > t\} - \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{y_i > t\} \right| \\ &= \left| \frac{1}{m} \sum_{i=1}^m \mathbf{1}\{x_i \leq t\} - \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{y_i \leq t\} \right| \end{aligned}$$

and similarly for $t \leq 0$, $|(\mathbb{P}_m - \mathbb{Q}_n)g_t^-|$ reduces to the same expression in the second line above. As we vary t from $-\infty$ to ∞ , this only changes at values $t \in Z_{(N)}$, which shows (6) and (1) are the same, i.e., Theorem 1 recovers the equivalence between (2) and (1).

Remark 2. For general $k \geq 0$, we can interpret (6) as a comparison between truncated k th order moments, between the empirical distributions P_m and Q_n . The test statistic T the maximum over all possible truncation locations t . The critical aspect here is *truncation*, which makes the higher-order KS test statistic a metric (recall Proposition 1). A comparison of moments, alone, would not be enough to ensure such a property.

Theorem 1 itself does not immediately lead to an algorithm for computing T , as the range of t considered in the suprema is infinite. However, through a bit more work, detailed in the next two subsections, we can obtain an exact linear-time algorithm for all $k \leq 5$, and a linear-time approximation for $k \geq 6$.

2.2 Linear-Time Algorithm for $k \leq 5$

The key fact that we will exploit is that the criterion in (6), as a function of t , is a piecewise polynomial of order k with knots in $Z_{(N)}$. Assume without a loss of generality that $z_1 < \dots < z_N$. Also assume without a loss of generality that $z_1 \geq 0$ (this simplifies notation, and the general case follows by the repeating the same arguments separately for the points in $Z_{(N)}$ on either side of 0). Define $c_i = \mathbf{1}\{z_i \in X_{(m)}\}/m - \mathbf{1}\{z_i \in Y_{(n)}\}/n$, $i \in [N]$, and

$$\phi_i(t) = \frac{1}{k!} \sum_{j=i}^N c_j (z_j - t)^k, \quad i \in [N]. \quad (7)$$

Then the statistic in (6) can be succinctly written as

$$T = \max_{i \in [N]} \sup_{t \in [z_{i-1}, z_i]} \phi_i(t), \quad (8)$$

where we let $z_0 = 0$ for convenience. Note each $\phi_i(t)$, $i \in [N]$ is a k th degree polynomial. We can compute a representation for these polynomials efficiently.

Lemma 1. Fix $k \geq 0$. The polynomials in (7) satisfy the recurrence relations

$$\phi_i(t) = \frac{1}{k!} c_i (z_i - t)^k + \phi_{i+1}(t), \quad i \in [N]$$

(where $\phi_{N+1} = 0$). Given the monomial expansion

$$\phi_{i+1}(t) = \sum_{\ell=0}^k a_{i+1,\ell} t^\ell,$$

we can compute an expansion for ϕ_i , with coefficients $a_{i\ell}$, $\ell \in \{0\} \cup [k]$, in $O(1)$ time. So we can compute all coefficients $a_{i,\ell}$, $i \in [N]$, $\ell \in \{0\} \cup [k]$ in $O(N)$ time.

To compute T in (8), we must maximize each polynomial ϕ_i over its domain $[z_{i-1}, z_i]$, for $i \in [N]$, and then compare maxima. Once we have computed a representation for these polynomials, as Lemma 1 ensures we can do in $O(N)$ time, we can use this to analytically maximize each polynomial over its domain, provided the order k is small enough. Of course, maximizing a polynomial over an interval can be reduced to computing the roots of its derivative, which is an analytic computation for any $k \leq 5$ (since the roots of any quartic have a closed-form, see, e.g., [Rosen 1995](#)). The next result summarizes.

Proposition 2. For any $0 \leq k \leq 5$, the test statistic in (8) can be computed in $O(N)$ time.

Maximizing a polynomial of degree $k \geq 6$ is not generally possible in closed-form. However, developments in semidefinite optimization allow us to approximate its maximum efficiently, investigated next.

2.3 Linear-Time Approximation for $k \geq 6$

Seminal work of [Shor \(1998\)](#); [Nesterov \(2000\)](#) shows that the problem of maximizing a polynomial over an interval can be cast as a semidefinite program (SDP). The number of variables in this SDP depends only on the polynomial order k , and all constraint functions are self-concordant. Using say an interior point method to solve this SDP, therefore, leads to the following result.

Proposition 3. Fix $k \geq 6$ and $\epsilon > 0$. For each polynomial in (7), we can compute an ϵ -approximation to its maximum in $c_k \log(1/\epsilon)$ time, for a constant $c_k > 0$ depending only on k . As we can compute a representation for all these polynomials in $O(N)$ time (Lemma 1), this means we can compute an ϵ -approximation to the statistic in (6) in $O(N \log(1/\epsilon))$ time.

Remark 3. Let T_ϵ denote the ϵ -approximation from Proposition 3. Under the null $P = Q$, we would need to have $\epsilon = o(1/\sqrt{N})$ in order for the approximation T_ϵ to share the asymptotic null distribution of T , as we will see in Section 3.3. Taking say, $\epsilon = 1/N$, the statistic $T_{1/N}$ requires $O(N \log N)$ computational time, and this is why in various places we make reference to a *nearly* linear-time approximation when $k \geq 6$.

2.4 Simple Linear-Time Approximation

We conclude this section by noting a simple approximation to (6) given by

$$T^* = \max \left\{ \max_{t \in Z_{(N)}^0, t \geq 0} |(\mathbb{P}_m - \mathbb{Q}_n)g_t^+|, \max_{t \in Z_{(N)}^0, t \leq 0} |(\mathbb{P}_m - \mathbb{Q}_n)g_t^-| \right\}, \quad (9)$$

where $Z_{(N)}^0 = \{0\} \cup Z_{(N)}$. Clearly, for $k = 0$ or 1 , the maximizing t in (6) must be one of the sample points $Z_{(N)}$, so $T^* = T$ and there is no approximation error in (9). For $k \geq 2$, we can control the error as follows.

Lemma 2. For $k \geq 2$, the statistics in (6), (9) satisfy

$$T - T^* = \frac{\delta_N}{(k-1)!} \left(\frac{1}{m} \sum_{i=1}^m |x_i|^{k-1} + \frac{1}{n} \sum_{i=1}^n |y_i|^{k-1} \right),$$

where δ_N is the maximum gap between sorted points in $Z_{(N)}^0$.

Remark 4. We would need to have $\delta_N = o_P(1/\sqrt{N})$ in order for T^* to share the asymptotic null of T , see again Section 3.3 (this is assuming that P has $k-1$ moments, so the sample moments concentrate for large enough N). This will not be true of δ_N , the maximum gap, in general. But it does hold when P is continuous, having compact support, and a density bounded from below on its support; here, in fact, $\delta_N = o_P(\log N/N)$ (see, e.g., Wang et al. 2014).

Although it does not have the strong guarantees of the approximation from Proposition 3, the statistic in (9) is simple and efficient—we must emphasize that it can be computed in $O(N)$ linear time, as a consequence of Lemma 1 (the evaluations of $\phi_i(t)$ at the sample points $t \in Z_{(N)}$ are the constant terms a_{i0} , $i \in [N]$ in their monomial expansions)—and is likely a good choice for most practical purposes.

3 ASYMPTOTIC NULL

To study the asymptotic null distribution of the proposed higher-order KS test, we will appeal to uniform central limit theorems (CLTs) from the empirical process theory literature, reviewed here for completeness.

For functions f, g in a class \mathcal{F} , let $\mathbb{G}_{P, \mathcal{F}}$ denote a Gaussian process indexed by \mathcal{F} with mean and covariance

$$\mathbb{E}(\mathbb{G}_{P, \mathcal{F}} f) = 0, \quad f \in \mathcal{F},$$

$$\text{Cov}(\mathbb{G}_{P, \mathcal{F}} f, \mathbb{G}_{P, \mathcal{F}} g) = \text{Cov}_{X \sim P}(f(X), g(X)), \quad f, g \in \mathcal{F}.$$

For functions l, u , let $[l, u]$ denote the set of functions $\{f : l(x) \leq f(x) \leq u(x), \text{ for all } x\}$. Call $[l, u]$ a *bracket* of size $\|u - l\|_2$, where $\|\cdot\|_2$ denotes the $L_2(P)$ norm, defined as

$$\|f\|_2^2 = \int f(x)^2 dP(x).$$

Finally, let $N_{[]}(\epsilon, \|\cdot\|_2, \mathcal{F})$ be the smallest number of ϵ -sized brackets that are required to cover \mathcal{F} . Define the bracketing integral of \mathcal{F} as

$$J_{[]}(\|\cdot\|_2, \mathcal{F}) = \int_0^1 \sqrt{\log N_{[]}(\epsilon, \|\cdot\|_2, \mathcal{F})} d\epsilon.$$

Note that this is finite when $\log N_{[]}(\epsilon, \|\cdot\|_2, \mathcal{F})$ grows slower than $1/\epsilon^2$. We now state an important uniform CLT from empirical process theory.

Theorem 2 (Theorem 11.1.1 in Dudley 1999). *If \mathcal{F} is a class of functions with finite bracketing integral, then when $P = Q$ and $m, n \rightarrow \infty$, the process*

$$\sqrt{\frac{mn}{m+n}} \{\mathbb{P}_m f - \mathbb{Q}_n f\}_{f \in \mathcal{F}}$$

converges weakly to the Gaussian process $\mathbb{G}_{P, \mathcal{F}}$. Hence,

$$\sqrt{\frac{mn}{m+n}} \sup_{f \in \mathcal{F}} |\mathbb{P}_m f - \mathbb{Q}_n f| \xrightarrow{d} \sup_{f \in \mathcal{F}} |\mathbb{G}_{P, \mathcal{F}} f|.$$

3.1 Bracketing Integral Calculation

To derive the asymptotic null of the higher-order KS test, based on its formulation in (5), and Theorem 2, we would need to bound the bracketing integral of \mathcal{F}_k . While there are well-known entropy (log covering) number bounds for related function classes (e.g., Birman and Solomyak 1967; Babenko 1979), and the conversion from covering to bracketing numbers is standard, these results unfortunately require the function class to be uniformly bounded in the sup norm, which is certainly not true of \mathcal{F}_k .

Note that the representer result in (6) can be written as $T = \rho(P_m, Q_n; \mathcal{G}_k)$, where

$$\mathcal{G}_k = \{g_t^+ : t \geq 0\} \cup \{g_t^- : t \leq 0\}. \quad (10)$$

We can hence instead apply Theorem 2 to \mathcal{G}_k , whose bracketing number can be bounded by direct calculation, assuming enough moments on P .

Lemma 3. Fix $k \geq 0$. Assume $\mathbb{E}_{X \sim P} |X|^{2k+\delta} \leq M < \infty$, for some $\delta > 0$. For the class \mathcal{G}_k in (10), there is a constant $C > 0$ depending only on k, δ such that

$$\log N_{[]}(\epsilon, \|\cdot\|_2, \mathcal{G}_k) \leq C \log \frac{M^{1 + \frac{\delta(k-1)}{2k+\delta}}}{\epsilon^{2+\delta}}.$$

3.2 Asymptotic Null for Higher-Order KS

Applying Theorem 2 and Lemma 3 to the higher-order KS test statistic (6) leads to the following result.

Theorem 3. *Fix $k \geq 0$. Assume $\mathbb{E}_{X \sim P} |X|^{2k+\delta} < \infty$, for some $\delta > 0$. When $P = Q$, the test statistic in (6) satisfies, as $m, n \rightarrow \infty$,*

$$\sqrt{\frac{mn}{m+n}} T \xrightarrow{d} \sup_{g \in \mathcal{G}_k} |\mathbb{G}_{P,k} g|,$$

where $\mathbb{G}_{P,k}$ is an abbreviation for the Gaussian process indexed by the function class \mathcal{G}_k in (10).

Remark 5. When $k = 0$, note that for $t \geq s \geq 0$, the covariance function is

$$\text{Cov}_{X \sim P}(1\{X > s\}, 1\{X > t\}) = F_P(s)(1 - F_P(t)),$$

where F_P denotes the CDF of P . For $s \leq t \leq 0$, the covariance function is again equal to $F_P(s)(1 - F_P(t))$. The supremum of this Gaussian process over $t \in \mathbb{R}$ is that of a Brownian bridge, so Theorem 3 recovers the well-known asymptotic null distribution of the KS test, which (remarkably) does not depend on P .

Remark 6. When $k \geq 1$, it is not clear how strongly the supremum of the Gaussian process from Theorem 3 depends on P ; it appears it must depend on the first k moments of P , but is not clear whether it *only* depends on these moments. Section 5 investigates empirically. Currently, we do not have a precise understanding of whether the asymptotic null is useable in practice, and we suggest using a permutation null instead.

3.3 Asymptotic Null Under Approximation

The approximation from Proposition 3 shares the same asymptotic null, provided $\epsilon > 0$ is small enough.

Corollary 1. *Fix $k \geq 0$. Assume $\mathbb{E}_{X \sim P} |X|^{2k+\delta} < \infty$, for some $\delta > 0$. When $P = Q$, as $m, n \rightarrow \infty$ such that m/n converges to a positive constant, the test statistic T_ϵ from Proposition 3 converges at a \sqrt{N} -rate to the supremum of the same Gaussian process in Theorem 3, provided $\epsilon = o(1/\sqrt{N})$.*

The approximation in (9) shares the same asymptotic null, provided P is continuous with compact support.

Corollary 2. *Fix $k \geq 0$. Assume that P is continuous, compactly supported, with density bounded from below on its support. When $P = Q$, as $m, n \rightarrow \infty$ such that m/n converges to a positive constant, the test statistic T^* in (9) converges at a \sqrt{N} -rate to the supremum of the same Gaussian process in Theorem 3.*

4 TAIL CONCENTRATION

We examine the convergence of our test statistics to their population analogs. In general, if the population-

level IPM $\rho(P, Q; \mathcal{F}_k)$ is large, then the concentration bounds below will imply that the empirical statistic $\rho(P_m, Q_n; \mathcal{F}_k)$ will be large for m, n sufficiently large, and the test will have power.

We first review the necessary machinery, again from empirical process theory. For $p \geq 1$, and a function f of a random variable $X \sim P$, recall the $L_p(P)$ norm is defined as $\|f\|_p = [\mathbb{E}(f(X)^p)]^{1/p}$. For $p > 0$, recall the *exponential Orlicz norm* of order p is defined as

$$\|f\|_{\Psi_p} = \inf \{t > 0 : \mathbb{E}[\exp(|X|^p/t^p)] - 1 \leq 1\}.$$

(These norms depend on the measure P , since they are defined in terms of expectations with respect to $X \sim P$, though this is not explicit in our notation.)

We now state an important concentration result.

Theorem 4 (Theorems 2.14.2 and 2.14.5 in [van der Vaart and Wellner 1996](#)). *Let \mathcal{F} be a class functions with an envelope function F , i.e., $f \leq F$ for all $f \in \mathcal{F}$. Define*

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

and abbreviate $J = J_{[]}(\|\cdot\|, \mathcal{F})$. For $p \geq 2$, if $\|F\|_p < \infty$, then for a constant $c_1 > 0$,

$$[\mathbb{E}(W^p)]^{1/p} \leq c_1 \left(\|F\|_2 J + n^{-1/2+1/p} \|F\|_p \right),$$

and for $0 < p \leq 1$, if $\|F\|_{\Psi_p} < \infty$, then for a constant $c_2 > 0$,

$$\|W\|_{\Psi_p} \leq c_2 \left(\|F\|_2 J + n^{-1/2} (1 + \log n)^{1/p} \|F\|_{\Psi_p} \right).$$

The two-sample test statistic $T = \rho(P_m, Q_n; \mathcal{G}_k)$ satisfies (following by a simple argument using convexity)

$$|T - \rho(P, Q; \mathcal{F}_k)| \leq \rho(P, P_m; \mathcal{F}_k) + \rho(Q, Q_n; \mathcal{F}_k).$$

The terms on the right hand side can each be bounded by Theorem 4, where we can use the envelope function $F(x) = |x|^k/k!$ for \mathcal{G}_k . Using Markov's inequality, we can then get a tail bound on the statistic.

Theorem 5. *Fix $k \geq 0$. Assume that P, Q both have p moments, where $p \geq 2$ and $p > 2k$. For the statistic in (6), for any $\alpha > 0$, with probability $1 - \alpha$,*

$$|T - \rho(P, Q; \mathcal{G}_k)| \leq c(\alpha) \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right),$$

where $c(\alpha) = c_0 \alpha^{-1/p}$, and $c_0 > 0$ is a constant. If P, Q both have finite exponential Orlicz norms of order $0 < p \leq 1$, then the above holds for $c(\alpha) = c_0 (\log(1/\alpha))^{1/p}$.

When we assume k moments, the population IPM for \mathcal{F}_k also has a representer in \mathcal{G}_k ; by Proposition 1, this implies $\rho(\cdot, \cdot; \mathcal{G}_k)$ is also a metric.

Corollary 3. Fix $k \geq 0$. Assuming P, Q both have k moments, $\rho(P, Q; \mathcal{F}_k) = \rho(P, Q; \mathcal{G}_k)$. Therefore, by Proposition 1, $\rho(\cdot, \cdot; \mathcal{G}_k)$ is a metric (over the space of distributions P, Q with k moments).

Putting this metric property together with Theorem 5 gives the following.

Corollary 4. Fix $k \geq 0$. For $\alpha_N = o(1)$ and $1/\alpha_N = o(N^{p/2})$, reject when the higher-order KS test statistic (6) satisfies $T > c(\alpha_N)(1/\sqrt{m} + 1/\sqrt{n})$, where $c(\cdot)$ is as in Theorem 5. For any P, Q that meet the moment conditions of Theorem 5, as $m, n \rightarrow \infty$ in such a way that m/n approaches a positive constant, we have type I error tending to 0, and power tending to 1, i.e., the higher-order KS test is asymptotically powerful.

5 NUMERICAL EXPERIMENTS

We present numerical experiments that examine the convergence of our test statistic to its asymptotic null, its power relative to other general purpose nonparametric tests, and its power when P, Q have densities with local differences. Experiments comparing to the MMD test with a polynomial kernel are deferred to the supplement, for space reasons.

Convergence to Asymptotic Null. In Figure 3, we plot histograms of finite-sample higher-order KS test statistics and their asymptotic null distributions, when $k = 1, 2$. We considered both $P = N(0, 1)$ and $P = \text{Unif}(-\sqrt{3}, \sqrt{3})$ (the uniform distribution standardized to have mean 0 and variance 1). For a total of 1000 repetitions, we drew two sets of samples from P , each of size $m = n = 2000$, then computed the test statistics. For a total of 1000 times, we also approximated the supremum of the Gaussian process from Theorem 3 via discretization. We see that the finite-sample statistics adhere closely to their asymptotic distributions. Interestingly, we also see that the distributions look roughly similar across all four cases considered. Future work will examine more thoroughly.

Comparison to General-Purpose Tests. In Figures 4 and 5, we compare the higher-order KS tests to the KS test, and other widely-used nonparametric tests from the literature: the kernel maximum mean discrepancy (MMD) test (Gretton et al., 2012) with a Gaussian kernel, the energy distance test (Szekely and Rizzo, 2004), and the Anderson-Darling test (Anderson and Darling, 1954). The simulation setup is the same as that in the introduction, where we considered P, Q with different variances, except here we study different means: $P = N(0, 1)$, $Q = N(0.2, 1)$, and different fourth moments: $P = N(0, 1)$, $Q = \text{Lap}(0, 1/\sqrt{2})$. The higher-order KS tests generally perform favorably, and

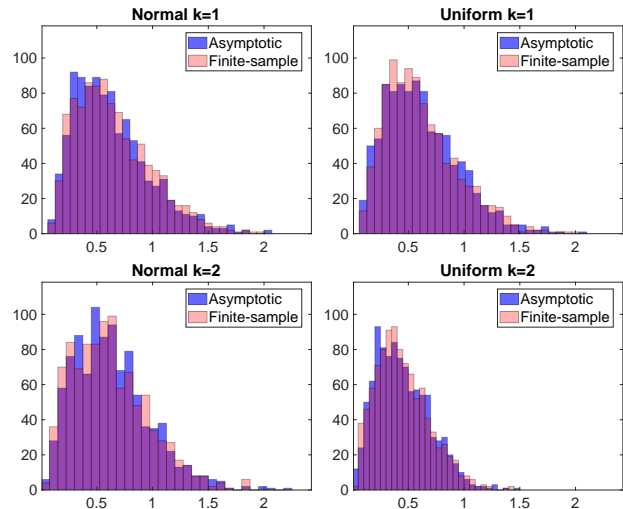


Figure 3: Histograms comparing finite-sample test statistics to their asymptotic null distribution.

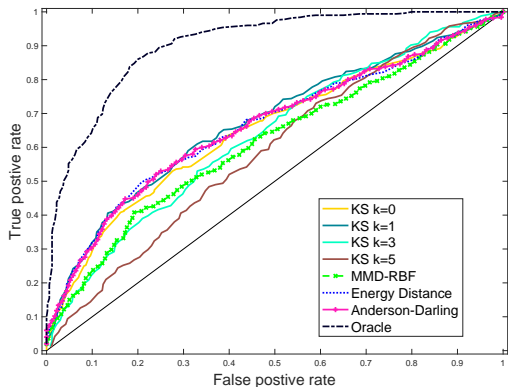
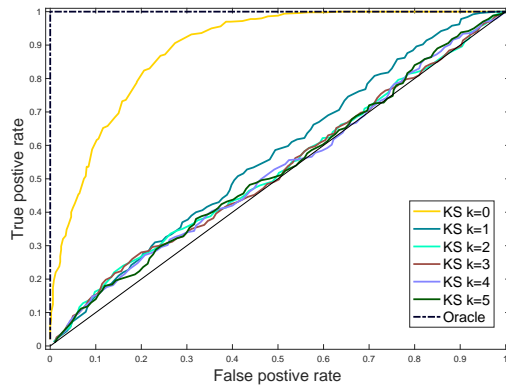
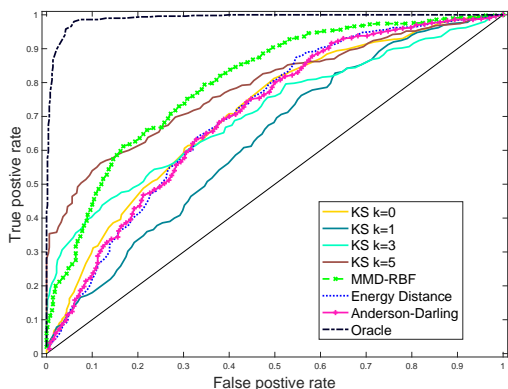
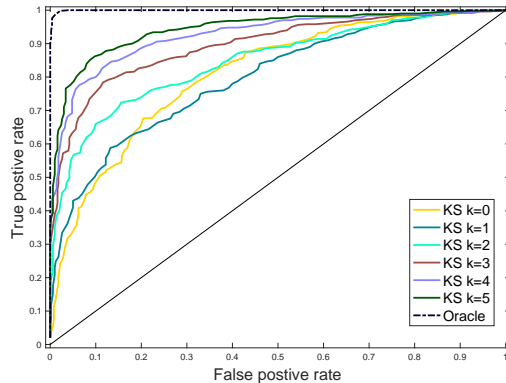
in each setting there is a choice of k that yields better power than KS. In the mean difference setting, this is $k = 1$, and the power degrades for $k = 3, 5$, likely because these tests are “smoothing out” the mean difference too much; see Proposition 4.

Local Density Differences. In Figures 6 and 7, we examine the higher-order KS tests and the KS test, in cases where P, Q have densities p, q such that $p - q$ has sharp local changes, and $m = n = 100$. Figure 6 displays a case where $p - q$ is piecewise constant with a few short departures from 0 (see the supplement for a plot). The KS test has large power, and the higher-order KS tests all perform poorly; in fact, the KS test here is more powerful than all of the commonly-used nonparametric tests that we tried (not shown).

Figure 7 displays a case where $p - q$ changes sharply in the right tail (see the supplement for a plot) and $m = n = 2000$. The power of the higher-order KS test appears to increase with k , likely because the witness functions are able to better concentrate on sharp departures for large k .

6 DISCUSSION

This paper began by noting the variational characterization of the classical KS test as an IPM with respect to functions of bounded total variation, and then proposed a generalization to higher-order total variation classes. This generalization was nontrivial, with subtleties arising in defining the right class of functions so that the statistic was finite and amenable for simplification via a representer result, challenges in computing the statistic efficiently, and challenges in studying asymptotic convergence and concentration due to the fact that the


Figure 4: ROC curves for $P = N(0, 1)$, $Q = N(0.2, 1)$.

Figure 6: ROC curves for piecewise constant $p - q$.

Figure 5: ROC curves for $P = N(0, 1)$, $Q = \text{Lap}(0, 1/\sqrt{2})$.

Figure 7: ROC curves for tail departure in $p - q$.

function class is not uniformly sup norm bounded. The resulting class of linear-time higher-order KS tests was shown empirically to be more sensitive to tail differences than the usual KS test, and to have competitive power relative to several other popular tests.

In future work, we intend to more formally study the power properties of our new higher-order tests relative to the KS test. The following is a lead in that direction. For $k \geq 1$, define I^k to be the k th order integral operator, acting on a function f , via

$$(I^k f)(x) = \int_0^x \int_0^{t_k} \cdots \int_0^{t_2} f(t_1) dt_2 \cdots dt_k.$$

Denote by F_P, F_Q the CDFs of the distributions P, Q . Notice that the population-level KS test statistic can be written as $\rho(P, Q; \mathcal{F}_0) = \|F_P - F_Q\|_\infty$, where $\|\cdot\|_\infty$ is the sup norm. Interestingly, a similar representation holds for the higher-order KS tests.

Proposition 4. *Assuming P, Q have k moments,*

$$\rho(P, Q; \mathcal{F}_k) = \|(I^k)^*(F_P - F_Q)\|_\infty,$$

where $(I^k)^*$ is the adjoint of the bounded linear operator I^k , with respect to the usual L_2 inner product. Further, if P, Q are supported on $[0, \infty)$, or their first k moments

match, then we have the more explicit representation

$$\rho(P, Q; \mathcal{F}_k) = \sup_{x \in \mathbb{R}} \left| \int_x^\infty \int_{t_k}^\infty \cdots \int_{t_2}^\infty (F_P - F_Q)(t_1) dt_2 \cdots dt_k \right|.$$

The representation in Proposition 4 could provide one avenue for power analysis. When P, Q are supported on $[0, \infty)$, or have k matching moments, the representation is particularly simple in form. This form confirms the intuition that detecting higher-order moment differences is hard: as k increases, the k -times integrated CDF difference $F_P - F_Q$ becomes smoother, and hence the differences are less accentuated.

In future work, we also intend to further examine the asymptotic null of the higher-order KS test (the Gaussian process from Theorem 3), and determine to what extent it depends on the underlying distribution P (beyond say, its first k moments). Lastly, some ideas in this paper seem extendable to the multivariate and graph settings, another direction for future work.

Acknowledgments. We thank Alex Smola for several early inspiring discussions. VS and RT were supported by NSF Grant DMS-1554123.

References

- Theodore W. Anderson and Donald A. Darling. Asymptotic theory of certain goodness of fit criteria based on stochastic processes. *Annals of Mathematical Statistics*, 23(2):193–212, 1952.
- Theodore W. Anderson and Donald A. Darling. A test of goodness of fit. *Journal of the American Statistical Association*, 49(268):765–769, 1954.
- Ery Arias-Castro, Bruno Pelletier, and Venkatesh Saligrama. Remember the curse of dimensionality: the case of goodness-of-fit testing in arbitrary dimension. *Journal of Nonparametric Statistics*, 30(2):448–471, 2018.
- K. Babenko. *Theoretical Foundations and Construction of Numerical Algorithms for the Problems of Mathematical Physics*. 1979. In Russian.
- Ludwig Baringhaus and Carsten Franz. On a new multivariate two-sample test. *Journal of Multivariate Analysis*, 88(1):190–206, 2004.
- Bhaswar B. Bhattacharya. *Power of graph-based two-sample tests*. PhD thesis, Stanford University, 2016.
- M. Birman and M. Solomyak. Piecewise-polynomial approximations of functions of the classes W_p^α . *Mathematics of the USSR-Sbornik*, 73(115):331–335, 1967. In Russian.
- Maurice C. Bryson. Heavy-tailed distributions: Properties and tests. *Technometrics*, 16(1):61–68, 1974.
- Fred Calitz. An alternative to the Kolmogorov-Smirnov test for goodness of fit. *Communications in Statistics: Theory and Methods*, 16(12):3519–3534, 1987.
- Richard M. Dudley. *Uniform Central Limit Theorems*. Cambridge University Press, 1999.
- 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.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Scholkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.
- Norbert Henze. A multivariate two-sample test based on the number of nearest neighbor type coincidences. *Annals of Statistics*, 16(2):772–783, 1988.
- Arnold Janssen. Global power functions of goodness of fit tests. *Annals of Statistics*, 28(1):239–253, 2000.
- Andrey Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *Giornale dell’Istituto Italiano degli Attuari*, 4:83–91, 1933.
- Regina Y. Liu and Kesar Singh. A quality index based on data depth and multivariate rank tests. *Journal of the American Statistical Association*, 88(421):252–260, 1993.
- Enno Mammen. Nonparametric regression under qualitative smoothness assumptions. *Annals of Statistics*, 19(2):741–759, 1991.
- David M. Mason and John H. Schuenemeyer. A modified Kolmogorov-Smirnov test sensitive to tail alternatives. *Annals of Statistics*, 11(3):933–946, 1983.
- Alfred Muller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.
- Yurii Nesterov. *Squared Functional Systems and Optimization Problems*, pages 405–440. Springer, 2000.
- Aaditya Ramdas, Sashank Reddi, Barnabas Pczos, Aarti Singh, and Larry Wasserman. On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions. *Twenty-Ninth Conference on Artificial Intelligence*, pages 3571–3577, 2015a.
- Aaditya Ramdas, Sashank Reddi, Barnabas Pczos, Aarti Singh, and Larry Wasserman. Adaptivity and computation-statistics tradeoffs for kernel and distance based high dimensional two sample testing. *arXiv preprint arXiv:1508.00655*, 2015b.
- Michael I. Rosen. Niels Hendrik Abel and equations of the fifth degree. *The American Mathematical Monthly*, 102(6):495–505, 1995.
- Paul R. Rosenbaum. An exact distribution-free test comparing two multivariate distributions based on adjacency. *Journal of the Royal Statistical Society: Series B*, 67(4):515–530, 2005.
- 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.
- Naum Z. Shor. *Nondifferentiable Optimization and Polynomial Problems*. Nonconvex Optimization and Its Applications. Springer, 1998.
- Nikolai Smirnov. Table for estimating the goodness of fit of empirical distributions. *Annals of Mathematical Statistics*, 19(2):279–281, 1948.
- Bharath K. Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Bernhard Scholkopf, and Gert R. G. Lanckriet. On integral probability metrics, ϕ -divergences and binary classification. *arXiv preprint arXiv:0901.2698*, 2009.
- Gabor J. Szekely and Maria L. Rizzo. Testing for equal distributions in high dimension. *InterStat*, 5(16.10):1249–1272, 2004.

Aad van der Vaart and Jon Wellner. *Weak Convergence*. Springer, 1996.

Yu-Xiang Wang, Alexander Smola, and Ryan J. Tibshirani. The falling factorial basis and its statistical applications. *International Conference on Machine Learning*, 31, 2014.