
A Witness Two-Sample Test

Jonas M. Kübler
Max Planck Institute
for Intelligent Systems
Tübingen, Germany

Wittawat Jitkrittum
Google Research

Bernhard Schölkopf
Max Planck Institute
for Intelligent Systems
Tübingen, Germany

Krikamol Muandet
Max Planck Institute
for Intelligent Systems
Tübingen, Germany

Abstract

The Maximum Mean Discrepancy (MMD) has been the state-of-the-art nonparametric test for tackling the two-sample problem. Its statistic is given by the difference in expectations of the witness function, a real-valued function defined as the mean of kernel evaluations on a set of basis points. Typically the kernel is optimized on a training set, and hypothesis testing is performed on a separate test set to avoid overfitting (i.e., control type-I error). That is, the test set is used to simultaneously estimate the expectations and define the basis points, while the training set only serves to select the kernel and is discarded. In this work, we propose to use the training set to also define the weights and the basis points for better data efficiency. We show that 1) the new test is consistent and has a well-controlled type-I error; 2) the optimal witness function is given by a precision-weighted mean in the reproducing kernel Hilbert space associated with the kernel; and 3) the test power of the proposed test is comparable or exceeds that of the MMD and other modern tests, as verified empirically on challenging synthetic and real problems (e.g., Higgs data).

1 INTRODUCTION

We tackle the classic *two-sample problem*: given two samples, do they differ significantly enough that we can conclude they originate from two different distributions? This is a common task in many life sciences such as bioinformatics and cancer diagnosis (Borgwardt

et al., 2006). To decide this, one can perform a *two-sample test*, whose goal is to reject the *null hypothesis* "the probability distributions are the same" in favor of the *alternative hypothesis* "the probability distributions are not the same" based on data (Lehmann and Romano, 2005). To quantitatively assess this, one defines a *test statistic* and estimates its value on the observed samples. If we know (or are able to simulate) the distribution of this test statistic under the null, we can reject the null if the observed value is significantly larger than what we would expect if the null was true. Traditional hypothesis tests have test statistics that are defined a priori. A simple example are *t*- or *z*-tests, which only test whether the empirical means of both samples differ significantly (Lehmann and Romano, 2005). However, such a simple approach is not sufficient to detect differences of distributions with the same mean but, for example, different variance, skewness, or kurtosis.

To detect any differences between two distributions we focus on two categories of tests closely tied to machine learning, but note that various other methods exist (Friedman and Rafsky, 1979; Chen and Friedman, 2017). The former first transforms data into a high-dimensional feature space based on a pre-defined feature map, e.g., kernel function. The test statistics can then be defined in terms of the embeddings of the two distributions in the feature space (Harchaoui et al., 2008b; Gretton et al., 2012a). The second approach instead learns to distinguish the two distributions by training a classifier, e.g., via a deep neural network. Based on the learned model, the test statistics is then computed on an independent set of samples, e.g., through data splitting (Friedman, 2003; Kim et al., 2021; Lopez-Paz and Oquab, 2017; Cheng and Cloninger, 2019).

The popular kernel two-sample test based on the *Maximum Mean Discrepancy* (MMD) in principle does not require data splitting and is completely determined a priori by a positive definite kernel function (Gretton et al., 2012a). However, recent research has shown that optimizing the kernel function on a held-out dataset

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

increases the power of the MMD-based tests (Gretton et al., 2012b; Sutherland et al., 2017; Liu et al., 2020; Kirchler et al., 2020). Thus most modern MMD-based tests are used as two-stage procedures with data splitting, although it is in principle possible to use the entire dataset for kernel selection and testing Fromont et al. (2012, 2013); Kübler et al. (2020).¹

To obtain maximally significant results in the testing phase, we advocate that in a “two-stage” two-sample test, it is more appropriate to learn a test statistic that is as problem-specific as possible. For the MMD tests, this means that we advocate to learn a one-dimensional witness function and not a kernel. To formalize this, we propose a general two-stage witness two-sample test (WiTS test). The introduced WiTS test has the following properties:

- The test statistic is the difference in means of a one-dimensional function called the *witness function* and is thus asymptotically normal under *both* the null and alternative hypotheses. This allows for a simple theoretical treatment (cf. Theorem 1 and Proposition 1).
- Compared to Sutherland et al. (2017) and Liu et al. (2020), the WiTS test has a simpler test power criterion as a training objective and test thresholds can be simulated more efficiently (cf. Section 3 & Eq. (7)).
- The WiTS tests empirically outperform the benchmark tests of Liu et al. (2020) and classification-based tests on challenging synthetic and real problems, e.g., Higgs data (cf. Figure 2).

The rest of the paper is organized as follows. Section 2 reviews MMD based two-sample tests with a focus on the witness function and discusses our motivation. We then present the general WiTS test framework in Section 3, followed by a specific example in Section 4. Next, we discuss related work in detail in Section 5. Finally, Section 6 provides the empirical results comparing the proposed WiTS tests to existing ones on several benchmark datasets. The code to reproduce the experiments is published under <https://github.com/jmkuebler/wits-test>.

2 BACKGROUND AND MOTIVATION

Notation and definitions. Let X, Y be random variables with probability distributions P and Q on $\mathcal{X} \subseteq \mathbb{R}^d$, respectively. In this work, we aim to test the null hypothesis $H_0 : P = Q$ against the alternative $H_1 : P \neq Q$ based on samples $\mathbb{X} = \{x_1, \dots, x_n\}$

¹Schrab et al. (2021) recently proposed an aggregated MMD two-sample test working without data splitting.

and $\mathbb{Y} = \{y_1, \dots, y_m\}$ drawn i.i.d. from P and Q , respectively. Rejecting H_0 although it is true creates a type-I error, whereas not rejecting the null when it is false creates a type-II error. Desirable testing procedures should minimize the type-II error rate, while controlling the type-I error rate at a significance level α (or below). When we consider data splitting, we use $\mathbb{X}_{\text{tr}}, \mathbb{X}_{\text{te}}$ and $\mathbb{Y}_{\text{tr}}, \mathbb{Y}_{\text{te}}$ to denote the disjoint training and test sets with $n = n_{\text{tr}} + n_{\text{te}}$, $m = m_{\text{tr}} + m_{\text{te}}$. We define the shorthands $[n] := \{1, \dots, n\}$, $\mathbb{Z} = \{\mathbb{X}, \mathbb{Y}\}$, $\mathbb{Z}_{\text{tr}} = \{\mathbb{X}_{\text{tr}}, \mathbb{Y}_{\text{tr}}\}$ and $\mathbb{Z}_{\text{te}} = \{\mathbb{X}_{\text{te}}, \mathbb{Y}_{\text{te}}\}$.

Although most of our analysis applies to more general function spaces, we will consider a reproducing kernel Hilbert space (RKHS) \mathcal{H} with positive definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ (Schölkopf and Smola, 2002). By the Riesz representation theorem, we have that $f(x) = \langle f, k(x, \cdot) \rangle$ for all $x \in \mathcal{X}$ and $f \in \mathcal{H}$. We assume that

$$\text{(A1): } \mathbb{E}[k(X, X)] < \infty, \mathbb{E}[k(Y, Y)] < \infty$$

holds. **(A1)** ensures the kernel mean embeddings of P and Q exist, i.e., $\mu_P = \mathbb{E}[k(X, \cdot)]$, $\mu_Q = \mathbb{E}[k(Y, \cdot)]$, and that we can write $\mathbb{E}[f(X)] = \langle f, \mu_P \rangle$ for all $f \in \mathcal{H}$ (Muandet et al., 2017). For a sample \mathbb{X} , we define the empirical mean embedding as $\mu_{\mathbb{X}} = \frac{1}{|\mathbb{X}|} \sum_{x \in \mathbb{X}} k(x, \cdot)$.

MMD and witness function. A popular class of two-sample tests are based on the *Maximum Mean Discrepancy* (MMD) (Gretton et al., 2012a). The MMD of two distributions with respect to the unit ball of \mathcal{H} is defined as (Gretton et al., 2012a, Eq. (1)): $\text{MMD} = \sup_{f \in \mathcal{H}, \|f\| \leq 1} \{\mathbb{E}[f(X)] - \mathbb{E}[f(Y)]\}$. The function that *witnesses* the MMD is $\text{argmax}_{f \in \mathcal{H}, \|f\| \leq 1} \{\mathbb{E}[f(X)] - \mathbb{E}[f(Y)]\} = (\mu_P - \mu_Q) / \|\mu_P - \mu_Q\|$ (Gretton et al., 2012a, Sec. 2.3). We define its unnormalized version as $h_k^{P,Q} = \mu_P - \mu_Q$ and obtain

$$\begin{aligned} \text{MMD}^2 &= \langle \mu_P - \mu_Q, \mu_P - \mu_Q \rangle = \langle \mu_P - \mu_Q, h_k^{P,Q} \rangle \\ &= \mathbb{E} \left[h_k^{P,Q}(X) \right] - \mathbb{E} \left[h_k^{P,Q}(Y) \right]. \end{aligned} \tag{1}$$

With a *characteristic* kernel (Sriperumbudur et al., 2010), $\mu_P = \mu_Q$ if and only if $P = Q$. Hence, the squared MMD (1) can be used to test the hypothesis $H_0 : P = Q$ against $H_1 : P \neq Q$.

MMD-BOOT test statistics. We can estimate the squared MMD (1) by replacing the witness $h_k^{P,Q}$ and the expectations in (1) with their empirical counterparts

$h_k^{\mathbb{Z}} = \mu_{\mathbb{X}} - \mu_{\mathbb{Y}}$ and

$$\begin{aligned} \widehat{\text{MMD}}_{\text{BOOT}}^2(\mathbb{Z}|k) &= \frac{1}{n} \sum_{x \in \mathbb{X}} h_k^{\mathbb{Z}}(x) - \frac{1}{m} \sum_{y \in \mathbb{Y}} h_k^{\mathbb{Z}}(y) \\ &= \left\langle \frac{1}{n} \sum_{x \in \mathbb{X}} k(x, \cdot) - \frac{1}{m} \sum_{y \in \mathbb{Y}} k(y, \cdot), h_k^{\mathbb{Z}}(\cdot) \right\rangle \\ &= \frac{1}{n^2} \sum_{x, x' \in \mathbb{X}} k(x, x') + \frac{1}{m^2} \sum_{y, y' \in \mathbb{Y}} k(y, y') \\ &\quad - \frac{2}{nm} \sum_{x \in \mathbb{X}, y \in \mathbb{Y}} k(x, y). \end{aligned} \quad (2)$$

The latter expression is a sum of V -statistics and up to the biased terms where $x = x'$ or $y = y'$ equals the unbiased U -statistic which is the standard MMD estimate (Gretton et al., 2012a). The witness itself depends on the same data \mathbb{Z} used to evaluate the test statistic (2) and compute the test threshold, the null distribution has to be simulated via permutation of the samples (or bootstrapping) (Gretton et al., 2012a). Thus, we refer to this approach as MMD-BOOT.²

OPT-MMD-BOOT test statistics. A drawback of MMD-BOOT is that the kernel k has to be chosen a priori before observing the data. Kernel choice, however, critically affects the performance of MMD based two-sample tests (Gretton et al., 2012b; Sutherland et al., 2017; Liu et al., 2020; Kübler et al., 2020; Jitkrittum et al., 2016). It is thus common to split the data into $\mathbb{Z} = (\mathbb{Z}_{\text{tr}}, \mathbb{Z}_{\text{te}})$ and optimize the kernel only on the held-out set \mathbb{Z}_{tr} . For the moment, without specifying how the kernel is optimized, we denote the resulting optimized kernel as k_{tr} with a subscript tr to indicate that it depends on the training data. After optimizing the kernel, the standard MMD-BOOT test is conducted on \mathbb{Z}_{te} with the optimized kernel k_{tr} (Sutherland et al., 2017; Liu et al., 2020). Hence, the empirical expectations and witness function in (2) are still dependent on the same data \mathbb{Z}_{te} , and the null distribution still has to be bootstrapped, for the same reason as in the case of MMD-BOOT. We will refer to this approach as OPT-MMD-BOOT with the test statistic

$$\begin{aligned} \widehat{\text{MMD}}_{\text{OPT-BOOT}}^2(\mathbb{Z}_{\text{te}}|k_{\text{tr}}) \\ = \frac{1}{n_{\text{te}}} \sum_{x \in \mathbb{X}_{\text{te}}} h_{k_{\text{tr}}}^{\mathbb{Z}_{\text{te}}}(x) - \frac{1}{m_{\text{te}}} \sum_{y \in \mathbb{Y}_{\text{te}}} h_{k_{\text{tr}}}^{\mathbb{Z}_{\text{te}}}(y). \end{aligned} \quad (3)$$

Our Motivation. This is the starting point of our investigations: Although the kernel is optimized, it is still

²Our naming convention should emphasize that the asymptotic distribution cannot be evaluated in closed-form and hence we necessarily need to simulate it. Note, however, that in practice often permutations are used (Sutherland et al., 2017) and it is not necessary to completely simulate the distribution from scratch.

a multidimensional representation of the data. While this makes the test statistic applicable to other problems (Liu et al., 2020; Kirchler et al., 2020), features that contain little information about the differences of P and Q will mainly add noise to the test statistic. Generally, the noisier the test statistic, the harder it is to obtain significant test results. Motivated by this drawback, we propose to formulate a test statistic that is more specific to the observed difference in \mathbb{Z}_{tr} . Being more specific to the training data (that is all we know about P and Q), comes at the risk of overfitting, which we mitigate via regularization and model selection (cf. Section 3). Specifically for MMD, after the kernel is optimized, *we define the witness directly on the training data* by replacing $h_{k_{\text{tr}}}^{\mathbb{Z}_{\text{te}}}$ with $h_{k_{\text{tr}}}^{\mathbb{Z}_{\text{tr}}} = \frac{1}{n_{\text{tr}}} \sum_{x \in \mathbb{X}_{\text{tr}}} k_{\text{tr}}(x, \cdot) - \frac{1}{m_{\text{tr}}} \sum_{y \in \mathbb{Y}_{\text{tr}}} k_{\text{tr}}(y, \cdot)$. We call this OPT-MMD-WITNESS:

$$\begin{aligned} \widehat{\text{MMD}}_{\text{OPT-WITNESS}}^2(\mathbb{Z}_{\text{te}}|h_{k_{\text{tr}}}^{\mathbb{Z}_{\text{tr}}}) \\ = \frac{1}{n_{\text{te}}} \sum_{x \in \mathbb{X}_{\text{te}}} h_{k_{\text{tr}}}^{\mathbb{Z}_{\text{tr}}}(x) - \frac{1}{m_{\text{te}}} \sum_{y \in \mathbb{Y}_{\text{te}}} h_{k_{\text{tr}}}^{\mathbb{Z}_{\text{tr}}}(y). \end{aligned} \quad (4)$$

This test statistic comes with numerous advantages. Firstly, the expectations (defined via \mathbb{Z}_{te}) are now independent of the witness function (defined via \mathbb{Z}_{tr}). Thus, the test statistic is asymptotically normal. Secondly, as we will see in the following sections, (4) allows us to compute asymptotic test thresholds in closed form and allows for a simpler derivation of a test power criterion than in the case of OPT-MMD-BOOT (Sutherland et al., 2017; Liu et al., 2020). Lastly, our empirical results suggest that OPT-MMD-WITNESS outperforms OPT-MMD-BOOT on datasets considered in Liu et al. (2020).

3 WITNESS TWO-SAMPLE TEST (WiTS TEST)

Similar to (4), the WiTS tests we propose are by design two-stage procedures: In *Stage I*, we learn the witness function h with the training data \mathbb{Z}_{tr} . This ensures that h is independent of the test data \mathbb{Z}_{te} , used in *Stage II* to define a test statistic

$$\hat{\tau}(\mathbb{Z}_{\text{te}}|h) \propto \frac{1}{n_{\text{te}}} \sum_{x \in \mathbb{X}_{\text{te}}} h(x) - \frac{1}{m_{\text{te}}} \sum_{y \in \mathbb{Y}_{\text{te}}} h(y). \quad (5)$$

We reject the null hypothesis $H_0 : P = Q$ if the observed value is larger than a test threshold. We start presenting Stage II and analyze the test's asymptotic power for a given function h . Then, we will use this test power criterion as the objective when optimizing the witness function in Stage I.

3.1 Stage II - Testing with the Witness Function

We start with a basic result on asymptotic normality of empirical means ((Serfling, 1980), Proof in App. A.1).

Theorem 1 (Asymptotic normality of WiTS test). *For a witness function $h : \mathcal{X} \rightarrow \mathbb{R}$, let $\sigma_P^2 := \text{Var}[h(X)]$ and $\sigma_Q^2 := \text{Var}[h(Y)]$ such that $0 < \sigma_P^2, \sigma_Q^2 < \infty$. Let $\{X_i\}_{i \in [n]}$ $\overset{i.i.d.}{\sim} P$, $\{Y_j\}_{j \in [m]}$ $\overset{i.i.d.}{\sim} Q$, and $c := \frac{n}{n+m} \in (0, 1)$ as $n + m \rightarrow \infty$. Denote by $\bar{h}_P := \mathbb{E}[h(X)]$ and $\bar{h}_Q := \mathbb{E}[h(Y)]$. We define the empirical means $\hat{h}_P^n := \frac{1}{n} \sum_{i \in [n]} h(X_i)$, $\hat{h}_Q^m := \frac{1}{m} \sum_{j \in [m]} h(Y_j)$ and denote the sample variance as $\hat{\sigma}_c^2(h) := \hat{\sigma}_P^2/c + \hat{\sigma}_Q^2/(1-c)$. Then*

$$\frac{\sqrt{n+m}}{\hat{\sigma}_c(h)} \left[\left(\hat{h}_P^n - \bar{h}_P \right) - \left(\hat{h}_Q^m - \bar{h}_Q \right) \right] \xrightarrow{d} \mathcal{N}(0, 1).$$

For any fixed h and for sufficiently large sample sizes, we can thus work with the asymptotic distribution of test statistics of the form $\tau(\cdot|h)$ in Eq. (5) to compute test thresholds and derive an asymptotic test-power objective for choosing h based on the training data \mathbb{Z}_{tr} in Stage I. Data splitting ensures that h is independent of \mathbb{Z}_{te} , which is necessary for Theorem 1 to hold. In the following, to make the comparison between different choices of h easier, we consider the standardized test statistic on the test samples \mathbb{Z}_{te}

$$\tau(\mathbb{Z}_{\text{te}}|h) = \frac{\frac{1}{n_{\text{te}}} \sum_{x \in \mathbb{X}_{\text{te}}} h(x) - \frac{1}{m_{\text{te}}} \sum_{y \in \mathbb{Y}_{\text{te}}} h(y)}{\hat{\sigma}_c(h)},$$

where $c = \frac{n_{\text{te}}}{n_{\text{te}} + m_{\text{te}}}$ and $\hat{\sigma}_c(h)$ is the empirical estimate of the pooled variance as in Theorem 1 based on \mathbb{Z}_{te} . To control the type-I error at a significance level α , we need to find a *test threshold* t_α such that $P(\tau(\mathbb{Z}_{\text{te}}|h) > t_\alpha | H_0) \leq \alpha$. By Theorem 1, we can define the threshold to be the $(1 - \alpha)$ quantile of the asymptotic null distribution. Under the null hypothesis we have $\bar{h}_P = \bar{h}_Q$ and obtain $t_\alpha = \Phi^{-1}(1 - \alpha)$ where Φ^{-1} denotes the inverse CDF of the standard normal.

Note that we only consider a "one-sided" test, since we choose h in stage I with the appropriate sign, i.e., such that it has larger expectation under \mathbb{X}_{tr} than under \mathbb{Y}_{tr} . A "two-sided" test ignores this and may lead to a reduction in test power.

We reject the null hypothesis $H_0 : P = Q$ if $\tau(\mathbb{Z}_{\text{te}}|h) > t_\alpha$. As an advantage of the asymptotic normality under the alternative and the closed form of the threshold of our test, we can write the asymptotic type-II error rate in closed form, similar as in Gretton et al. (2012b,

Sec. 3):

$$P(\tau(\mathbb{Z}_{\text{te}}|h) < t_\alpha) \approx \Phi \left(\Phi^{-1}(1 - \alpha) - \sqrt{n_{\text{te}} + m_{\text{te}}} \frac{\bar{h}_P - \bar{h}_Q}{\sigma_c(h)} \right). \quad (6)$$

An important consideration in designing a hypothesis test is test consistency. A hypothesis test is called consistent, if for a fixed alternative hypothesis, its test power converges to one as sample size goes to infinity. With (6), we can characterize for which functions h the statistic τ_h leads to a consistent test.

Proposition 1 (Consistency of WiTS test). *Assume $0 < \sigma_c(h) < \infty$, where $\sigma_c(h)$ is defined in Theorem 1. A WiTS test based on h is consistent against a fixed alternative hypothesis $P \neq Q$ if and only if $\bar{h}_P > \bar{h}_Q$.*

Proposition 1 ensures that, for a given alternative hypothesis, our proposed test will eventually (in the limit of the sample size) reject the null hypothesis H_0 when it is false. Associated with this notion is the *test power*, the probability that the test rejects H_0 when it is false; this quantity is equivalent to $1 - \text{type-II error}$. Defining the *signal-to-noise* ratio $\text{SNR}(h) = \frac{\bar{h}_P - \bar{h}_Q}{\sigma_c(h)}$, it follows from (6) that the asymptotic test power of our test is

$$\beta_h \approx 1 - \Phi \left(\Phi^{-1}(1 - \alpha) - \sqrt{n_{\text{te}} + m_{\text{te}}} \text{SNR}(h) \right). \quad (7)$$

Since Φ increases monotonically, the test power grows monotonically with the *signal-to-noise* ratio (SNR).

3.2 Stage I - Finding an Optimal Witness

We now propose an objective to find an optimal witness function. Based on our test power consideration, we argue that in the first stage one should find a witness by maximizing a, possibly regularized, empirical estimate of the SNR in (7). Let \mathcal{F} be a function class containing candidates for the witness. We propose using the witness \hat{h}_λ defined as

$$\hat{h}_\lambda = \underset{f \in \mathcal{F}}{\text{argmax}} \frac{\bar{f}_{\mathbb{X}_{\text{tr}}} - \bar{f}_{\mathbb{Y}_{\text{tr}}}}{\sigma_{c,\lambda}^{\mathbb{Z}_{\text{tr}}}(f)}, \quad (8)$$

$$\text{with } \bar{f}_{\mathbb{X}_{\text{tr}}} = \frac{1}{n_{\text{tr}}} \sum_{x \in \mathbb{X}_{\text{tr}}} f(x), \quad \bar{f}_{\mathbb{Y}_{\text{tr}}} = \frac{1}{m_{\text{tr}}} \sum_{y \in \mathbb{Y}_{\text{tr}}} f(y),$$

and $\sigma_{c,\lambda}^{\mathbb{Z}_{\text{tr}}}(f) = ((\sigma_c^{\mathbb{Z}_{\text{tr}}}(f)^2 + \lambda \Omega(f))^{\frac{1}{2}}$, where $\sigma_c^{\mathbb{Z}_{\text{tr}}}(f)$ corresponds to $\hat{\sigma}_c(h)$ defined in Theorem 1 and Ω is a regularizer. We remark that the optimal witness is generally not uniquely defined since the SNR is invariant to rescaling the function. Correctly rejecting H_0 when it is false is at the core of hypothesis testing. Our choice of maximizing the SNR in (7) is in line with this principle: it leads to a test that maximizes

the asymptotic test power. By contrast, while other objectives such as classification loss (Kim et al., 2021; Lopez-Paz and Oquab, 2017), softmax loss (Cheng and Cloninger, 2019), or the MMD statistic itself (Gretton et al., 2012a), can be used to learn the witness function, their relationship to the test power may be indirect.

OPT-MMD-Witness. A closely related objective to our SNR in (7) was used in previous work (Sutherland et al., 2017; Liu et al., 2020) to find a good kernel for a OPT-MMD-BOOT test, see (3). For a given kernel k , Liu et al. (2020, Eq.(3)) derive the training objective as $J(P, Q|k) = \text{MMD}^2(P, Q|k) / \sigma_{H_1}(P, Q|k)$ where $\sigma_{H_1}^2(P, Q|k)$ is the asymptotic variance of the MMD estimate under the alternative hypothesis. In Appendix A.5, we examine this quantity in more detail, and show that $J(P, Q|k) = 1/\sqrt{2} \text{SNR}(h_k^{P,Q})$. For a given class of kernels and corresponding (empirical) MMD witnesses, this implies that selecting the optimal witness according to our SNR criterion leads to the same function as first optimizing the kernel with the J criterion and defining the witness afterwards.

Model Selection and Optimization. The choice of function class \mathcal{F} and regularization parameter λ affects the learned witness in (8). We therefore recommend that practitioners use standard tools for model selection such as cross-validation (CV) for finding suitable “hyperparameters” and to validate that the learned witness actually has a high SNR. CV ensures that the witness actually learns the differences between P and Q and does not solely overfit the training data. Model-selection on \mathbb{Z}_{tr} is legit since in Stage II we only use \mathbb{Z}_{te} , which are independent of \mathbb{Z}_{tr} . While this is also possible in classifier two-sample tests (Lopez-Paz and Oquab, 2017), in the standard MMD-BOOT this is not done.

Our objective (7) can be used with a variety of function classes \mathcal{F} . For instance, \mathcal{F} can be defined based on an RKHS, or parameterized by a deep neural network. Note that optimization methods to maximize (8) are generally function class specific, and may require an iterative procedure. However, when \mathcal{F} is an RKHS, we can derive the closed-form solution to (8), as shall be explained in Section 4. Algorithm 1 shows the general procedure for the two-stage WiTS test.

Permutation-based Thresholds. For our theoretical analysis we used the asymptotic threshold. However, the witness is also chosen in a data-dependent manner. Thus, we generally recommend to simulate the threshold via *permutations* in order to ensure type-I error control at finite sample size. In this case, for simplicity and ease of implementation, we compute the test statistic without normalization and simply take

the difference in means. We first compute the value of the witness function on all points in \mathbb{Z}_{te} and store it in an array. Then we compute the simplified test statistic by taking the difference in means of \mathbb{X}_{te} and \mathbb{Y}_{te} (as computed from the array that stores all the witness evaluations). We then iterate over $B \in \mathbb{N}$ permutation runs to estimate the p -value of the computed test statistic. For each run, we permute the array storing the witness evaluations, and then compute the difference in means of the first n_{te} and the last m_{te} entries. If this is larger or equal than the test statistic on the original partition, this contributes $1/B$ to the p -value. After all permutations, if the p -value is smaller than α , we reject (see Alg. 1). This correctly controls type-I errors, as under the null, the initial partition can be thought of as being itself a random permutation of the data. Since for this procedure we only need to compute the witness once on each data point the overall cost is $\mathcal{O}((n_{\text{te}} + m_{\text{te}})B)$. Note that simulating the null for MMD-BOOT instead has cost $\mathcal{O}((n_{\text{te}} + m_{\text{te}})^2 B)$ (Liu et al., 2020, Sec. 5).

4 KFDA-WITNESS

In this section, we consider the function class in (8) to be an RKHS, and show that this choice leads to a closed form solution for the optimal witness. To start, let \mathcal{H} be an RKHS associated with a positive definite kernel k (see Section 2). Additionally to the mean embeddings μ_P, μ_Q , we define the (centered) covariance operator $\Sigma_P = \mathbb{E}[k(X, \cdot) \otimes k(X, \cdot)] - \mu_P \otimes \mu_P$ (analogously for Q) whose existence is ensured by Assumption (A1) (Muandet et al., 2017, Sec. 3). For any function in the RKHS we then have $\mathbb{E}[f(X)] = \langle \mu_P, f \rangle$ and $\text{Var}[f(X)] = \langle f, \Sigma_P f \rangle$, and analogously for Q . We define the pooled covariance operator $\Sigma = \frac{\Sigma_P}{c} + \frac{\Sigma_Q}{1-c}$. Then for all $f \in \mathcal{H}$ with non-zero variance we have

$$\text{SNR}(f) = \frac{\langle \mu_P - \mu_Q, f \rangle}{\langle f, \Sigma f \rangle^{\frac{1}{2}}}, \quad (9)$$

where SNR is defined in (7). This objective corresponds to Kernel Fisher discriminant analysis (KFDA)’s learning objective (Mika et al., 1999). For singular covariance operator the SNR can diverge, and for infinite-dimensional RKHS, the empirical estimation of the covariance operator is ill-posed. In the following, we therefore consider a regularized ($\lambda > 0$) version of (9) and call its solution (*regularized*) *KFDA witness*:

$$h_\lambda = \operatorname{argmax}_{f \in \mathcal{H}} \frac{\langle \mu_P - \mu_Q, f \rangle}{\langle f, (\Sigma + \lambda I) f \rangle^{\frac{1}{2}}}. \quad (10)$$

The solution of (10) is given by the solution to the generalized eigenvalue problem $(\Sigma + \lambda I)h_\lambda = \gamma(\mu_P -$

Algorithm 1 WiTS test with KFDA-WITNESS

<p>1: Input: $\mathbb{X}, \mathbb{Y}, \alpha$, paramGrid, r 2: $\mathbb{X}_{\text{tr}}, \mathbb{X}_{\text{te}}, \mathbb{Y}_{\text{tr}}, \mathbb{Y}_{\text{te}} \leftarrow \text{RANDOMSPLIT}(\mathbb{X}, \mathbb{Y}, r)$ 3: # Optionally perform model selection 4: $k, \lambda \leftarrow \text{GRIDSEARCHCV}(\text{paramGrid}, \mathbb{Z}_{\text{tr}})$ 5: # Stage I - Optimize Witness 6: $h \leftarrow \text{KFDAWITNESS}(\mathbb{Z}_{\text{tr}}, k, \lambda)$ \triangleright App. Alg.2 7: # Stage II - Test 8: return: $\text{WITNESSTEST}(\mathbb{Z}_{\text{te}}, h, \alpha)$</p>	<p>9: function $\text{WITNESSTEST}(\mathbb{Z}_{\text{te}}, h(\cdot), \alpha, B = 200)$ 10: $h_{\mathbb{Z}_{\text{te}}} \leftarrow [h(z) \text{ for } z \text{ in } \mathbb{Z}_{\text{te}}]$ 11: $\tau \leftarrow \text{MEAN}(h_{\mathbb{Z}_{\text{te}}}[n_{\text{te}}]) - \text{MEAN}(h_{\mathbb{Z}_{\text{te}}}[n_{\text{te}} :])$ 12: $p \leftarrow 0$ \triangleright simulate p-value via permutations 13: for i in $[B]$ do 14: $h_{\mathbb{Z}_{\text{te}}} \leftarrow \text{PERMUTE}(h_{\mathbb{Z}_{\text{te}}})$ 15: if $\text{MEAN}(h_{\mathbb{Z}_{\text{te}}}[n_{\text{te}}]) - \text{MEAN}(h_{\mathbb{Z}_{\text{te}}}[n_{\text{te}} :]) \geq \tau$ then 16: $p \leftarrow p + 1/B$ 17: if $p \leq \alpha$ then return: 1 else return: 0</p>
--	--

μ_Q) (Mika, 2003, Sec.3.2), thus

$$h_\lambda = \gamma(\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q), \quad (11)$$

where $\gamma > 0$ is an arbitrary positive constant we fix to 1, unless stated otherwise. We will refer to the test with the witness function h_λ as the KFDA-WITNESS test.

Next, we show how we can estimate the KFDA-witness with the training data.

Estimation of the KFDA Witness. Let $\mathbb{Z}_{\text{tr}} = \{x_1, \dots, x_{n_{\text{tr}}}, y_1, \dots, y_{m_{\text{tr}}}\}$ denote the pooled training sample and K denote the kernel matrix such that $K_{ij} = k(z_i, z_j)$ for $i, j \in [n_{\text{tr}} + m_{\text{tr}}]$. Further, we define $\delta = (\frac{1}{n_{\text{tr}}}, \dots, \frac{1}{n_{\text{tr}}}, -\frac{1}{m_{\text{tr}}}, \dots, -\frac{1}{m_{\text{tr}}})^\top \in \mathbb{R}^{n_{\text{tr}} + m_{\text{tr}}}$. For $l \in \{n_{\text{tr}}, m_{\text{tr}}\}$, we define the idempotent centering matrix $P_l = I_l - l^{-1} \mathbf{1}_l \mathbf{1}_l^\top$, where I_l denotes the identity operator and $\mathbf{1}_l$ the l dimensional vector with all ones. With this we define the $(n_{\text{tr}} + m_{\text{tr}}) \times (n_{\text{tr}} + m_{\text{tr}})$ matrix $N_c = \begin{pmatrix} \frac{1}{c} P_{n_{\text{tr}}} & 0 \\ 0 & \frac{1}{1-c} P_{m_{\text{tr}}} \end{pmatrix}$. Using the representer theorem (Schölkopf et al., 2001), we can empirically estimate the KFDA witness (more detail in App. A.3) as

$$\hat{h}_\lambda(\cdot) = \sum_{i=1}^{n_{\text{tr}} + m_{\text{tr}}} \hat{\alpha}_i k(z_i, \cdot), \quad (12)$$

$$\hat{\alpha} = \left(\frac{KN_c K}{n_{\text{tr}} + m_{\text{tr}}} + \lambda K \right)^{-1} K \delta. \quad (13)$$

$\hat{h}_\lambda(\cdot)$ can be viewed as a precision-weighted (inverse covariance) mean of the embeddings of the basis points \mathbb{Z}_{tr} in the RKHS. Since $\mu_{\mathbb{X}_{\text{tr}}}, \mu_{\mathbb{Y}_{\text{tr}}}$, and $\hat{\Sigma}$ are consistent estimates of μ_P, μ_Q , and Σ , for fixed regularization, we have $\hat{h}_\lambda \rightarrow h_\lambda = (\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q)$ (see Appendix A.4). For the asymptotic witness h_λ we can compute the difference in expectation under P and Q in closed form: $\bar{h}_{\lambda, P} - \bar{h}_{\lambda, Q} = \langle \mu_P - \mu_Q, h_\lambda \rangle = \langle \mu_P - \mu_Q, (\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q) \rangle$. This difference is positive, and hence by Proposition 1 we obtain a consistent WiTS test, if and only if $\mu_P \neq \mu_Q$. We can

ensure this for arbitrary $P \neq Q$ by using a *characteristic* kernel (Sriperumbudur et al., 2010), the same condition as for MMD based tests.

Despite asymptotic consistency, the test power at finite sample size depends on the *splitting ratio* $r \in (0, 1)$, i.e., $n_{\text{tr}} = \lceil rn \rceil$ and $n_{\text{te}} = n - n_{\text{tr}}$ and accordingly for the sample from Q . Based on our experimental results, we observe that, for a fixed kernel k , fixed regularization $\lambda > 0$ and sufficiently large sample size, the splitting ratio $r = 1/2$ appears to give the highest test power in many cases, compared to other values of r . Generally, identifying the optimal splitting ratio remains an open problem. We observe (middle panel of Fig. 1) that if we include model selection in stage I, it is favorable to use more than half of the data for the first stage, i.e., $r > 1/2$. However, since we cannot quantify how much "more" data we should use, we generally recommend using a 50/50 split.

The cost of computing the exact solution \hat{a} in (12) is $\mathcal{O}((n_{\text{tr}} + m_{\text{tr}})^2)$ in space (storing the kernel matrix) and $\mathcal{O}((n_{\text{tr}} + m_{\text{tr}})^3)$ time (matrix inversion). In Appendix C, we adopt recent advances in large-scale kernel machines (Rudi et al., 2017; Meanti et al., 2020) to obtain approximate solutions with lower time and space complexity and thus scale to large datasets. Using the Nyström approximation (Williams and Seeger, 2000) to approximate the solution and approximately solving it with conjugate gradient, we obtain a complexity of $\mathcal{O}((n_{\text{tr}} + m_{\text{tr}})Mt + M^3)$ in time and $\mathcal{O}(M^2)$ in space, where M denotes the number of Nyström centers and t the number of conjugate gradient iterations. For stage II we then only need $(n_{\text{te}} + m_{\text{te}})M$ kernel evaluations to compute the test statistic. This makes our approach scalable to large-scale dataset.³

Connection of OPT-MMD-WITNESS and KFDA-WITNESS. To emphasize the relationship between optimizing the MMD and using KFDA, consider a fixed kernel k and denote by \mathcal{A} the set of bounded positive

³After acceptance of this work, Chatalic et al. (2022) proposed a Nyström approximation of the kernel mean embedding to speed up the MMD estimation.

operators on \mathcal{H}_k . We consider the nonparametric class of kernels $\mathcal{K} = \{k_A | k_A(x, y) = \langle Ak(x, \cdot), Ak(y, \cdot) \rangle, A \in \mathcal{A}\}$. For this class of kernels, we show in App. A.6 that using OPT-MMD-WITNESS leads to the same witness function as using KFDA-WITNESS.

KFDA-BOOT. It turns out that KFDA-like test statistics were considered before (Harchaoui et al., 2008b), but in settings without data splitting. Indeed, for a fixed k and $\lambda > 0$, we can use the whole data, i.e., \mathbb{X}, \mathbb{Y} for learning the witness $(\hat{\Sigma} + \lambda)^{-1}(\mu_{\mathbb{X}} - \mu_{\mathbb{Y}})$ and computing the test statistic (empirical mean difference). The test statistic thus is $\tau_{\text{KFDA-BOOT}} = \langle \mu_{\mathbb{X}} - \mu_{\mathbb{Y}}, (\hat{\Sigma} + \lambda)^{-1}(\mu_{\mathbb{X}} - \mu_{\mathbb{Y}}) \rangle$, and we call its population version $\text{KFDA}^2(P, Q | k, \lambda)$. This, is the test statistic as studied by Harchaoui et al. (2008b). As for MMD-BOOT, the same data is used for estimating the witness and computing the mean difference, hence Theorem 1 does not hold anymore. We thus need to bootstrap the null distribution via permutations of the samples; thus, we refer to it as KFDA-BOOT. KFDA-BOOT has similar drawbacks as MMD-BOOT: 1. simulating the null distribution via permutations has cost $\mathcal{O}((n + m)^3 B)$ for $B \in \mathbb{N}$ draws from the null distribution; and 2. we have to fix k and λ a priori, and their choices strongly affect the test power. Harchaoui et al. (2008b) do not provide guidance for how to choose k and λ .

5 RELATED WORK

Besides the kernel-based tests we discussed so far, Chwialkowski et al. (2015) proposed tests based on *smooth characteristic functions* (SCF), and projected *mean embeddings* (ME) of the distributions where the mean embeddings are projected to J -dimensional Euclidean vectors for $J \in \mathbb{N}$. In fact, the normalized ME statistic in (Chwialkowski et al., 2015, Eq. 13) can be seen as a variant of the KFDA where the function classes is restricted by the J projection directions. Note that for a finite-dimensional RKHS and without regularization, KFDA-BOOT corresponds to the Hotelling’s T^2 statistic (Hotelling, 1931). Jitkrittum et al. (2016) improve this approach by optimizing the features in the first stage. However, they also discard the training data after learning the J projection directions. Kirchler et al. (2020) propose to learn a deep finite-dimensional representation of the data and to use this for a subsequent MMD or KFDA test. However, their training objective does not directly maximize the test power (Kirchler et al., 2020, Sec. 3.1.1). Liu et al. (2020) propose a deep version of OPT-MMD-BOOT. They learn a deep-kernel (MMD-D) of the form $k_{\omega}(x, x') = [(1 - \epsilon)\kappa(\phi_{\omega}(x), \phi_{\omega}(x')) + \epsilon]q(x, x')$, where $\epsilon \in (0, 1)$, κ and q are Gaussian kernels and ϕ_{ω}

is a deep representation optimized via the criterion J , see App. A.5. They also consider a version called MMD-O which is $k_{\omega}(x, x') = \kappa(\phi_{\omega}(x), \phi_{\omega}(x'))$ and conclude that learning a full kernel (they advocate MMD-D) is better than learning a one-dimensional representation.

Most of the aforementioned works focus on developing a practical testing procedure for a specific dataset at hand. However, there also exist more theoretical work on the statistical optimality of different kernel-based approaches. Balasubramanian et al. (2021) show that a *moderated* MMD approach (which is related to KFDA) leads to optimal rates when testing against local alternatives. A similar discussion can be found in the long version of Harchaoui et al. (2008a, Sec.5.1). This resonates our findings, that a witness based on KFDA is more powerful than simply using the MMD witness. Furthermore, Li and Yuan (2019) show how the choice of scaling parameter in Gaussian kernels affects the statistical optimality. However, such theoretically optimal tests oftentimes are unpractical to use. Balasubramanian et al. (2021), for examples requires, the eigendecomposition of the kernel function, which generally is hard to obtain. Furthermore, without data splitting also these works cannot find a good kernel function.

Since our proposed witness function is one-dimensional, it is closely related to classification based two-sample tests (Friedman, 2003; Kim et al., 2021; Lopez-Paz and Oquab, 2017; Cheng and Cloninger, 2019; Cai et al., 2020). Lopez-Paz and Oquab (2017) proposed learning a deep classifier and using its classification accuracy as test statistic. We refer to this as C2ST-S, where S stands for sign. The method has two drawbacks. First, classification loss does not optimize the 0-1 loss, whereas we directly maximize test power (Lopez-Paz and Oquab, 2017, Remark 2). Second, it only uses the sign of the classification function and thus neglects information by weighting all points equally. Cheng and Cloninger (2019) address the second issue by considering the network’s output before thresholding the function into a classifier. They train with a softmax loss, which also does not directly address test power. The connections of these methods to kernel-based tests were also thoroughly discussed by Liu et al. (2020) and, in accordance, we refer to the approach of Cheng and Cloninger (2019) as C2ST-L.

6 EXPERIMENTS

We empirically assess the test power of the proposed WiTS tests in two settings. First, we perform instructive experiments to highlight the differences of the methods summarized in Table 1. Second, we perform benchmark experiments on two challenging datasets

Table 1: Overview of kernel-based two-sample tests. A PRIORI means that the kernel/regularization is chosen independently of the data. The present work proposes the "witness" methods.

METHOD	KERNEL CHOICE	REG. λ	WITNESS OBJ.	WITNESS ESTIM.	TEST DATA	THRESHOLD
KFDA-WITNESS (PROPOSED)	CV	CV	SNR	Z_{TR}	Z_{TE}	ANALYTIC
KFDA-BOOT(HARCHAOU ET AL., 2008B)	A PRIORI	A PRIORI	SNR	Z (IMPLICIT)	Z	BOOTSTRAP
MMD-BOOT(GRETTON ET AL., 2012A)	A PRIORI	-	MMD	Z (IMPLICIT)	Z	BOOTSTRAP
OPT-MMD-WITNESS (PROPOSED)	J WITH Z_{TR}	-	MMD	Z_{TR}	Z_{TE}	ANALYTIC
OPT-MMD-BOOT(SUTHERLAND ET AL., 2017)	J WITH Z_{TR}	-	MMD	Z_{TE} (IMPLICIT)	Z_{TE}	BOOTSTRAP

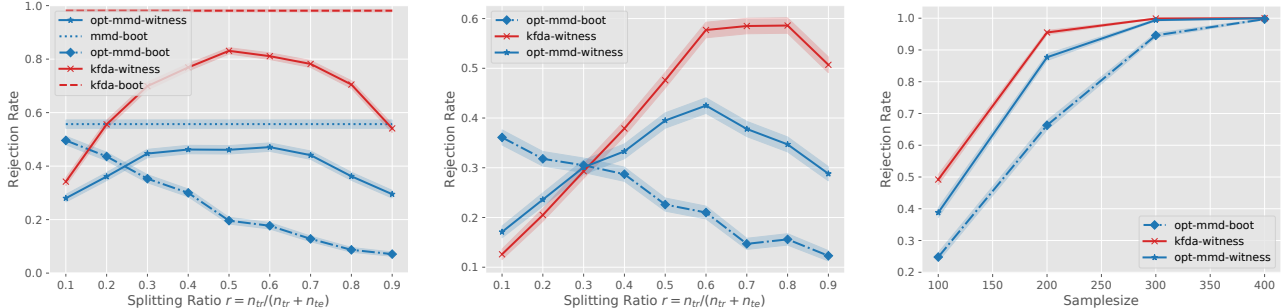


Figure 1: Instructive experiments on "Blobs" dataset. **Left:** Fixed kernel and fixed regularization for sample size $n = m = 100$. **Middle:** For multiple candidate kernels (\mathcal{K}_{10}) kernel optimization becomes more important and the difference of KFDA-WITNESS and OPT-MMD-WITNESS becomes smaller. Further, OPT-MMD-WITNESS already outperforms OPT-MMD-BOOT. **Right:** Same kernels as in the middle figure and $r = 1/2$. All the tests are consistent, i.e., converge to power equal 1.

and compare the performance of the introduced WiTS tests (KFDA-WITNESS and OPT-MMD-WITNESS) to the benchmarks (MMD-D, MMD-O, ME, SCF, C2ST-S, C2ST-L) introduced in Section 5. For the benchmarks, we reuse the implementation provided by Liu et al. (2020) without changing any hyperparameters. Throughout our experiments we set the level $\alpha = 0.05$. App. B contains experiments for correct type-I error control. The shaded regions contain \pm one standard deviation of the estimates.⁴

Instructive experiments. In Figure 1, we consider a **Blobs** dataset (Gretton et al., 2012b) where P and Q are mixtures of nine anisotropic 2-d Gaussians with Q having the covariance matrix rotated by an angle $\theta = \pi/4$, see Figure 5 in the appendix. For the left panel of Fig. 1, we consider a single Gaussian kernel $k_\sigma(x, x') = \exp(-\|x - x'\|^2/\sigma^2)$ with bandwidth $\sigma = 0.2$ and a regularization parameter for the KFDA methods of $\lambda = 10^{-2}$ (in the appendix we show the effect of the regularization in Fig. 4. Note that for $\lambda \rightarrow \infty$, KFDA and MMD methods coincide). We showcase the effect of varying splitting ratios r when the

kernel is fixed a-priori (thus we can apply MMD-BOOT and KFDA-BOOT). With fixed kernel, OPT-MMD-BOOT essentially discards the training data. We estimate the test power (rejection rate) with fixed overall sample size $n = m = 100$. We observe that the witness methods achieve highest power for a 50/50 split, given a fixed kernel and fixed regularization. We also observe that the boot approaches outperform the witness methods in this case.

However, in practice, it is unlikely that we can pick a powerful kernel and regularization *a priori*. Therefore, for the middle panel of Figure 1, we optimize the kernel function over a class of kernels \mathcal{K}_{10} consisting of ten Gaussian kernels with bandwidths on a logarithmic range from 10^{-3} to 10^1 . Additionally, for KFDA-WITNESS we cross-validate over five candidate regularizations on a log range from 10^{-4} to 10^3 . In this case, the witness methods attain the highest power at a splitting ratio $r > 1/2$, and OPT-MMD-WITNESS outperforms OPT-MMD-BOOT for the majority of splitting ratios and also globally. For the right panel, we use the same setting, but fix the splitting ratio at $r = 1/2$ and vary the sample size. As we expect, all tests are consistent and we observe that both WiTS test approaches outperform OPT-MMD-BOOT at a 50/50 split.

⁴Note that in Fig. 2 we used different approaches to estimate the rejection rates, see Appendix B. This explains that at the same rejection rate we can have differently large errors.

Benchmark Experiments. Liu et al. (2020) bench-

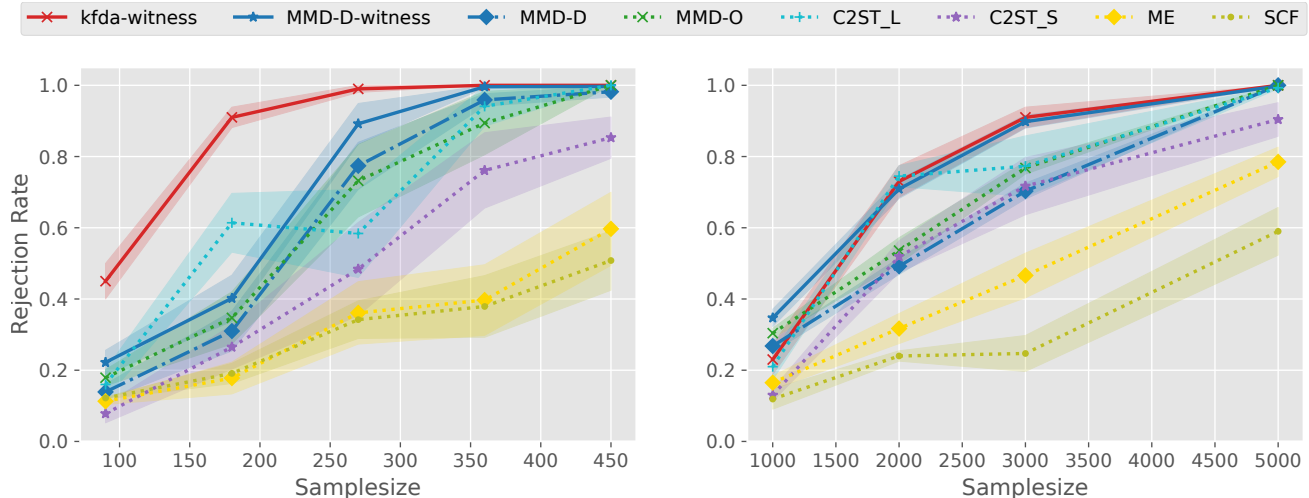


Figure 2: Benchmark experiments adapted from Liu et al. (2020) **Left:** Blobs, **Right:** HIGGS. Computing the MMD witness after kernel optimization and performing a witness test (MMD-D-WITNESS) improves the test power over MMD-D. Directly learning the KFDA-WITNESS also leads to high power.

marked several deep classification two-sample tests (C2ST-L, C2ST-C) against MMD with an optimized deep kernel (MMD-D, MMD-O) and the optimized tests (ME, SCF) of Jitkrittum et al. (2016). We implement OPT-MMD-WITNESS on top of their proposed method MMD-D, which optimizes a deep kernel (Liu et al., 2020, Section 5). Therefore after the kernel optimization, we use the training data to define the MMD witness function (Eq. (4)) and then proceed with WITNESSTEST from Algorithm 1. We also run KFDA-WITNESS with grid search over the same kernels and regularization as for the previous experiments. We run the experiments on two benchmarks. First, an adopted **Blobs** problem, with multiple different covariances (Liu et al., 2020, Figure 1) (see Figure 5 in the appendix), introduced to show the limitations of MMD with translation-invariant kernels. Second, the **Higgs** dataset (Baldi et al., 2014) where "we compare the jet ϕ -momenta distribution ($d = 4$) of the background process, P , which lacks Higgs bosons, to the corresponding distribution Q for the process that produces Higgs bosons" (cited from Liu et al. (2020)). For the Higgs dataset we consider sample sizes larger than a thousand per class. To speed up the computation of the KFDA-WITNESS, we approximate the solution with $M = 500$ Nyström centers, see Appendix C, which underlines the scalability of our approach. For both datasets we observe higher power of the WiTS tests we propose, see Figure 2. We emphasize that we used the implementation of Liu et al. (2020), without changing the deep architecture or any hyperparameters.

7 CONCLUSION

We introduced a principled approach to learn optimal witness functions for two-sample testing. The approach consists of two-stages: First, we learn a witness on a subset of the observations by maximizing a test-power criterion. In the second stage, we simply test whether the witness function attains the same mean on the test samples, and efficiently simulate the null distribution via permutations. We further showed how to adopt recent tests based on optimized Maximum Mean Discrepancy into a witness two-sample test. Liu et al. (2020) advocated optimizing a (deep) kernel in the training stage. Our experiments show, however, that explicitly learning a one-dimensional witness can perform better than learning a high-dimensional representation (a kernel function) in the training stage.

Our results extend beyond kernel methods since we derive a principled objective to train a one-dimensional function optimal for two-sample testing. This objective and the proposed testing procedure can be applied with any function class. The proposed framework thus not only allows domain experts to perform two-sample tests with the models most suitable to the data at hand, but can also easily incorporate model selection techniques developed for classification and regression tasks to optimize for the best parameter settings.

Acknowledgments This work was in part supported by the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039B) and the Machine Learning Cluster of Excellence number 2064/1 – Project 390727645.

References

- K. Balasubramanian, T. Li, and M. Yuan. On the optimality of kernel-embedding based goodness-of-fit tests. *Journal of Machine Learning Research*, 22(1): 1–45, 2021.
- P. Baldi, P. Sadowski, and D. Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5(1):1–9, 2014.
- K. M. Borgwardt, A. Gretton, M. J. Rasch, H.-P. Kriegel, B. Schölkopf, and A. J. Smola. Integrating structured biological data by Kernel Maximum Mean Discrepancy. *Bioinformatics*, 22(14):49–57, 2006.
- H. Cai, B. Goggin, and Q. Jiang. Two-sample test based on classification probability. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 13(1): 5–13, 2020.
- A. Chatalic, N. Schreuder, A. Rudi, and L. Rosasco. Nyström on kernel mean embeddings. *arXiv:2201.13055*, 2022.
- H. Chen and J. H. Friedman. A new graph-based two-sample test for multivariate and object data. *Journal of the American Statistical Association*, 112(517):397–409, 2017.
- X. Cheng and A. Cloninger. Classification logit two-sample testing by neural networks. *arXiv:1909.11298*, 2019.
- K. Chwialkowski, A. Ramdas, D. Sejdinovic, and A. Gretton. Fast two-sample testing with analytic representations of probability measures. In *NeurIPS*, 2015.
- J. H. Friedman. On multivariate goodness of fit and two sample testing. *Stanford Linear Accelerator Center-PUB-10325*, 2003.
- J. H. Friedman and L. C. Rafsky. Multivariate Generalizations of the Wald-Wolfowitz and Smirnov Two-Sample Tests. *The Annals of Statistics*, 7(4):697 – 717, 1979.
- M. Fromont, B. Laurent, M. Lerasle, and P. Reynaud-Bouret. Kernels based tests with non-asymptotic bootstrap approaches for two-sample problems. In *COLT*, 2012.
- M. Fromont, B. Laurent, and P. Reynaud-Bouret. The two-sample problem for Poisson processes: Adaptive tests with a nonasymptotic wild bootstrap approach. *The Annals of Statistics*, 41(3):1431 – 1461, 2013.
- K. Fukumizu, F. R. Bach, and A. Gretton. Statistical convergence of kernel CCA. In *NeurIPS*, 2005.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13: 723–773, 2012a.
- A. Gretton, D. Sejdinovic, H. Strathmann, S. Balakrishnan, M. Pontil, K. Fukumizu, and B. K. Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In *NeurIPS*, 2012b.
- Z. Harchaoui, F. Bach, and E. Moulines. Testing for homogeneity with kernel fisher discriminant analysis. *arXiv:0804.1026*, 2008a.
- Z. Harchaoui, F. R. Bach, and E. Moulines. Testing for homogeneity with kernel Fisher discriminant analysis. In *NeurIPS*, 2008b.
- H. Hotelling. The generalization of student’s ratio. *The Annals of Mathematical Statistics*, 2(3):360–378, 1931.
- W. Jitkrittum, Z. Szabó, K. P. Chwialkowski, and A. Gretton. Interpretable distribution features with maximum testing power. In *NeurIPS*, 2016.
- I. Kim, A. Ramdas, A. Singh, and L. Wasserman. Classification accuracy as a proxy for two-sample testing. *The Annals of Statistics*, 49(1):411 – 434, 2021.
- M. Kirchler, S. Khorasani, M. Kloft, and C. Lippert. Two-sample testing using deep learning. In *AISTATS*, 2020.
- J. M. Kübler, W. Jitkrittum, B. Schölkopf, and K. Muandet. Learning kernel tests without data splitting. In *NeurIPS*, 2020.
- E. L. Lehmann and J. P. Romano. *Testing statistical hypotheses*. Springer Texts in Statistics. Springer, third edition, 2005.
- T. Li and M. Yuan. On the optimality of gaussian kernel based nonparametric tests against smooth alternatives. *arXiv:1909.03302*, 2019.
- F. Liu, W. Xu, J. Lu, G. Zhang, A. Gretton, and D. J. Sutherland. Learning deep kernels for nonparametric two-sample tests. In *ICML*, 2020.
- D. Lopez-Paz and M. Oquab. Revisiting classifier two-sample tests. In *ICLR*, 2017.
- G. Meanti, L. Carratino, L. Rosasco, and A. Rudi. Kernel methods through the roof: Handling billions of points efficiently. In *NeurIPS*, 2020.
- S. Mika. *Kernel Fisher Discriminants*. Doctoral thesis, Technische Universität Berlin, Berlin, 2003.
- S. Mika, G. Rätsch, J. Weston, B. Schölkopf, and K. R. Müller. Fisher discriminant analysis with kernels. In *Neural Networks for Signal Processing IX*, pages 41–48, 1999.
- K. Muandet, K. Fukumizu, B. Sriperumbudur, and B. Schölkopf. *Kernel Mean Embedding of Distributions: A Review and Beyond*, volume 10 of *Foundations and Trends in Machine Learning*. 2017.

- A. Rudi, L. Carratino, and L. Rosasco. Falcon: An optimal large scale kernel method. In *NeurIPS*, 2017.
- B. Schölkopf and A. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2002.
- B. Schölkopf, R. Herbrich, and A. J. Smola. A generalized representer theorem. In *COLT*, 2001.
- A. Schrab, I. Kim, M. Albert, B. Laurent, B. Guedj, and A. Gretton. Mmd aggregated two-sample test. *arXiv:2110.15073*, 2021.
- R. J. Serfling. *Approximation theorems of mathematical statistics*. John Wiley & Sons, 1980.
- B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet. Hilbert Space Embeddings and Metrics on Probability Measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.
- D. J. Sutherland, H.-Y. Tung, H. Strathmann, S. De, A. Ramdas, A. Smola, and A. Gretton. Generative models and model criticism via optimized maximum mean discrepancy. In *ICLR*, 2017.
- C. K. I. Williams and M. W. Seeger. Using the Nyström method to speed up kernel machines. In *NeurIPS*, 2000.

A PROOFS

A.1 Proof of Theorem 1

Proof. Theorem 1 follows by the application of the CLT; see, e.g., Theorem A, Chapter 1.9.1 in Serfling (1980). The CLT implies $\sqrt{n+m}(\hat{h}_P^n - \bar{h}_P) = \sqrt{n/c}(\hat{h}_P^n - \bar{h}_P) \xrightarrow{d} \mathcal{N}(0, \sigma_P^2/c)$, analogously for Q and the variances add up. Since $\hat{\sigma}_c^2(h) \xrightarrow{p} \sigma_c := \sigma_P^2/c + \sigma_Q^2/(1-c)$, the result follows from Slutsky's theorem. \square

A.2 Proof of Proposition 1

Proof. Since we assume $\sigma_c(h) > 0$, it follows that

$$\lim_{n_{te}+m_{te} \rightarrow \infty} \Phi \left(\Phi^{-1}(1-\alpha) - \sqrt{n_{te}+m_{te}} \frac{\bar{h}_P - \bar{h}_Q}{\sigma_c(h)} \right) = 0, \quad (14)$$

i.e., the asymptotic rate of type-II errors goes to zero, if and only if $\bar{h}_P > \bar{h}_Q$. \square

A.3 Derivation of Equation (12)

We use the following definitions: Let $Z = \{x_1, \dots, x_{n_{tr}}, y_1, \dots, y_{m_{tr}}\}$ denote the pooled training sample and K denote the kernel matrix such that $K_{ij} = k(z_i, z_j)$ for $i, j \in [n_{tr} + m_{tr}]$. Let us define $G \in \mathcal{H}^{n_{tr}+m_{tr}}$ such that $G_i = k(z_i, \cdot)$. And we write $K = G^\top G$. Further we define $v_1 = (\frac{1}{n_{tr}}, \dots, \frac{1}{n_{tr}}, 0, \dots, 0)^\top \in \mathbb{R}^{n_{tr}+m_{tr}}$, $v_2 = (0, \dots, 0, \frac{1}{m_{tr}}, \dots, \frac{1}{m_{tr}})^\top \in \mathbb{R}^{n_{tr}+m_{tr}}$, and $\delta = v_1 - v_2$. For $l = n_{tr}, m_{tr}$ we define the idempotent centering operator $P_l = I_l - l^{-1} \mathbf{1}_l \mathbf{1}_l^\top$, where I denotes the identity operator and $\mathbf{1}_l$ the l dimensional vector with all ones. With this we define the $(n_{tr} + m_{tr}) \times (n_{tr} + m_{tr})$ matrix $N_c = \begin{pmatrix} \frac{1}{c} P_{n_{tr}} & 0 \\ 0 & \frac{1}{1-c} P_{m_{tr}} \end{pmatrix}$. With the preceding definitions, we obtain $\hat{\mu}_P - \hat{\mu}_Q = G\delta$, $\hat{\Sigma} = \frac{1}{n_{tr}+m_{tr}} G N_c G^\top$.

Starting from (10) we estimate the KFDDA witness based on the empirical estimates of μ_P, μ_Q, Σ , i.e.,

$$\hat{h}_\lambda = \operatorname{argmax}_{f \in \mathcal{H}} \frac{\langle \mu_{\mathbb{X}_{tr}} - \mu_{\mathbb{Y}_{tr}}, f \rangle}{\langle f, (\hat{\Sigma} + \lambda I) f \rangle^{\frac{1}{2}}}. \quad (15)$$

We first show a *representer Theorem* for KFDDA (Mika, 2003, Sec. 3.4.3). Therefore, we decompose possible candidate functions $f = f_1 + f_2 \in \mathcal{H}$ into a part f_1 that lies in the span of the training data $\mathcal{S}_{tr} = \operatorname{span}(\{k(z_i, \cdot) | i \in [n_{tr} + m_{tr}]\})$ and f_2 which lies in the span's orthogonal complement. Thus, by definition, we have $\langle f_2, k(z_i, \cdot) \rangle = 0$ for all $i \in [n_{tr} + m_{tr}]$. Since $\mu_{\mathbb{X}_{tr}}$ and $\mu_{\mathbb{Y}_{tr}}$ are within \mathcal{S}_{tr} , we have $\langle \mu_{\mathbb{X}_{tr}} - \mu_{\mathbb{Y}_{tr}}, f \rangle = \langle \mu_{\mathbb{X}_{tr}} - \mu_{\mathbb{Y}_{tr}}, f_1 \rangle$. Similarly, since $\hat{\Sigma}$ is only defined via the training samples in Z , $\hat{\Sigma}$ maps functions from \mathcal{S}_{tr} to \mathcal{S}_{tr} and we have $\hat{\Sigma} f_2 = 0$. Thus for the denominator of (15) we get

$$\langle f, (\hat{\Sigma} + \lambda I) f \rangle = \langle f_1, (\hat{\Sigma} + \lambda I) f_1 \rangle + \lambda \|f_2\|^2 \geq \langle f_1, (\hat{\Sigma} + \lambda I) f_1 \rangle. \quad (16)$$

We have shown that the nominator of (15) stays constant, if we add a function f_2 that is not in \mathcal{S}_{tr} and the denominator can only grow. This implies that the maximum in (15) is attained for a function in \mathcal{S}_{tr} and we can expand it as $\hat{h}_\lambda(\cdot) = \sum_{i=1}^{n_{tr}+m_{tr}} \hat{\alpha}_i k(z_i, \cdot)$. Hence the solution is

$$\hat{\alpha} = \operatorname{argmax}_{\alpha \in \mathbb{R}^{n_{tr}+m_{tr}}} \frac{\langle \mu_{\mathbb{X}_{tr}} - \mu_{\mathbb{Y}_{tr}}, \sum_{i=1}^{n_{tr}+m_{tr}} \alpha_i k(z_i, \cdot) \rangle}{\langle \sum_{i=1}^{n_{tr}+m_{tr}} \alpha_i k(z_i, \cdot), (\hat{\Sigma} + \lambda I) \sum_{i=1}^{n_{tr}+m_{tr}} \alpha_i k(z_i, \cdot) \rangle^{\frac{1}{2}}} \quad (17)$$

$$= \operatorname{argmax}_{\alpha \in \mathbb{R}^{n_{tr}+m_{tr}}} \frac{\delta^\top K \alpha}{\left(\alpha^\top \left(\frac{K N_c K}{n_{tr}+m_{tr}} + \lambda K \right) \alpha \right)^{\frac{1}{2}}}. \quad (18)$$

The solution to this is (Mika, 2003, Sec. 3.2)⁵

$$\left(\frac{K N_c K}{n_{tr} + m_{tr}} + \lambda K \right) \hat{\alpha} = K \delta \quad \iff \quad \hat{\alpha} = \left(\frac{K N_c K}{n_{tr} + m_{tr}} + \lambda K \right)^{-1} K \delta. \quad (19)$$

⁵For a sanity check, simply compute the gradient of (17) and set it to zero.

A.4 Convergence of \hat{h}_λ

We will show that $\hat{h}_\lambda \rightarrow h_\lambda = (\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q)$ in probability.

Proof. First, we observe that

$$\begin{aligned} \hat{h}_\lambda - h_\lambda &= (\hat{\Sigma} + \lambda I)^{-1}(\mu_{\mathbb{X}_{\text{tr}}} - \mu_{\mathbb{Y}_{\text{tr}}}) - (\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q) \\ &= (\hat{\Sigma} + \lambda I)^{-1}(\mu_{\mathbb{X}_{\text{tr}}} - \mu_{\mathbb{Y}_{\text{tr}}}) - (\hat{\Sigma} + \lambda I)^{-1}(\mu_P - \mu_Q) \\ &\quad + (\hat{\Sigma} + \lambda I)^{-1}(\mu_P - \mu_Q) - (\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q) \\ &= (\hat{\Sigma} + \lambda I)^{-1}[(\mu_{\mathbb{X}_{\text{tr}}} - \mu_{\mathbb{Y}_{\text{tr}}}) - (\mu_P - \mu_Q)] + \left[(\hat{\Sigma} + \lambda I)^{-1} - (\Sigma + \lambda I)^{-1} \right] (\mu_P - \mu_Q). \end{aligned}$$

Thus it follows that

$$\begin{aligned} \|\hat{h}_\lambda - h_\lambda\|_{\mathcal{H}} &\leq \|(\hat{\Sigma} + \lambda I)^{-1}[(\mu_{\mathbb{X}_{\text{tr}}} - \mu_{\mathbb{Y}_{\text{tr}}}) - (\mu_P - \mu_Q)]\|_{\mathcal{H}} \\ &\quad + \|[(\hat{\Sigma} + \lambda I)^{-1} - (\Sigma + \lambda I)^{-1}](\mu_P - \mu_Q)\|_{\mathcal{H}} \\ &= (A) + (B). \end{aligned}$$

Probabilistic bound on (A). By the triangle inequality,

$$\begin{aligned} \|(\hat{\Sigma} + \lambda I)^{-1}[(\mu_{\mathbb{X}_{\text{tr}}} - \mu_{\mathbb{Y}_{\text{tr}}}) - (\mu_P - \mu_Q)]\|_{\mathcal{H}} &\leq \|(\hat{\Sigma} + \lambda I)^{-1}\| \|(\mu_{\mathbb{X}_{\text{tr}}} - \mu_{\mathbb{Y}_{\text{tr}}}) - (\mu_P - \mu_Q)\|_{\mathcal{H}} \\ &\leq \|(\hat{\Sigma} + \lambda I)^{-1}\| (\|\mu_{\mathbb{X}_{\text{tr}}} - \mu_P\|_{\mathcal{H}} + \|\mu_Q - \mu_{\mathbb{Y}_{\text{tr}}}\|_{\mathcal{H}}). \end{aligned}$$

By the spectral theorem, $\|(\hat{\Sigma} + \lambda I)^{-1}\| = \sup_{\hat{l}_k \in (\hat{l}_k)_{k=1}^{\infty}} \frac{1}{\hat{l}_k + \lambda} \leq 1/\lambda$ where $(\hat{l}_k)_{k=0}^{\infty}$ are the eigenvalues of $\hat{\Sigma}$ and by definition non-negative. Then, the \sqrt{n} -convergence of (A) follows from the \sqrt{n} -convergence of the kernel mean embeddings $\|\mu_{\mathbb{X}_{\text{tr}}} - \mu_P\|_{\mathcal{H}} = \mathcal{O}_p(n_{\text{tr}}^{-1/2})$ and $\|\mu_Q - \mu_{\mathbb{Y}_{\text{tr}}}\|_{\mathcal{H}} = \mathcal{O}_p(m_{\text{tr}}^{-1/2})$; see, e.g., Muandet et al. (2017, Theorem 3.4). That is, $(A) = \mathcal{O}_p(\min(n_{\text{tr}}, m_{\text{tr}})^{-1/2})$.

Probabilistic bound on (B). Using the identity $C^{-1} - D^{-1} = C^{-1}(D - C)D^{-1}$, we can rewrite (B) as

$$\begin{aligned} \|[(\hat{\Sigma} + \lambda I)^{-1} - (\Sigma + \lambda I)^{-1}](\mu_P - \mu_Q)\|_{\mathcal{H}} &= \|(\hat{\Sigma} + \lambda I)^{-1}(\hat{\Sigma} - \Sigma)(\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q)\|_{\mathcal{H}} \\ &\leq \|(\hat{\Sigma} + \lambda I)^{-1}\| \|\hat{\Sigma} - \Sigma\| \|(\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q)\|_{\mathcal{H}} \\ &\leq \|(\hat{\Sigma} + \lambda I)^{-1}\| \|\hat{\Sigma} - \Sigma\|_{\text{HS}} \|(\Sigma + \lambda I)^{-1}(\mu_P - \mu_Q)\|_{\mathcal{H}}, \end{aligned}$$

where we used that the operator norm is upper bounded by the Hilbert-Schmidt norm. Let $n := n_{\text{tr}} + m_{\text{tr}}$. Then, since $\|(\hat{\Sigma} + \lambda I)^{-1}\| \leq 1/\lambda$, the \sqrt{n} -convergence of (B) follows from the \sqrt{n} -convergence of the covariance operator, i.e., $\|\hat{\Sigma} - \Sigma\|_{\text{HS}} = \mathcal{O}_p(n^{-1/2})$ (Fukumizu et al., 2005, Lemma 4). That is, $(B) = \mathcal{O}_p((n_{\text{tr}} + m_{\text{tr}})^{-1/2})$.

Combining the rates of (A) and (B) yields the overall rate of convergence: $\|\hat{h}_\lambda - h_\lambda\|_{\mathcal{H}} = \mathcal{O}_p(\min(n_{\text{tr}}, m_{\text{tr}})^{-1/2})$. \square

A.5 Witness objective vs. kernel optimization objective in MMD tests

In MMD-based two sample tests, the most common estimate of the MMD is the U-statistic estimate, defined as (Gretton et al., 2012a)

$$\widehat{\text{MMD}}_u^2 = \frac{1}{n(n+1)} \sum_{i \neq j} H_{ij}, \quad (20)$$

with $H_{ij} = \langle k(x_i, \cdot) - k(y_i, \cdot), k(x_j, \cdot) - k(y_j, \cdot) \rangle$. The objective function used in Sutherland et al. (2017); Liu et al. (2020) bases on the asymptotic variance of the estimator under the alternative hypothesis. If the population value of MMD^2 is positive, then the distribution of the estimate is asymptotically normal (Serfling, 1980, Section 5.5.1), $\sqrt{n} \left(\widehat{\text{MMD}}_u^2 - \text{MMD}^2 \right) \xrightarrow{d} \mathcal{N}(0, \sigma_{H_1}^2)$, with $\sigma_{H_1}^2 = 4(\mathbb{E}[H_{12}H_{13}] - \mathbb{E}[H_{12}]^2)$ (Liu et al., 2020). This can be used

to derive an asymptotic test power criterion, which is given as the signal-to-noise ratio $J = \frac{\text{MMD}^2}{\sigma_{H_1}}$ (Sutherland et al., 2017, Sec. 2.1).

We show, that the power criterion $J = \frac{\text{MMD}^2}{\sigma_{H_1}}$ corresponds to the SNR criterion we derived in (8). It is an easy exercise to show that

$$\begin{aligned} \sigma_{H_1}^2 &= 4 \left(\mathbb{E}_{X \sim P} \left[\langle \mu_P - \mu_Q, k(X, \cdot) \rangle^2 \right] + \mathbb{E}_{Y \sim Q} \left[\langle \mu_P - \mu_Q, k(Y, \cdot) \rangle^2 \right] \right. \\ &\quad \left. - \langle \mu_P - \mu_Q, \mu_P \rangle^2 - \langle \mu_P - \mu_Q, \mu_Q \rangle^2 \right). \end{aligned}$$

Recalling the definition of the covariance operator $\Sigma_P = \mathbb{E} [k(X, \cdot) \otimes k(X, \cdot)] - \mu_P \otimes \mu_P$, we obtain

$$\begin{aligned} \sigma_{H_1}^2 &= 4 \langle \mu_P - \mu_Q, (\Sigma_P + \Sigma_Q)(\mu_P - \mu_Q) \rangle = 2 \langle \mu_P - \mu_Q, (2\Sigma_P + 2\Sigma_Q)(\mu_P - \mu_Q) \rangle \\ &= 2 \langle \mu_P - \mu_Q, \Sigma(\mu_P - \mu_Q) \rangle, \end{aligned}$$

where we used $\Sigma = \Sigma_P/c + \Sigma_Q/(1-c)$ and $c = 1/2$ for balanced samples.

Using $h_k^{P,Q} = \mu_P - \mu_Q$, we have

$$J(P, Q|k) = \frac{\text{MMD}^2}{\sigma_{H_1}} = \frac{\langle \mu_P - \mu_Q, \mu_P - \mu_Q \rangle}{\sqrt{2} \langle \mu_P - \mu_Q, \Sigma(\mu_P - \mu_Q) \rangle^{\frac{1}{2}}} = \frac{\langle \mu_P - \mu_Q, h_k^{P,Q} \rangle}{\sqrt{2} \langle h_k^{P,Q}, \Sigma h_k^{P,Q} \rangle^{\frac{1}{2}}} \quad (21)$$

$$= \frac{1}{\sqrt{2}} \text{SNR}(h_k^{P,Q}). \quad (22)$$

A.6 MMD of nonparametrically optimized kernel corresponds to KFDA

Consider a fixed kernel k and denote by \mathcal{A} the set of bounded positive operators on \mathcal{H}_k . For the nonparametric class of kernels $\mathcal{K} = \{k_A | k_A(x, y) = \langle Ak(x, \cdot), Ak(y, \cdot) \rangle, A \in \mathcal{A}\}$ using OPT-MMD-WITNESS leads to exactly the same witness function as using KFDA-WITNESS.

Proof. Writing inner products in the original RKHS with kernel k for kernel k_A we have the regularized J criterion

$$J_A^\lambda = \frac{\langle A(\mu_P - \mu_Q), A(\mu_P - \mu_Q) \rangle}{\langle A(\mu_P - \mu_Q), A(\Sigma + \lambda I)AA(\mu_P - \mu_Q) \rangle^{\frac{1}{2}}}.$$

We define $\delta_A := A^2(\mu_P - \mu_Q)$ and obtain

$$J_A^\lambda = \frac{\langle \mu_P - \mu_Q, \delta_A \rangle}{\langle \delta_A, (\Sigma + \lambda I)\delta_A \rangle^{\frac{1}{2}}}, \quad (23)$$

which looks almost like (10). The solution to (10) is (11) which implies that $\tilde{A}_\lambda = (\Sigma + \lambda I)^{-\frac{1}{2}}$ defines the optimal kernel

$$\begin{aligned} \tilde{k}_\lambda(x, x') &:= \langle (\Sigma + \lambda I)^{-\frac{1}{2}} k(x, \cdot), (\Sigma + \lambda I)^{-\frac{1}{2}} k(x', \cdot) \rangle_{\mathcal{H}} \\ &= \langle k(x, \cdot), (\Sigma + \lambda I)^{-1} k(x', \cdot) \rangle_{\mathcal{H}}. \end{aligned}$$

Based on the empirical estimates the MMD witness of the optimized kernel would be (expressed in terms of the original kernel k)

$$h_{\tilde{k}_\lambda}^{\text{Ztr}} = (\hat{\Sigma} + \lambda I)^{-1} (\mu_{\text{Xtr}} - \mu_{\text{Ytr}}) = \hat{h}_\lambda, \quad (24)$$

i.e., the witness of OPT-MMD-WITNESS coincides with the KFDA-WITNESS in the original RKHS. \square

B FURTHER EXPERIMENTS AND DETAILS

This section provides supplementary information on our experiments. We provide code upon personal request.

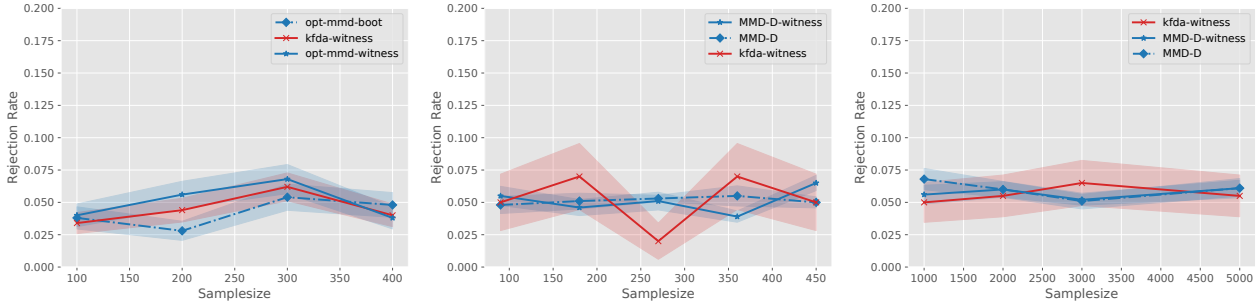


Figure 3: Rejection Rates for true null hypothesis (Type I error) at $\alpha = 0.05$. **Left:** Standard Blobs dataset (500 iterations). **Middle:** Blobs dataset of Liu et al. (2020), KFDA-WITNESS is only average over 100 trials the others over 10×100 , therefore KFDA-WITNESS has higher variance. **Right:** Higgs dataset

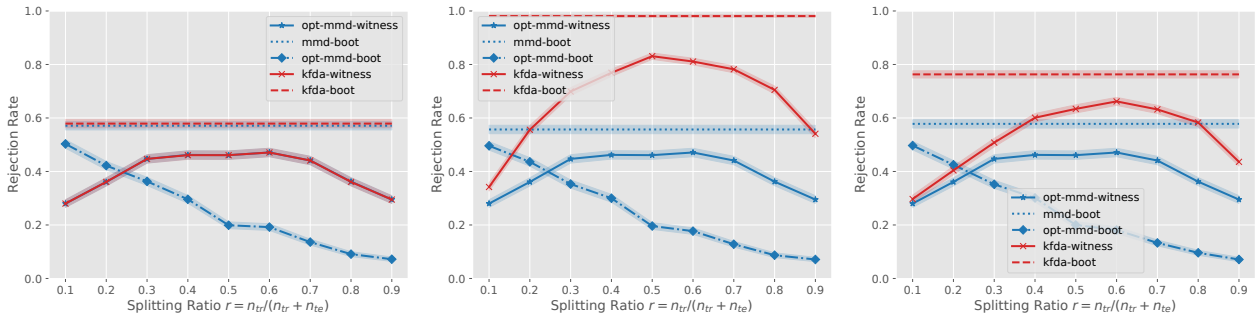


Figure 4: **Effect of regularization on KFDA.** We consider the same setting as in the left panel of Fig. 1 (fixed kernel and fixed regularization and $n = m = 100$) but for different regularization. **Left** ($\lambda = 10^3$): For large regularization KFDA converges to MMD. **Middle** ($\lambda = 10^{-2}$): For a good regularization the KFDA approaches clearly outperform the corresponding MMD approaches. **Right** ($\lambda = 10^{-4}$): If the regularization is too small for a given sample size (here $n = 100$), then KFDA overfits in the training phase, which leads to a reduction in test power.

Datasets. We used two different versions of the Blobs dataset. We show random draws for both cases in Figure 5. For the benchmark experiments we also used the Higgs dataset (Baldi et al., 2014), which is part of the *UCI Machine Learning Repository* (<https://archive.ics.uci.edu/ml/datasets/HIGGS>). We used a version that is ready for Python usage provided by Liu et al. (2020) (https://drive.google.com/open?id=1sHIIFCoHbauk6Mkb6e8a_tp1qnvuUOCc). To ensure the comparability we follow the implementation of Liu et al. (2020) and draw samples from the Higgs dataset *without replacement*.

Effect of regularization of KFDA-WITNESS. In the left panel of Figure 1, we chose a fixed regularization $\lambda = 10^{-2}$ for the KFDA methods. In Figure 4, we show the effect of choosing a bad regularization. If the regularization is too large (left), then KFDA coincides with MMD. On the other hand, if the regularization is too small (right), then the effect of inaccurately estimating the covariance operator might as well lead to a reduced test power. For good performance it is thus important to chose a suitable regularization. This can be automated by including a model selection procedure, such as cross-validation, in the training stage.

Estimation of Rejection Rates. For the instructive experiments in Figure 1 we estimate the rejection rates by repeating the whole two-stage procedure 1000 times. For the benchmark experiments we use 100 iterations of the two-stage procedure for KFDA-WITNESS. For all the other methods in the benchmark experiments, we follow the implementation of Liu et al. (2020) and estimate the rejection rates by running the first stage ten times and estimating the rejection rate over 100 independent test sets for each run of the first stage. The reason for this is, that the first stage is quite slow (training a neural network).

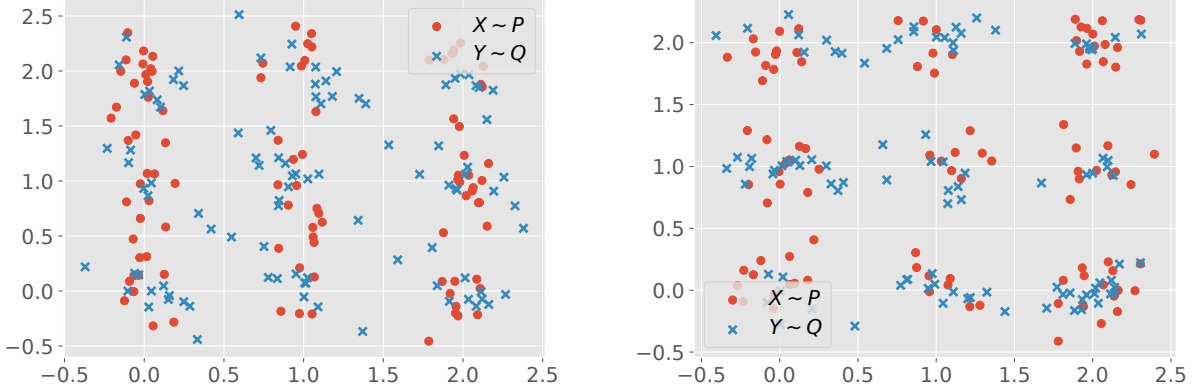


Figure 5: **Left:** Draws from Blobs dataset for the instructive experiments. The distributions are mixtures of nine Gaussians, with anisotropic covariance (but the same covariance matrix across blobs). The covariance matrix of Q is rotated by $\theta = \pi/4$ relative to the covariance matrix of P . To simulate the null hypothesis we use $\theta = 0$, which corresponds to drawing both samples from P . **Right:** Blobs dataset used for Figure 2 as suggested by Liu et al. (2020, Figure 1). In this case, P has isotropic Gaussian, the blobs in Q are anisotropic and have different covariance matrices. To simulate the null hypothesis, we draw both samples from P .

Algorithm 2 Pseudocode for the FdaFalkon algorithm. Adopted for KFDA from (Meanti et al., 2020)

<pre> 1: function FDAFALKON($Z, \mathbf{y}, k, \lambda, m, t$) 2: $Z_m, \mathbf{y}_m \leftarrow \text{RANDOMSUBSAMPLE}((Z, \mathbf{y}), m)$ 3: $T, A \leftarrow \text{PRECONDITIONER}(Z_m, \mathbf{y}_m, \lambda)$ 4: function LINOP(β) 5: $\mathbf{v} \leftarrow A^{-1}\beta$ 6: $\mathbf{c} \leftarrow k(Z_m, Z)NN^\top k(Z, Z_m)T^{-1}\mathbf{v}$ 7: return $A^{-\top}(T^{-\top}\mathbf{c} + \lambda n\mathbf{v})$ 8: $R \leftarrow A^{-\top}T^{-\top}k(Z_m, Z)\mathbf{y}$ 9: $\beta \leftarrow \text{CONJUGATEGRADIENT}(\text{LINOP}, R, t)$ 10: return $T^{-1}A^{-1}\beta, Z_m$ </pre>	<pre> 13: function PRECONDITIONER($Z_m, \mathbf{y}_m, \lambda$) 14: $K_{mm} \leftarrow k(Z_m, Z_m)$ 15: $T \leftarrow \text{chol}(K_{mm})$ 16: $K_{mm} \leftarrow \frac{1}{m}TN_mN_mT^\top + \lambda\mathbf{I}$ 17: $A \leftarrow \text{chol}(K_{mm})$ 18: return T, A 19: function KFDAWITNESS($Z_{\text{tr}}, k, \lambda$) 20: $Z \leftarrow \text{CONCATENATE}(Z_{\text{tr}})$ 21: $\mathbf{y} = [1] * \text{LEN}(X_{\text{tr}}) + [-1] * \text{LEN}(Y_{\text{tr}})$ 22: $m = \text{LEN}(Z)$ ▷ # Nyström centers 23: $\alpha, Z \leftarrow \text{FDAFALKON}(Z, \mathbf{y}, k, \lambda, m)$ 24: return $h_\lambda = \sum_{i=1}^m \alpha_i k(z_i, \cdot)$ </pre>
--	--

Type-I errors. We report Type-I errors for all three different datasets in Figure 3.

C APPROXIMATE COMPUTATION OF THE KFDA WITNESS

In this section we will use n instead of n_{tr} and m instead of m_{tr} to keep the notation more concise. In A.3, we showed that the exact solution for the estimate of the KFDA witness is given by

$$\hat{h}_\lambda(\cdot) = \sum_{i=1}^{n+m} \hat{\alpha}_i k(z_i, \cdot), \quad (25)$$

$$\hat{\alpha} = \left(\frac{KN_c K}{n+m} + \lambda K \right)^{-1} K \delta. \quad (26)$$

Remark 1. *The problem with computing the KFDA witness is that a naive implementation scales cubically with the pooled sample size. In this section, we thus derive an approach that builds on recent results, that show that one can essentially get optimal convergence guarantees while only using $\mathcal{O}((n+m)^{3/2})$ time. Therefore two steps are needed. First, the solution is approximated with $M = \mathcal{O}((n+m)^{\frac{1}{2}})$ Nyström centers. Second the solution with for the Nyström centers is found via conjugate gradient, where a preconditioner is computed again with only M datapoints.*

We take an approach similar to Rudi et al. (2017); Meanti et al. (2020). We will thus explicitly assume that the function h has the parametric form

$$h_{\tilde{\alpha}}(x) = \sum_{m=1}^M \tilde{\alpha}_m k(x, \tilde{z}_m), \quad (27)$$

with $M = \{\tilde{z}_1, \dots, \tilde{z}_M\} \subseteq \{x_1, \dots, x_n, y_1, \dots, y_m\}$ (we overload notation and use M to denote the set itself as well as its size). We take the notation introduced in Section 4 and constrain to the case $c = \frac{1}{2}$. In this case we can use $N = \begin{pmatrix} P_n & 0 \\ 0 & P_m \end{pmatrix} = \frac{N_c}{2}$, instead of N_c . Note that this only affects the scaling of the solution (if we also scale λ accordingly), which is unimportant for WiTS tests. Using N instead of N_c has the advantage that N itself is idempotent $N = NN^\top$, which makes the following easier. Nevertheless, it is straightforward to use the below algorithm for any $c \in (0, 1)$, simply by using $N_c = \begin{pmatrix} \frac{1}{\sqrt{c}}P_n & 0 \\ 0 & \frac{1}{\sqrt{1-c}}P_m \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{c}}P_n & 0 \\ 0 & \frac{1}{\sqrt{1-c}}P_m \end{pmatrix}$.

In the following we denote with K_{ZM} the $(n+m) \times M$ matrix of entries $k(z_i, \tilde{z}_j)$ and K_{MZ} its transpose. We can then rewrite the terms in our objective

$$\langle \hat{\mu}_P - \hat{\mu}_Q, h_{\tilde{\alpha}} \rangle = \delta^\top K_{ZM} \tilde{\alpha}, \quad (28)$$

$$\begin{aligned} & \langle h_{\tilde{\alpha}}, (\hat{\Sigma} + \lambda \mathbf{1}) h_{\tilde{\alpha}} \rangle \\ &= \tilde{\alpha}^\top \left(\frac{1}{n+m} K_{MZ} N N^\top K_{ZM} + \lambda K_{MM} \right) \tilde{\alpha}. \end{aligned} \quad (29)$$

Let us define $R_{MZ} := K_{MZ} N$. This is a $M \times (n+m)$ matrix. Note that N is the sum of the identity and two 1-sparse matrices, hence computing R_{MZ} requires only $\mathcal{O}((n+m) \cdot M)$ operations.

With our considerations from above we can write the optimal coefficients as

$$\tilde{\alpha}^* = (R_{MZ} R_{MZ}^\top + (n+m)\lambda K_{MM})^{-1} K_{MZ} \delta, \quad (30)$$

$$\Leftrightarrow (R_{MZ} R_{MZ}^\top + (n+m)\lambda K_{MM}) \tilde{\alpha}^* = K_{MZ} \delta \quad (31)$$

Computing $R_{MZ} R_{MZ}^\top$ explicitly costs $\mathcal{O}((n+m)M^2)$ operations and would thus dominate the cost of our previous operations. However, (31) is now exactly in the same form as Eq. (8) in Rudi et al. (2017). Thus from this point onwards we can build on their results to efficiently find a solution.

The key idea of Rudi et al. (2017) is to find an efficient way to precondition the system of linear equations in (31). In analogy, we propose to use the following preconditioner

$$B B^\top = \left(\frac{n+m}{M} R_{MM} R_{MM}^\top + \lambda(n+m) K_{MM} \right)^{-1}, \quad (32)$$

where $R_{MM} := K_{MM} N_M$ and N_M is defined in analogy to N but only with the M Nyström centers. The preconditioner (32) thus corresponds to the ideal preconditioner of the problem without Nyström approximation but only M points to start with.

Using this preconditioner we use t conjugate gradient steps to solve

$$B^\top (R_{MZ} R_{MZ}^\top + (n+m)\lambda K_{MM}) B \beta = B^\top K_{MZ} \delta. \quad (33)$$

If $\hat{\beta}$ is the approximate solution after t steps, we obtain an approximate solution as

$$\hat{\alpha} = B \hat{\beta}. \quad (34)$$

The algorithm is described in Algorithm 2 and has overall complexity of $\mathcal{O}((n_{\text{tr}} + m_{\text{tr}})Mt + M^3)$ in time and $\mathcal{O}(M^2)$.