

Spectral Regularized Kernel Goodness-of-Fit Tests

Omar Hagrass
Bharath K. Sriperumbudur
Bing Li
Department of Statistics
Pennsylvania State University
University Park, PA, 16802 USA

OIH3@PSU.EDU
BKS18@PSU.EDU
BXL9@PSU.EDU

Editor: Krishnakumar Balasubramanian

Abstract

Maximum mean discrepancy (MMD) has enjoyed a lot of success in many machine learning and statistical applications, including non-parametric hypothesis testing, because of its ability to handle non-Euclidean data. Recently, it has been demonstrated in Balasubramanian et al. (2021) that the goodness-of-fit test based on MMD is not minimax optimal while a Tikhonov regularized version of it is, for an appropriate choice of the regularization parameter. However, the results in Balasubramanian et al. (2021) are obtained under the restrictive assumptions of the *mean element* being zero, and the *uniform boundedness condition* on the eigenfunctions of the integral operator. Moreover, the test proposed in Balasubramanian et al. (2021) is not practical as it is not computable for many kernels. In this paper, we address these shortcomings and extend the results to general spectral regularizers that include Tikhonov regularization.

Keywords: Goodness-of-fit test, maximum mean discrepancy, reproducing kernel Hilbert space, covariance operator, U-statistics, Bernstein's inequality, minimax separation, adaptivity, permutation test, spectral regularization

1 Introduction

Given $\mathbb{X}_n := (X_i)_{i=1}^n \stackrel{i.i.d.}{\sim} P$, where P is defined on a measurable space \mathcal{X} , a goodness-of-fit test involves testing $H_0 : P = P_0$ against $H_1 : P \neq P_0$, where P_0 is a fixed known distribution. This is a classical and well-studied problem in statistics for which many tests have been proposed, including the popular ones such as the χ^2 -test and Kolmogorov-Smirnov test (Lehmann and Romano, 2006). However, many of these tests either rely on strong distributional assumptions or cannot handle non-Euclidean data that naturally arise in many modern applications.

A non-parametric testing framework that has gained a lot of popularity over the last decade is based on the notion of reproducing kernel Hilbert space (RKHS) (Aronszajn, 1950) embedding of probability distributions (Smola et al. 2007, Sriperumbudur et al. 2009, Muandet et al. 2017). The power of this framework lies in its ability to handle data that is not necessarily Euclidean. This framework involves embedding a probability measure P into an RKHS, \mathcal{H} through the corresponding mean element, i.e.,

$$\mu_P = \int_{\mathcal{X}} K(\cdot, x) dP(x) \in \mathcal{H},$$

where $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is the unique reproducing kernel (r.k.) associated with \mathcal{H} , with P satisfying $\int_{\mathcal{X}} \sqrt{K(x, x)} dP(x) < \infty$. Using this embedding, a pseudo-metric can be defined on the space of probability measures, called the *maximum mean discrepancy* (MMD) (Gretton et al. 2012, Gretton et al. 2006), as

$$D_{\text{MMD}}(P, Q) = \|\mu_P - \mu_Q\|_{\mathcal{H}},$$

which has the following variational representation (Gretton et al. 2012, Sriperumbudur et al. 2010),

$$D_{\text{MMD}}(P, Q) := \sup_{f \in \mathcal{H}: \|f\|_{\mathcal{H}} \leq 1} \int_{\mathcal{X}} f(x) d(P - Q)(x).$$

We refer the interested reader to (Sriperumbudur et al. 2010, Sriperumbudur 2016, Simon-Gabriel and Schölkopf 2018) for more details about D_{MMD} . Thus given some fixed P_0 , a consistent goodness-of-fit test can be conducted by using the following estimator of D_{MMD}^2 as a test statistic, i.e.,

$$\begin{aligned} \hat{D}_{\text{MMD}}^2(P, P_0) &:= \frac{1}{n(n-1)} \sum_{i \neq j} \langle K(\cdot, X_i) - \mu_0, K(\cdot, X_j) - \mu_0 \rangle_{\mathcal{H}} \\ &= \frac{1}{n(n-1)} \sum_{i \neq j} K(X_i, X_j) - \frac{2}{n} \sum_{i=1}^n \mu_0(X_i) + \|\mu_0\|_{\mathcal{H}}^2, \end{aligned}$$

and using the $1 - \alpha$ quantile of the asymptotic null distribution of $\hat{D}_{\text{MMD}}^2(P, P_0)$ as the critical level (Balasubramanian et al., 2021, Theorem 1), while assuming $\mu_0 := \mu_{P_0}$ and $\|\mu_0\|_{\mathcal{H}}^2$ are computable in closed form. The asymptotic null distribution of $\hat{D}_{\text{MMD}}^2(P, P_0)$ does not have a simple closed form—the distribution is that of an infinite sum of weighted chi-squared random variables with the weights being the eigenvalues of an integral operator associated with the kernel K w.r.t. the distribution P_0 (Serfling, 2009). Assuming $\mu_0 = 0$, recently, (Balasubramanian et al., 2021) showed this test based on \hat{D}_{MMD} to be not optimal in the minimax sense and modified it to achieve a minimax optimal test. Li and Yuan (2019) constructed an optimal test by using the Gaussian kernel on $\mathcal{X} = \mathbb{R}^d$ (the test and analysis can be extended to translation invariant kernels on \mathbb{R}^d using the ideas in Schrab et al., 2021) by allowing the bandwidth of the kernel to shrink to zero as $n \rightarrow \infty$ —this is in contrast to the \hat{D}_{MMD} test where the bandwidth or the kernel parameter is fixed and does not depend on n . By relaxing the requirement of $\mathcal{X} = \mathbb{R}^d$, Balasubramanian et al. (2021) studied the question of optimality for general domains by proposing a regularized test statistic,

$$D_{\lambda}^2(P, P_0) = \sum_{j \geq 1} \frac{\lambda_j}{\lambda_j + \lambda} (\mathbb{E}_P \phi_j)^2, \quad (1)$$

assuming $\mathbb{E}_P \phi_j = 0$ for all j , where $(\lambda_j)_{j \geq 1}$ and $(\phi_j)_{j \geq 1}$ are the eigenvalues and eigenfunctions of an integral operator associated with the kernel K w.r.t. the distribution P_0 , and $\lambda > 0$ is the regularization parameter. Under some regularity conditions, they showed the asymptotic null distribution of an appropriately normalized version of (1) to be the standard normal distribution, using which a minimax optimal goodness-of-fit test was constructed (Balasubramanian et al. 2021, Theorems 2–4). However, this test is impractical and limited

for two reasons: (i) The test requires knowledge of the eigenvalues and eigenfunctions which are only known for a few (K, P_0) pairs, and (ii) $\mathbb{E}_{P_0}\phi_j = 0$ for all j implies that $\mu_0 = 0$, a condition that is not satisfied by any characteristic translation invariant kernels on \mathbb{R}^d , including the Gaussian kernel (Sriperumbudur et al., 2010, 2011). To address these issues, in this paper, we follow an operator theoretic approach and construct a generalized version of (1) based on the idea of spectral regularization that includes (1) as a special case while relaxing the requirement of $\mathbb{E}_{P_0}\phi_j = 0$ for all j —hence resolving (ii)—, and establish its minimax optimality. Moreover, under an additional assumption of P_0 being samplable, i.e., extra samples are available from P_0 , we propose a practical test (i.e., computable) that is also minimax optimal, thereby resolving the issue mentioned in (i).

Before introducing our contributions, we will first introduce the minimax framework pioneered by Burnashev (1979) and Ingster (1987, 1993) to study the optimality of tests, which is essential to understand our contributions and their connection to the results of (Balasubramanian et al., 2021; Li and Yuan, 2019). Let $\phi(\mathbb{X}_n)$ be any test that rejects H_0 when $\phi = 1$ and fails to reject H_0 when $\phi = 0$. Denote the class of all such asymptotic (*resp.* exact) α -level tests to be Φ_α (*resp.* $\Phi_{n,\alpha}$). The Type-II error of a test $\phi \in \Phi_\alpha$ (*resp.* $\in \Phi_{n,\alpha}$) w.r.t. \mathcal{P}_Δ is defined as

$$R_\Delta(\phi) = \sup_{P \in \mathcal{P}_\Delta} \mathbb{E}_{P^n}[1 - \phi],$$

where

$$\mathcal{P}_\Delta := \{P \in \mathcal{C} : \rho^2(P, P_0) \geq \Delta\},$$

is the class of Δ -separated alternatives in the probability metric (or divergence) ρ , with Δ being referred to as the *separation boundary* or *contiguity radius*. Of course, the interest is in letting $\Delta \rightarrow 0$ as $n \rightarrow \infty$ (i.e., shrinking alternatives) and analyzing R_Δ for a given test, ϕ , i.e., whether $R_\Delta(\phi) \rightarrow 0$. In the asymptotic setting, the *minimax separation* or *critical radius* Δ^* is the fastest possible order at which $\Delta \rightarrow 0$ such that $\liminf_{n \rightarrow \infty} \inf_{\phi \in \Phi_\alpha} R_{\Delta^*}(\phi) \rightarrow 0$, i.e., for any Δ such that $\Delta/\Delta^* \rightarrow \infty$, there is no test $\phi \in \Phi_\alpha$ that is consistent over \mathcal{P}_Δ . A test is *asymptotically minimax optimal* if it is consistent over \mathcal{P}_Δ with $\Delta \asymp \Delta^*$. On the other hand, in the non-asymptotic setting, the minimax separation Δ^* is defined as the minimum possible separation, Δ such that $\inf_{\phi \in \Phi_{n,\alpha}} R_\Delta(\phi) \leq \delta$, for $0 < \delta < 1 - \alpha$. A test $\phi \in \Phi_{n,\alpha}$ is called *minimax optimal* if $R_\Delta(\phi) \leq \delta$ for some $\Delta \asymp \Delta^*$. In other words, there is no other α -level test that can achieve the same power with a better separation boundary.

Balasubramanian et al. (2021) consider \mathcal{P}_Δ as

$$\mathcal{P}_\Delta = \left\{ P : \frac{dP}{dP_0} - 1 \in \mathcal{F}(\nu; M), \chi^2(P, P_0) = \left\| \frac{dP}{dP_0} - 1 \right\|_{L^2(P_0)}^2 \geq \Delta \right\}, \quad (2)$$

where $\nu > 0$, and

$$\begin{aligned} \mathcal{F}(\nu; M) := \{f \in L^2(P_0) : & \text{for any } R > 0, \exists f_R \in \mathcal{H} \text{ such that } \|f_R\|_{\mathcal{H}} \leq R, \\ & \text{and } \|f - f_R\|_{L^2(P_0)} \leq MR^{-1/\nu}\}. \end{aligned}$$

\mathcal{P}_Δ in (2) denotes the class of alternatives that are Δ -separated from P_0 in the χ^2 -divergence—alternately, the squared $L^2(P_0)$ norm of the likelihood ratio, $dP/dP_0 - 1$ is lower bounded

by Δ —, while satisfying a smoothness condition. The smoothness condition is imposed on the likelihood ratio and is defined through the rate of approximation of a function in $L^2(P_0)$ by an element in an RKHS ball of radius R . The faster the approximation rate—controlled by ν —, the smoother the function being approximated. $\mathcal{F}(\nu; M)$ is a subspace of a real interpolation space obtained by interpolating \mathcal{H} and $L^2(P_0)$. Particularly, $\nu = 0$ corresponds to an RKHS ball of radius R . Note that (2) requires $P \ll P_0$ (i.e., P is absolutely continuous w.r.t. P_0) so that the Radon-Nikodym derivative dP/dP_0 is well defined. Define

$$\tilde{\mathcal{P}}_\Delta := \left\{ P : \frac{dP}{dP_0} - 1 \in \text{Ran}(L_K^{\frac{1}{2\nu+2}}), \chi^2(P, P_0) = \left\| \frac{dP}{dP_0} - 1 \right\|_{L^2(P_0)}^2 \geq \Delta \right\}, \quad (3)$$

where $L_K : L^2(P_0) \rightarrow L^2(P_0)$, $f \mapsto \int_{\mathcal{X}} K(\cdot, x) f(x) dP_0(x)$ is an integral operator defined by K , and $\text{Ran}(A)$ denotes the range space of A . It follows from (Cucker and Zhou, 2007, Theorem 4.1) that

$$\text{Ran}(L_K^{\frac{1}{2\nu+2}}) \subset \mathcal{F} \left(\nu; 2^{\frac{2\nu+2}{\nu}} \|L_K^{-\frac{1}{2\nu+2}} (dP/dP_0 - 1)\|_{L^2(P_0)}^{\frac{2\nu+2}{\nu}} \right),$$

and if P_0 is non-degenerate, then

$$\mathcal{F}(\nu; M) \subset \text{Ran}(L_K^{\frac{1}{2\nu+2} - \varepsilon}), \forall \varepsilon > 0, \text{ i.e., } \mathcal{F}(\nu; M) \subset \text{Ran}(L_K^\eta), \forall 0 \leq \eta < \frac{1}{2\nu+2}.$$

In this work, we employ an operator theoretic perspective to the goodness-of-fit test problem involving \mathcal{T} (see Section 4 for details), which is a centered version of the integral operator L_K . The centered version is needed to do away with the assumption of $\mu_0 = 0$, which is assumed in Balasubramanian et al., 2021. Therefore, we choose \mathcal{P}_Δ similar to the form in (3) but with L_K replaced by \mathcal{T} . We write it as

$$\mathcal{P} := \mathcal{P}_{\theta, \Delta} := \left\{ P : \frac{dP}{dP_0} - 1 \in \text{Ran}(\mathcal{T}^\theta), \chi^2(P, P_0) = \left\| \frac{dP}{dP_0} - 1 \right\|_{L^2(P_0)}^2 \geq \Delta \right\}, \quad (4)$$

for $\theta > 0$. Note that θ and ν are in inverse proportion to each other and $\theta = \frac{1}{2}$ yields $\text{Ran}(\mathcal{T}^\theta) = \tilde{\mathcal{H}}$, with $0 < \theta < \frac{1}{2}$ yielding interpolation spaces and $\theta = 0$ corresponds to $L^2(P_0)$, where $\tilde{\mathcal{H}}$ is the RKHS induced by the centered kernel,

$$\bar{K}(x, y) = \langle K(\cdot, x) - \mu_0, K(\cdot, y) - \mu_0 \rangle_{\mathcal{H}}.$$

The explicit representation of $\text{Ran}(\mathcal{T}^\theta)$ typically relies on the kernel K and the distribution P_0 . If the kernel K has a Mercer decomposition with respect to eigenfunctions that constitute an orthonormal basis for $L^2(P_0)$, then $\text{Ran}(\mathcal{T}^\theta)$ comprises functions within the span of these orthonormal basis functions. For instance, in the following examples, we present an explicit representation of $\text{Ran}(\mathcal{T}^\theta)$ when P_0 is a uniform distribution on (i) $[0, 1]$, (ii) \mathbb{S}^2 , a unit sphere, and (iii) when P_0 is a standard Gaussian distribution on \mathbb{R} . In this context, $\text{Ran}(\mathcal{T}^\theta)$ can be expressed in terms of Fourier basis in Example 1, spherical harmonic basis in Example 2 and Hermite polynomials basis in Example 3.

Example 1 (Uniform distribution on $[0, 1]$) Let P_0 be the uniform distribution defined on $[0, 1]$ with

$$K(x, y) = a_0 + \sum_{k \neq 0} |k|^{-\beta} e^{\sqrt{-1}2\pi kx} e^{-\sqrt{-1}2\pi ky}, \quad a_0 \geq 0, \quad \beta > 1. \quad (5)$$

Then

$$\text{Ran}(\mathcal{T}^\theta) = \left\{ \sum_{k \neq 0} a_k e^{\sqrt{-1}2\pi kx} : \sum_{k \neq 0} a_k^2 k^{2\theta\beta} < \infty \right\}.$$

Note that the s -order Sobolev space defined on $[0, 1]$ is given by

$$\mathcal{W}^{s,2} := \left\{ f(x) = \sum_{k \in \mathbb{Z}} a_k e^{\sqrt{-1}2\pi kx}, \quad x \in [0, 1] : \sum_{k \in \mathbb{Z}} (1 + k^2)^s a_k^2 < \infty \right\}.$$

Since $\sum_{k \neq 0} k^{2\theta\beta} a_k^2 \leq \sum_{k \in \mathbb{Z}} (1 + k^2)^{\theta\beta} a_k^2$, it follows that $\mathcal{W}^{s,2} \subset \text{Ran}(\mathcal{T}^\theta)$. This means, if $u := \frac{dP}{dP_0} - 1 \in \mathcal{W}^{s,2}$, then $u \in \text{Ran}(\mathcal{T}^\theta)$ with $\theta = \frac{s}{\beta}$. An example of a kernel that follows the form in (5) is the periodic spline kernel, represented as $\tilde{K}(x, y) = \frac{(-1)^{r-1}}{(2r)!} B_{2r}([x - y])$, where B_r denotes the Bernoulli polynomial, which is generated by the generating function $\frac{te^{tx}}{e^t - 1} = \sum_{r=0}^{\infty} B_r(x) \frac{t^r}{r!}$, and $[t]$ denotes the fractional part of t . Then using the formula $B_{2r}(x) = \frac{(-1)^{r-1}(2r)!}{(2\pi)^{2r}} \sum_{k \neq 0}^{\infty} |k|^{-2r} e^{\sqrt{-1}2\pi kx}$, it can be demonstrated that $\tilde{K}(x, y) = \sum_{k \neq 0} (2\pi|k|)^{-2r} e^{\sqrt{-1}2\pi kx} e^{-\sqrt{-1}2\pi ky}$ (see Wahba, 1990, page 21 for details).

Example 2 (Uniform distribution on \mathbb{S}^2) Let P_0 be a uniform distribution on $\mathcal{X} = \mathbb{S}^2$, where \mathbb{S}^2 denotes the unit sphere. Let

$$K(x, y) := \sum_{k=1}^{\infty} \sum_{j=-k}^k \lambda_k Y_k^j(\theta_x, \phi_x) Y_k^j(\theta_y, \phi_y). \quad (6)$$

where $x = (\sin \theta_x \cos \phi_x, \sin \theta_x \sin \phi_x, \cos \theta_x)$, $y = (\sin \theta_y \cos \phi_y, \sin \theta_y \sin \phi_y, \cos \theta_y)$ with $0 < \theta_x, \theta_y < \pi$, $0 < \phi_x, \phi_y < 2\pi$, and

$$Y_k^j(\theta, \phi) := \sqrt{\frac{(2k+1)(k-j)!}{4\pi(k+j)!}} p_k^j(\cos \theta) e^{\sqrt{-1}j\phi},$$

with $p_k^j(x) := (-1)^j (1-x^2)^{\frac{j}{2}} \frac{d^j p_k(x)}{dx^j}$ and $p_k(x) := \frac{1}{k! 2^k} \frac{d^k (x^2 - 1)^k}{dx^k}$. Here $(Y_k^j(\theta, \phi))_{j,k}$ denote the spherical harmonics which form an orthonormal basis in $L^2(\mathbb{S}^2)$. If $\sum_{k=1}^{\infty} (2k+1)\lambda_k < \infty$, then

$$\text{Ran}(\mathcal{T}^\theta) = \left\{ \sum_{k=1}^{\infty} \sum_{j=-k}^k a_{kj} Y_k^j(\theta_x, \phi_x) : \sum_{k=1}^{\infty} \sum_{j=-k}^k a_{kj}^2 \lambda_k^{-2\theta} < \infty \right\}.$$

Many common kernels take the form in (6). For example, Minh et al. (2006, Theorem 2 and 3) provide explicit expressions for the eigenvalues corresponding to Gaussian and polynomial kernels on the sphere.

Example 3 (Gaussian distribution with Mehler kernel on \mathbb{R}) Let P_0 be a standard Gaussian distribution on \mathbb{R} and K be the Mehler kernel, i.e.,

$$K(x, y) := \frac{1}{\sqrt{1 - \rho^2}} \exp \left(-\frac{\rho^2(x^2 + y^2) - 2\rho xy}{2(1 - \rho^2)} \right),$$

for $0 < \rho < 1$. Then

$$\text{Ran}(\mathcal{T}^\theta) = \left\{ \sum_{k=1}^{\infty} a_k \gamma_k(x) : \sum_{k=1}^{\infty} a_k^2 e^{-2k\theta \log \rho} < \infty \right\},$$

where $\gamma_k(x) = \frac{H_k(x)}{\sqrt{k!}}$, and $H_k(x) = (-1)^k e^{x^2/2} \frac{d^k}{dx^k} e^{-x^2/2}$ is the Hermite polynomial. Steinwart and Scovel (2012, Theorem 4.6) provides an interpretation of $\text{Ran}(\mathcal{T}^\theta)$ as a real interpolation of $L^2(P_0)$ and $\tilde{\mathcal{H}}$. Therefore, with the kernel being fixed, the influence of P_0 on $\text{Ran}(\mathcal{T}^\theta)$ can be understood as follows. Suppose $P_{0,\sigma_i} := N(0, \sigma_i^2)$, $i = 1, 2$. It is easy to verify that $L^2(P_{0,\sigma_1}) \subset L^2(P_{0,\sigma_2})$ if $\sigma_2 < \sigma_1$, which implies that $\text{Ran}(\mathcal{T}_1^\theta) \subseteq \text{Ran}(\mathcal{T}_2^\theta)$, where \mathcal{T}_i is the integral operator defined w.r.t. P_{0,σ_i} , $i = 1, 2$. Based on this intuition, in the context of this example, choosing P_0 as a Gaussian distribution with variance larger (resp. smaller) than 1 yields a smaller (resp. larger) range space than that mentioned above.

With this background, we now present our contributions.

1.1 Contributions

The main contributions of the paper are as follows:

(i) First, in Theorem 1, we show that the test based on $\hat{D}_{\frac{-2\theta}{2\theta+1}}^2$ (we refer to it as the MMD test) cannot achieve a separation boundary better than $n^{\frac{-2\theta}{2\theta+1}}$ w.r.t. \mathcal{P} defined in (4). This is an extension and generalization of (Balasubramanian et al., 2021, Theorem 1), which only shows such a claim for $\theta = \frac{1}{2}$ in an asymptotic setting, assuming $\mu_0 = 0$ and the *uniform boundedness condition*, $\sup_i \|\phi_i\|_\infty < \infty$, where $(\phi_i)_i$ are the eigenfunctions of \mathcal{T} . In contrast, Theorem 1 presents the result both by assuming and not assuming the uniform boundedness condition. Note that the uniform boundedness condition $\sup_i \|\phi_i\|_\infty < \infty$ is not satisfied in many scenarios (of course, it is satisfied in Example 1). For example, as illustrated in Minh et al. (2006, Theorem 5), for $\mathcal{X} = \mathbb{S}^{d-1}$, representing the d -dimensional unit sphere, $\sup_i \|\phi_i\|_\infty = \infty$ for all $d \geq 3$ when using any kernel of the form $K(x, y) = f(\langle x, y \rangle_2)$, where $x, y \in \mathcal{X}$ and f is continuous (see Example 2). The Gaussian kernel on \mathbb{S}^{d-1} serves as an instance of such a kernel. Moreover, the condition $\mu_0 = 0$ is not satisfied by any characteristic kernel on general domain \mathcal{X} and therefore excludes popular kernels such as Gaussian, Matérn, and inverse multiquadric on \mathbb{R}^d . Relaxing these two assumptions allows a large class of (K, P_0) pairs to be handled by Theorem 1.

(ii) Note that the separation boundary of the MMD test depends only on the smoothness of $dP/dP_0 - 1$, which is determined by θ but is completely oblivious to the *intrinsic dimensionality* of the RKHS, $\tilde{\mathcal{H}}$, which is controlled by the decay rate of the eigenvalues of \mathcal{T} . To this end, by taking into account the intrinsic dimensionality of $\tilde{\mathcal{H}}$, we show in Theorem 2 that the minimax separation w.r.t. \mathcal{P} is $n^{-\frac{4\theta\beta}{4\theta\beta+1}}$ for $\theta > \frac{1}{2}$ if $\lambda_i \asymp i^{-\beta}$, $\beta > 1$,

i.e., the eigenvalues of \mathcal{T} decay at a polynomial rate β , and is $\sqrt{\log n}/n$ if $\lambda_i \asymp e^{-i}$, i.e., exponential decay. These results clearly establish the non-optimality of the MMD-based test. Theorem 2, which is non-asymptotic, generalizes the asymptotic version of (Balasubramanian et al., 2021, Theorem 4) without requiring the uniform boundedness condition and also recovers it under the uniform boundedness condition, while not requiring $\mu_0 = 0$ for both these results. Moreover, even under the uniform boundedness condition, while (Balasubramanian et al., 2021, Theorem 4) provides a bound on the minimax separation for $\frac{1}{2} > \theta > \frac{1}{2\beta}$, we improve this range in Theorem 2 by showing the minimax separation for $\theta > \frac{1}{4\beta}$.

(iii) In Section 4, we employ an operator theoretic perspective to the regularization idea presented in Balasubramanian et al. (2021) that allows us to generalize the idea to handle general spectral regularizers, without requiring $\mu_0 = 0$. More precisely, we propose a statistic of the form $\eta_\lambda(P, P_0) := \left\| g_\lambda^{1/2}(\Sigma_0)(\mu_P - \mu_{P_0}) \right\|_{\mathcal{H}}^2$, which when $g_\lambda(x) = (x + \lambda)^{-1}$, i.e., Tikhonov regularization, and $\mu_0 = 0$ reduces to the regularized statistic in (1). Here Σ_0 corresponds to the centered covariance operator w.r.t. P_0 . Assuming μ_0 and Σ_0 are computable, we propose a spectral regularized test based on η_λ and provide sufficient conditions on g_λ for the test to be minimax optimal w.r.t. \mathcal{P} (see Theorems 3, 4 and Corollaries 5, 6). Compared to the results in (Balasubramanian et al., 2021), we provide general sufficient conditions on the separation boundary for any bounded kernel and show the minimax optimality in the non-asymptotic setting for a wider range of θ , both with and without the uniform boundedness condition (see Theorem 4). However, the drawback of the test is that one needs first to compute the eigenvalues and eigenfunctions of Σ_0 which is not possible for many (K, P_0) pairs. Thus we refer to this test as the *Oracle test*.

(iv) To address the shortcomings with the Oracle test, in Section 4.2, we assume that P_0 is samplable, i.e., P_0 can be sampled to generate new samples. Based on these samples, we propose a test statistic defined in (11) that involves using the estimators of μ_0 and Σ_0 in η_λ . We show that such a test statistic can be computed only through matrix operations and by solving a finite-dimensional eigensystem (see Theorem 7). We present two approaches to compute the critical level of the test. In Section 4.3, we compute the critical level based on a concentration inequality and refer to the corresponding test as *spectral regularized concentration test* (SRCT), and in Section 4.4, we employ permutation testing (e.g., Lehmann and Romano 2006, Pesarin and Salmaso 2010, Kim et al. 2022), which we refer to as the *spectral regularized permutation test* (SRPT), leading to a critical level that is easy to compute (see Theorems 8 and 10). We show that both these tests are minimax optimal w.r.t. \mathcal{P} (see Theorems 9 and 11). Note that under these additional samples from P_0 , a goodness-of-fit test can be seen as a two-sample test, and therefore SRCT and SRPT can be interpreted as two-sample tests. Recently, Hagrass et al. (2024) developed a *spectral regularized kernel two-sample test* (SR2T) and showed it to be minimax optimal for a suitable class of alternatives. In this work, we show that SRCT and SRPT have better separation rates than those of SR2T for the range of θ , where all these tests are not minimax.

(v) The minimax optimal separation rate in the proposed tests (SRCT and SRPT) is tightly controlled by the choice of the regularization parameter, λ , which in turn depends on the unknown parameters, θ and β (in the case of the polynomial decay of the eigenvalues of

\mathcal{T}). Therefore, to make these tests completely data-driven, in Section 4.5, we present an adaptive version of both tests by aggregating tests over different λ (see Theorems 14 and 16) and show the resulting tests to be minimax optimal up to a $\sqrt{\log n}$ factor in case of the SRCT (see Theorem 15) and $\log \log n$ factor in case of SRPT (see Theorem 17). In contrast, (Balasubramanian et al., 2021, Theorem 5) considers an adaptive and asymptotic version of the Oracle test under $\mu_0 = 0$ and the uniform boundedness condition, where it only adapts over θ assuming β is known, with β being the polynomial decay rate of the eigenvalues of \mathcal{T} .

(vi) Through numerical simulations on benchmark data, in Section 5, we demonstrate the superior performance of the proposed spectral regularized tests in comparison to the MMD test based on $\hat{D}_{\text{MMD}}(P, P_0)$, Energy test (Szekely and Rizzo, 2004) based on the energy distance, Kolmogorov-Smirnov (KS) test (Puritz et al., 2022; Fasano and Franceschini, 1987), and SR2T.

1.2 Comparison to Hagrass et al. (2024)

As mentioned in (iv) of Section 1.1, the proposed goodness-of-fit tests (SRCT and SRPT) can be seen as two-sample tests because of access to additional samples from P_0 . Similar to the two-sample test SR2T proposed in Hagrass et al. (2024), these tests also employ the spectral regularization approach of Hagrass et al. (2024) and their analysis uses many technical results developed in Hagrass et al. (2024). Therefore, to emphasize the conceptual and technical novelty of our work, in this section, we compare and contrast our setup and results to that of Hagrass et al. (2024).

(i) **Alternative space:** In this paper, the alternative space, \mathcal{P}_Δ shown in (4) involves a smoothness condition that is defined with respect to the function $u := \frac{dP}{dP_0} - 1$. In contrast, the smoothness condition in Hagrass et al. (2024) was defined through $\frac{dP}{dR} - 1$, where $R = \frac{P+P_0}{2}$. The separation boundary in this paper is measured in χ^2 -distance, i.e., $\chi^2(P, P_0)$ compared to the Hellinger distance between P and P_0 (which is topologically equivalent to $\chi^2(P, R)$) as in Hagrass et al. (2024). Since the χ^2 -divergence dominates the Hellinger distance, the notion of separation considered in this paper is stronger than the one considered in Hagrass et al. (2024). These changes were made to leverage the knowledge of P_0 in the goodness-of-fit problem (which is not available in the two-sample problem), resulting in a better separation boundary than that achieved by the test proposed in Hagrass et al. (2024).

(ii) **Estimation of the covariance operator, Σ_0 :** In Hagrass et al. (2024), the covariance operator Σ_0 is defined with respect to the average probability measure $R := \frac{P+P_0}{2}$, which means two sets of samples are required to estimate it and therefore, the estimation error is controlled by the minimum of sizes of two sets of samples. However, in this paper, since we are considering a goodness-of-fit problem where the null distribution P_0 is known, we can utilize this knowledge by defining the covariance operator Σ_0 with respect to P_0 , which means the estimation error is controlled only by the samples from P_0 . Since we do not have any budget constraints on sampling from P_0 , the estimation error of Σ_0 can be controlled at a desired level for a large enough sample size. Therefore, we investigate the required number of i.i.d. samples s to be drawn from P_0 to estimate Σ_0 to achieve a similar

separation boundary as the oracle test, which assumes Σ_0 is exactly known in closed form. Both in this work and Hagrass et al. (2024), while the separation rates are determined by the minimum size of the two sets of samples, since the sample size associated with P_0 in this work can be chosen to be large enough, the separation rate will be controlled only by the one sample size. Therefore this work yields better separation rates than those in Hagrass et al. (2024)—also see (iv) in Section 1.1—as it should be since a goodness-fit-testing problem is simpler than a two-sample testing problem.

(iii) **Spectral regularized concentration test (SRCT):** While SRPT proposed in this paper shares in principle the similar ideas of permutation testing as in Hagrass et al. (2024), the proposed SRCT involves a concentration inequality based test threshold that was not considered in Hagrass et al. (2024). While the analysis of SRPT uses multiple technical results developed in Hagrass et al. (2024)—of course, with some deviations because of a different alternate space and estimator for the covariance operator—the analysis of SRCT requires establishing new technical results for the estimation error bounds between Σ_0 and $\hat{\Sigma}_0$ (see Lemmas A.4 and A.5), and the estimation error between $\hat{\mathcal{N}}_2(\lambda)$ and $\mathcal{N}_2(\lambda)$, where $\mathcal{N}_2(\lambda) := \|\Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}\|_{\mathcal{L}^2(\mathcal{H})}$ and $\Sigma_{0,\lambda} := \Sigma_0 + \lambda \mathbf{I}$ (see Lemmas A.6 and A.7).

2 Definitions & Notation

For a topological space \mathcal{X} , $L^r(\mathcal{X}, \mu)$ denotes the Banach space of r -power ($r \geq 1$) μ -integrable function, where μ is a finite non-negative Borel measure on \mathcal{X} . For $f \in L^r(\mathcal{X}, \mu) =: L^r(\mu)$, $\|f\|_{L^r(\mu)} := (\int_{\mathcal{X}} |f|^r d\mu)^{1/r}$ denotes the L^r -norm of f . $\mu^n := \mu \times \dots \times \mu$ is the n -fold product measure. \mathcal{H} denotes a reproducing kernel Hilbert space with a reproducing kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. $[f]_{\sim}$ denotes the equivalence class associated with $f \in L^r(\mathcal{X}, \mu)$, that is the collection of functions $g \in L^r(\mathcal{X}, \mu)$ such that $\|f - g\|_{L^r(\mu)} = 0$. For two measures P and Q , $P \ll Q$ denotes that P is dominated by Q which means, if $Q(A) = 0$ for some measurable set A , then $P(A) = 0$. Let H_1 and H_2 be abstract Hilbert spaces. $\mathcal{L}(H_1, H_2)$ denotes the space of bounded linear operators from H_1 to H_2 . For $S \in \mathcal{L}(H_1, H_2)$, S^* denotes the adjoint of S . $S \in \mathcal{L}(H) := \mathcal{L}(H, H)$ is called self-adjoint if $S^* = S$. For $S \in \mathcal{L}(H)$, $\text{Tr}(S)$, $\|S\|_{\mathcal{L}^2(H)}$, and $\|S\|_{\mathcal{L}^\infty(H)}$ denote the trace, Hilbert-Schmidt and operator norms of S , respectively. For $x, y \in H$, $x \otimes_H y$ is an element of the tensor product space of $H \otimes H$ which can also be seen as an operator from $H \rightarrow H$ as $(x \otimes_H y)z = x\langle y, z \rangle_H$ for any $z \in H$.

For constants a and b , $a \lesssim b$ (resp. $a \gtrsim b$) denotes that there exists a positive constant c (resp. c') such that $a \leq cb$ (resp. $a \geq c'b$). $a \asymp b$ denotes that there exists positive constants c and c' such that $cb \leq a \leq c'b$. We denote $[\ell]$ for $\{1, \dots, \ell\}$.

3 Non-optimality of \hat{D}_{MMD}^2 test

Assuming $\mu_0 = 0$, (Balasubramanian et al., 2021) established the non-optimality of the MMD-based goodness-of-fit test. In this section, we extend this result in two directions by not assuming $\mu_0 = 0$ and by considering the setting of non-asymptotic minimax compared to the asymptotic minimax setting of Balasubramanian et al. (2021). The key to achieving these extensions is an operator representation of D_{MMD}^2 , which we obtain below. To this end, we make the following assumption throughout the paper.

(A_0) $(\mathcal{X}, \mathcal{B})$ is a second countable (i.e., completely separable) space endowed with Borel σ -algebra \mathcal{B} . (\mathcal{H}, K) is an RKHS of real-valued functions on \mathcal{X} with a continuous reproducing kernel K satisfying $\sup_x K(x, x) \leq \kappa$.

The continuity of K ensures that $K(\cdot, x) : \mathcal{X} \rightarrow \mathcal{H}$ is Bochner-measurable for all $x \in \mathcal{X}$, which along with the boundedness of K ensures that μ_P and μ_{P_0} are well-defined (Dinicleanu, 2000). Also, the separability of \mathcal{X} along with the continuity of K ensures that \mathcal{H} is separable (Steinwart and Christmann 2008, Lemma 4.33). Therefore,

$$\begin{aligned} D_{\text{MMD}}^2(P, P_0) &= \|\mu_P - \mu_{P_0}\|_{\mathcal{H}}^2 = \left\langle \int_{\mathcal{X}} K(\cdot, x) d(P - P_0)(x), \int_{\mathcal{X}} K(\cdot, x) d(P - P_0)(x) \right\rangle_{\mathcal{H}} \\ &= \left\langle \int_{\mathcal{X}} K(\cdot, x) u(x) dP_0(x), \int_{\mathcal{X}} K(\cdot, x) u(x) dP_0(x) \right\rangle_{\mathcal{H}}, \end{aligned} \quad (7)$$

where $u = \frac{dP}{dP_0} - 1$. As done in (Hagrass et al., 2024), by defining $\mathfrak{I} : \mathcal{H} \rightarrow L^2(P_0)$, $f \mapsto [f - \mathbb{E}_{P_0} f]_{\sim}$, where $\mathbb{E}_{P_0} f = \int_{\mathcal{X}} f(x) dP_0(x)$, it follows from (Sriperumbudur and Sterge, 2022, Proposition C.2) that $\mathfrak{I}^* : L^2(P_0) \rightarrow \mathcal{H}$, $f \mapsto \int K(\cdot, x) f(x) dP_0(x) - \mu_{P_0} \mathbb{E}_{P_0} f$. Also, it follows from (Sriperumbudur and Sterge, 2022, Proposition C.2) that $\mathcal{T} = \Upsilon - (1 \otimes_{L^2(P_0)} 1) \Upsilon - \Upsilon (1 \otimes_{L^2(P_0)} 1) + (1 \otimes_{L^2(P_0)} 1) \Upsilon (1 \otimes_{L^2(P_0)} 1)$, where $\Upsilon : L^2(P_0) \rightarrow L^2(P_0)$, $f \mapsto \int K(\cdot, x) f(x) dP_0(x)$ and $\mathcal{T} := \mathfrak{I} \mathfrak{I}^* : L^2(P_0) \rightarrow L^2(P_0)$. Note that \mathcal{T} is a trace class operator, and thus compact since K is bounded. Also, \mathcal{T} is self-adjoint and positive, and therefore spectral theorem (Reed and Simon, 1980, Theorems VI.16, VI.17) yields that

$$\mathcal{T} = \sum_{i \in I} \lambda_i \tilde{\phi}_i \otimes_{L^2(P_0)} \tilde{\phi}_i,$$

where $(\lambda_i)_i \subset \mathbb{R}^+$ are the eigenvalues and $(\tilde{\phi}_i)_i$ are the orthonormal system of eigenfunctions (strictly speaking classes of eigenfunctions) of \mathcal{T} that span $\text{Ran}(\mathcal{T})$ with the index set I being either countable in which case $\lambda_i \rightarrow 0$ or finite. In this paper, we assume that the set I is countable, i.e., infinitely many eigenvalues. Since $\tilde{\phi}_i$ represents an equivalence class in $L^2(P_0)$, by defining $\phi_i := \frac{\mathfrak{I}^* \tilde{\phi}_i}{\lambda_i}$, it is clear that $\mathfrak{I} \phi_i = [\phi_i - \mathbb{E}_{P_0} \phi_i]_{\sim} = \tilde{\phi}_i$ and $\phi_i \in \mathcal{H}$. Throughout the paper, ϕ_i refers to this definition.

Using these definitions, it follows from (7) that

$$D_{\text{MMD}}^2(P, P_0) = \langle \mathfrak{I}^* u, \mathfrak{I}^* u \rangle_{\mathcal{H}} = \langle \mathcal{T} u, u \rangle_{L^2(P_0)} = \sum_{i \geq 1} \lambda_i \langle u, \tilde{\phi}_i \rangle_{L^2(P_0)}^2.$$

The above expression was already obtained by (Balasubramanian et al., 2021, p. 6) but through Mercer's representation of K . Here we obtain it alternately through the operator representation, which will turn out to be crucial for the rest of the paper. This representation also highlights the limitation of D_{MMD} that D_{MMD} might not capture the difference between P and P_0 if they differ in the higher Fourier coefficients of u , i.e., $\langle u, \tilde{\phi}_i \rangle_{L^2(P_0)}$ for large i , since $(\lambda_i)_i$ is a decreasing sequence. On the other hand, $\chi^2(P || P_0) = \|u\|_{L^2(P_0)}^2 = \sum_{i \geq 1} \langle u, \tilde{\phi}_i \rangle_{L^2(P_0)}^2$ if $u \in \text{span}\{\tilde{\phi}_i : i \in I\}$, does not suffer from such an issue. The following result shows that the test based on \hat{D}_{MMD}^2 cannot achieve a separation boundary of order better than $n^{\frac{-2\theta}{2\theta+1}}$.

Theorem 1 (Separation boundary of MMD test) *Let $n \geq 2$ and*

$$\sup_{P \in \mathcal{P}} \left\| \mathcal{T}^{-\theta} u \right\|_{L^2(P_0)} < \infty.$$

Then for any $\alpha > 0$, $\delta > 0$, $P_{H_0}\{\hat{D}_{\text{MMD}}^2 \geq \gamma\} \leq \alpha$,

$$\inf_{P \in \mathcal{P}} P_{H_1}\{\hat{D}_{\text{MMD}}^2 \geq \gamma\} \geq 1 - \delta,$$

where $\gamma = \frac{4\kappa}{\sqrt{\alpha n}}$,

$$\Delta_n := \Delta = c(\alpha, \delta) n^{\frac{-2\theta}{2\theta+1}},$$

and $c(\alpha, \delta) \asymp \max\{\alpha^{-1/2}, \delta^{-1}\}$. Furthermore if $\Delta_n < d_\alpha n^{\frac{-2\theta}{2\theta+1}}$ for some $d_\alpha > 0$ and one of the following holds: (i) $\theta \geq \frac{1}{2}$, (ii) $\sup_i \|\phi_i\|_\infty < \infty$, $\theta > 0$, then for any decay rate of $(\lambda_i)_i$, there exists \tilde{k}_δ such that for all $n \geq \tilde{k}_\delta$,

$$\inf_{P \in \mathcal{P}} P_{H_1}\{\hat{D}_{\text{MMD}}^2 \geq \gamma\} < \delta.$$

Remark 1 *Note that the above theorem also holds asymptotically if the testing threshold γ is chosen as the $(1 - \alpha)$ -quantile of the asymptotic distribution of \hat{D}_{MMD}^2 under H_0 , thereby extending (Balasubramanian et al., 2021, Theorem 1), which only considers $\theta = \frac{1}{2}$ but assuming $\mu_0 = 0$. In fact, the result holds for any threshold that converges in probability to such an asymptotic quantile.*

By providing the minimax separation rate w.r.t. \mathcal{P} , the following result demonstrates the non-optimality of the MMD test presented in Theorem 1.

Theorem 2 (Minimax separation boundary) *If $\lambda_i \asymp i^{-\beta}$, $\beta > 1$, then there exists $c(\alpha, \delta)$ such that if*

$$\Delta_n \leq c(\alpha, \delta) n^{\frac{-4\theta\beta}{4\theta\beta+1}}, \quad 0 \leq \delta \leq 1 - \alpha,$$

then $R_{\Delta_n}^ := \inf_{\phi \in \Phi_{n,\alpha}} R_{\Delta_n, M}(\phi) \geq \delta$, provided one the following holds: (i) $\theta \geq \frac{1}{2}$, (ii) $\sup_i \|\phi_i\|_\infty < \infty$, $\theta \geq \frac{1}{4\beta}$, where $R_{\Delta_n}(\phi) := \sup_{P \in \mathcal{P}} \mathbb{E}_{P^n}[1 - \phi]$.*

Suppose $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$, $\theta > 0$. Then there exists $c(\alpha, \delta, \theta)$ and k such that if

$$\Delta_n \leq c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n},$$

and $n \geq k$, then for any $0 \leq \delta \leq 1 - \alpha$, $R_{\Delta_n}^ \geq \delta$.*

Remark 2 (i) *Since $\inf_{\beta > 1} \frac{4\theta\beta}{4\theta\beta+1} = \frac{4\theta}{4\theta+1} > \frac{2\theta}{2\theta+1}$ and $1 > \frac{2\theta}{2\theta+1}$ for any $\theta > 0$, it follows that the separation boundary of MMD is larger than the minimax separation boundary w.r.t. \mathcal{P} irrespective of the decay rate of the eigenvalues of \mathcal{T} .*

(ii) *In the setting of polynomial decay, Theorem 2 generalizes (Balasubramanian et al., 2021, Theorem 4) in two ways: (a) When the uniform boundedness condition holds, the range of θ for which the minimax separation rate holds is extended from $\frac{1}{2\beta} < \theta < \frac{1}{2}$ to $\theta > \frac{1}{4\beta}$, and (b) minimax separation is also obtained without assuming the uniform boundedness condition.*

(iii) The uniform boundedness condition, $\sup_i \|\phi_i\|_\infty < \infty$ does not hold in general. For example, the Gaussian kernel on S^{d-1} , $d \geq 3$, does not satisfy the uniform boundedness condition (Minh et al., 2006, Theorem 5), while the Gaussian kernel on \mathbb{R}^d for any d satisfies the uniform boundedness condition (Steinwart et al., 2006). In this paper, we provide results both with and without the uniform boundedness condition to understand its impact on the behavior of the test. Such a condition has also been used in the analysis of the impact of regularization in kernel learning (see Mendelson and Neeman 2010, p. 531).

4 Spectral regularized MMD test

In this section, we propose a spectral regularized version of the MMD test and show it to be minimax optimal w.r.t. \mathcal{P} . The proposed test statistic is based on the *spectral regularized discrepancy*, which is defined as

$$\eta_\lambda(P, P_0) := \langle \mathcal{T}g_\lambda(\mathcal{T})u, u \rangle_{L^2(P_0)}, \quad (8)$$

where $u = \frac{dP}{dP_0} - 1$, $g_\lambda : (0, \infty) \rightarrow (0, \infty)$ is a *spectral regularizer* that satisfies $\lim_{\lambda \rightarrow 0} xg_\lambda(x) \asymp 1$ (more concrete assumptions on g_λ will be introduced later), and

$$g_\lambda(\mathcal{B}) := \sum_{i \geq 1} g_\lambda(\tau_i)(\psi_i \otimes_H \psi_i) + g_\lambda(0) \left(\mathbf{I} - \sum_{i \geq 1} \psi_i \otimes_H \psi_i \right),$$

with \mathcal{B} being any compact, self-adjoint operator defined on a separable Hilbert space, H . Here $(\tau_i, \psi_i)_i$ are the eigenvalues and eigenfunctions of \mathcal{B} which enjoys the spectral representation, $\mathcal{B} = \sum_i \tau_i \psi_i \otimes_H \psi_i$. The well known *Tikhonov regularizer*, $(\mathcal{B} + \lambda \mathbf{I})^{-1}$, which is used in Balasubramanian et al. (2021), is obtained as a special case by choosing $g_\lambda(x) = (x + \lambda)^{-1}$.

The key idea in proposing η_λ is based on the intuition that $\mathcal{T}g_\lambda(\mathcal{T}) \approx \mathbf{I}$ for sufficiently small λ so that $\eta_\lambda(P, P_0) \approx \|u\|_{L^2(P_0)}^2$, and therefore does not suffer from the limitation of $D_{\text{MMD}}^2(P, P_0)$ as aforementioned in Section 3 (see Lemma A.2). Using η_λ , in the following, we present details about the construction of the test statistic and the test. First, we provide an alternate representation for η_λ which is useful to construct the test statistic. Define the *covariance operator*,

$$\begin{aligned} \Sigma_0 &:= \Sigma_{P_0} = \int_{\mathcal{X}} (K(\cdot, x) - \mu_{P_0}) \otimes_{\mathcal{H}} (K(\cdot, x) - \mu_{P_0}) dP_0(x) \\ &= \frac{1}{2} \int_{\mathcal{X} \times \mathcal{X}} (K(\cdot, x) - K(\cdot, y)) \otimes_{\mathcal{H}} (K(\cdot, x) - K(\cdot, y)) dP_0(x) dP_0(y), \end{aligned}$$

which is a positive, self-adjoint, trace-class operator. It can be shown (Sriperumbudur and Sterge, 2022, Proposition C.2) that $\Sigma_0 = \mathfrak{I}^* \mathfrak{I} : \mathcal{H} \rightarrow \mathcal{H}$. Using this representation in (8) yields

$$\begin{aligned} \eta_\lambda(P, P_0) &= \langle \mathcal{T}g_\lambda(\mathcal{T})u, u \rangle_{L^2(P_0)} = \langle \mathfrak{I} \mathfrak{I}^* g_\lambda(\mathfrak{I} \mathfrak{I}^*)u, u \rangle_{L^2(P_0)} \stackrel{(\dagger)}{=} \langle \mathfrak{I} g_\lambda(\mathfrak{I}^* \mathfrak{I}) \mathfrak{I}^* u, u \rangle_{L^2(P_0)} \\ &= \langle g_\lambda(\Sigma_0) \mathfrak{I}^* u, \mathfrak{I}^* u \rangle_{\mathcal{H}} = \langle g_\lambda(\Sigma_0)(\mu_P - \mu_{P_0}), \mu_P - \mu_{P_0} \rangle_{\mathcal{H}} \\ &= \left\| g_\lambda^{1/2}(\Sigma_0)(\mu_P - \mu_{P_0}) \right\|_{\mathcal{H}}^2, \end{aligned} \quad (9)$$

where (\dagger) follows from (Hagras et al. (2024), Lemma A.8(i) by replacing Σ_{PQ} by Σ_0) that $\mathcal{T}g_\lambda(\mathcal{T}) = \mathcal{I}g_\lambda(\Sigma_0)\mathcal{I}^*$. Throughout the paper, we assume that g_λ satisfies the following:

$$\begin{aligned} (A_1) \quad & \sup_{x \in \Gamma} |xg_\lambda(x)| \leq C_1, & (A_2) \quad & \sup_{x \in \Gamma} |\lambda g_\lambda(x)| \leq C_2, \\ (A_3) \quad & \sup_{\{x \in \Gamma: xg_\lambda(x) < B_3\}} |B_3 - xg_\lambda(x)|x^{2\varphi} \leq C_3\lambda^{2\varphi}, & (A_4) \quad & \inf_{x \in \Gamma} g_\lambda(x)(x + \lambda) \geq C_4, \end{aligned}$$

where $\Gamma := [0, \kappa]$, $\varphi \in (0, \xi]$ and the constant ξ is called the *qualification* of g_λ . C_1, C_2, C_3, B_3 and C_4 are finite positive constants (all independent of $\lambda > 0$). Note that these assumptions are quite standard in the inverse problem literature (see e.g., Engl et al., 1996) and spectral regularized kernel ridge regression (Bauer et al., 2007), except for (A_3) , which is replaced by a stronger version—the stronger version of (A_3) takes supremum over whole Γ . Recently, however, in a two-sample testing scenario, (Hagras et al., 2024, Section 4.2) use (A_3) . The less restrictive assumption (A_3) implies that higher qualifications are possible for the same function g_λ in the testing problem compared to the known qualifications in the literature of inverse problems and spectral regularized kernel ridge regression. For instance, consider the function $g_\lambda(x) = \frac{1}{x+\lambda}$ corresponding to Tikhonov regularization. In this case, the stronger condition used in literature $\sup_{x \in \Gamma} |1 - xg_\lambda(x)|x^{2\varphi} \leq C_3\lambda^{2\varphi}$ is satisfied only for $\varphi \in (0, \frac{1}{2}]$. However, (A_3) holds for any $\varphi > 0$, indicating infinite qualification with no saturation at $\varphi = \frac{1}{2}$ with $B_3 = \frac{1}{2}$ and $C_3 = 1$.

Define $\Sigma_{0,\lambda} := \Sigma_0 + \lambda \mathbf{I}$,

$$\mathcal{N}_1(\lambda) := \text{Tr}(\Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}), \quad \text{and} \quad \mathcal{N}_2(\lambda) := \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})},$$

which capture the intrinsic dimensionality (or degrees of freedom) of \mathcal{H} , with $\mathcal{N}_1(\lambda)$ being quite heavily used in the analysis of kernel ridge regression (e.g., Caponnetto and Vito 2007). Based on these preliminaries, in the following, we present an Oracle goodness-of-fit test.

4.1 Oracle Test

Using the samples $(X_i)_{i=1}^n$, a U -statistic estimator of η_λ defined in (9) can be written as

$$\hat{\eta}_\lambda = \frac{1}{n(n-1)} \sum_{i \neq j} \left\langle g_\lambda^{1/2}(\Sigma_0)(K(\cdot, X_i) - \mu_0), g_\lambda^{1/2}(\Sigma_0)(K(\cdot, X_j) - \mu_0) \right\rangle_{\mathcal{H}},$$

which when $\mu_0 = 0$ and $g_\lambda(x) = (x + \lambda)^{-1}$ reduces to the moderated MMD statistic proposed in Balasubramanian et al. (2021). The following result provides an α -level test based on $\hat{\eta}_\lambda$.

Theorem 3 (Critical region–Oracle) *Let $n \geq 2$. Suppose (A_0) – (A_2) hold. Then for any $\alpha > 0$ and any $\lambda > 0$,*

$$P_{H_0} \{ \hat{\eta}_\lambda \geq \gamma \} \leq \alpha,$$

where $\gamma = \frac{2(C_1 + C_2)\mathcal{N}_2(\lambda)}{n\sqrt{\alpha}}$.

Unfortunately, the test is not practical as the critical value, γ , and the test statistic depend on the eigenvalues and eigenfunctions of Σ_0 , which are not easy to compute for many (K, P_0)

pairs. Therefore, we call the above test the *Oracle test*. The following result analyzes the power of the Oracle test and presents sufficient conditions on the separation boundary to achieve the desired power.

Theorem 4 (Separation boundary–Oracle) *Suppose (A_0) – (A_3) . Let*

$$\sup_{P \in \mathcal{P}} \left\| \mathcal{T}^{-\theta} u \right\|_{L^2(P_0)} < \infty,$$

$\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} \geq \lambda = d_\theta \Delta_n^{\frac{1}{2\theta}}$, $d_\theta > 0$, where d_θ is a constant that depends on θ . For any $0 < \delta \leq 1$, if Δ_n satisfies

$$\frac{\Delta_n^{\frac{2\tilde{\theta}+1}{2\theta}}}{\mathcal{N}_2 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)} \gtrsim \frac{d_\theta^{-1} \delta^{-2}}{n^2}, \quad \frac{\Delta_n}{\mathcal{N}_2 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)} \gtrsim \frac{(\alpha^{-1/2} + \delta^{-1})}{n}.$$

then

$$\inf_{P \in \mathcal{P}} P_{H_1} \{ \hat{\eta}_\lambda \geq \gamma \} \geq 1 - \delta, \quad (10)$$

where $\gamma = \frac{2(C_1 + C_2)\mathcal{N}_2(\lambda)}{n\sqrt{\alpha}}$, and $\tilde{\theta} = \min(\theta, \xi)$. Furthermore, suppose $C := \sup_i \|\phi_i\|_\infty < \infty$. Then (10) holds when the above conditions on Δ_n are replaced by

$$\frac{\Delta_n}{\mathcal{N}_1 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)} \gtrsim \frac{\delta^{-2}}{n^2}, \quad \frac{\Delta_n}{\mathcal{N}_2 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)} \gtrsim \frac{(\alpha^{-1/2} + \delta^{-1})}{n}.$$

The following corollaries to Theorem 4 investigate the separation boundary of the test under the polynomial and exponential decay condition on the eigenvalues of Σ_0 .

Corollary 5 (Polynomial decay–Oracle) *Suppose $\lambda_i \asymp i^{-\beta}$, $\beta > 1$. Then for any $\delta > 0$,*

$$\inf_{P \in \mathcal{P}} P_{H_1} \{ \hat{\eta}_\lambda \geq \gamma \} \geq 1 - \delta,$$

when

$$\Delta_n = \begin{cases} c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\theta\beta+1}}, & \tilde{\theta} > \frac{1}{2} - \frac{1}{4\beta} \\ c(\alpha, \delta) n^{-\frac{8\tilde{\theta}\beta}{4\theta\beta+2\beta+1}}, & \tilde{\theta} \leq \frac{1}{2} - \frac{1}{4\beta} \end{cases},$$

with $c(\alpha, \delta) \gtrsim (\alpha^{-1/2} + \delta^{-2})$. Furthermore, if $\sup_i \|\phi_i\|_\infty < \infty$, then

$$\Delta_n = c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\theta\beta+1}}.$$

Corollary 6 (Exponential decay–Oracle) *Suppose $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$. Then for any $\delta > 0$, there exists $k_{\alpha, \delta}$ such that for all $n \geq k_{\alpha, \delta}$,*

$$\inf_{P \in \mathcal{P}} P_{H_1} \{ \hat{\eta}_\lambda \geq \gamma \} \geq 1 - \delta,$$

when

$$\Delta_n = \begin{cases} c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n}, & \tilde{\theta} > \frac{1}{2} \\ c(\alpha, \delta, \theta) \left(\frac{\sqrt{\log n}}{n} \right)^{\frac{4\tilde{\theta}}{2\theta+1}}, & \tilde{\theta} \leq \frac{1}{2} \end{cases},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\theta}}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$. Furthermore, if $\sup_i \|\phi_i\|_\infty < \infty$, then

$$\Delta_n = c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\theta}}, \frac{1}{2\theta}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$.

Remark 3 (i) Observe that larger qualification ξ (defined in Assumption (A₃) in Section 4) corresponds to a smaller separation boundary. Therefore, it is important to work with regularizers with infinite qualification, such as Tikhonov and Showalter. It has to be noted that the Tikhonov regularizer has infinite qualification as per (A₃) but has a qualification of $\frac{1}{2}$ w.r.t. the stronger version of (A₃).

(ii) Suppose g_λ has infinite qualification, $\xi = \infty$, then $\tilde{\theta} = \theta$. Comparing Corollary 5 (resp. Corollary 6) and Theorem 2 shows that the spectral regularized test based on $\hat{\eta}_\lambda$ is minimax optimal w.r.t. \mathcal{P} in the ranges of θ as given in Theorem 2 if the eigenvalues of \mathcal{T} decay polynomially (resp. exponentially). Outside these ranges of θ , the optimality of the test remains an open question since we do not have a minimax separation covering these ranges of θ .

(iii) Corollary 5 recovers the minimax separation rate in Balasubramanian et al. (2021) under the uniform boundedness condition but without assuming $\mu_0 = 0$. Furthermore, it also presents the separation rate for the regularized MMD test without assuming both the uniform boundedness condition and $\mu_0 = 0$, and shows a phase transition in the separation rate depending on the value of $\tilde{\theta}$.

4.2 Two-sample statistic

The Oracle test statistic requires the knowledge of μ_0 and Σ_0 for it to be computable. Though P_0 is known, Σ_0 and μ_0 are not known in closed form in general for many (K, P_0) pairs. To address this issue, in this section, we assume that P_0 is samplable, i.e., a set of i.i.d. samples from P_0 are available or can be generated. To this end, let us say $m + s$ i.i.d. samples are available from P_0 of which $(Y_i^0)_{i=1}^s \stackrel{i.i.d.}{\sim} P_0$ are used to estimate Σ_0 and $(X_i^0)_{i=1}^m \stackrel{i.i.d.}{\sim} P_0$ are used to estimate μ_0 , with $(Y_i^0)_{i=1}^s \stackrel{i.i.d.}{=} (X_i^0)_{i=1}^m$. Note that we do not use all $m + s$ samples to estimate both μ_0 and Σ_0 . Instead, we do sample splitting so that the estimators of Σ_0 and μ_0 are decoupled, which will turn out to be critical for the analysis. Based on this, η_λ can be estimated as a two-sample U -statistic (Hoeffding, 1992) as

$$\hat{\eta}_\lambda^{TS} := \frac{1}{n(n-1)} \frac{1}{m(m-1)} \sum_{1 \leq i \neq j \leq n} \sum_{1 \leq i' \neq j' \leq N} h(X_i, X_j, X_{i'}^0, X_{j'}^0), \quad (11)$$

where

$$h(X_i, X_j, X_{i'}^0, X_{j'}^0) := \left\langle g_\lambda^{1/2}(\hat{\Sigma}_0)(K(\cdot, X_i) - K(\cdot, X_{i'}^0)), g_\lambda^{1/2}(\hat{\Sigma}_0)(K(\cdot, X_j) - K(\cdot, X_{j'}^0)) \right\rangle_{\mathcal{H}},$$

and

$$\hat{\Sigma}_0 := \frac{1}{2s(s-1)} \sum_{i \neq j}^s (K(\cdot, Y_i^0) - K(\cdot, Y_j^0)) \otimes_{\mathcal{H}} (K(\cdot, Y_i^0) - K(\cdot, Y_j^0)),$$

is a one-sample U -statistic estimator of Σ_0 based on $(Y_i^0)_{i=1}^s$. Note that $\hat{\eta}_\lambda^{TS}$ is not exactly a U -statistic since it involves $\hat{\Sigma}_0$, but conditioned on $(Y_i^0)_{i=1}^s$, one can see that it is exactly a two-sample U -statistic. By expanding the inner product in h and writing (11) as

$$\begin{aligned} \hat{\eta}_\lambda^{TS} &= \frac{1}{n(n-1)} \sum_{i \neq j} \left\langle g_\lambda^{1/2}(\hat{\Sigma}_0) K(\cdot, X_i), g_\lambda^{1/2}(\hat{\Sigma}_0) K(\cdot, X_j) \right\rangle_{\mathcal{H}} \\ &\quad + \frac{1}{m(m-1)} \sum_{i \neq j} \left\langle g_\lambda^{1/2}(\hat{\Sigma}_0) K(\cdot, X_i^0), g_\lambda^{1/2}(\hat{\Sigma}_0) K(\cdot, X_j^0) \right\rangle_{\mathcal{H}} \\ &\quad - \frac{2}{nm} \sum_{i,j} \left\langle g_\lambda^{1/2}(\hat{\Sigma}_0) K(\cdot, X_i), g_\lambda^{1/2}(\hat{\Sigma}_0) K(\cdot, X_j^0) \right\rangle_{\mathcal{H}}, \end{aligned}$$

the following result shows that $\hat{\eta}_\lambda^{TS}$ can be computed only through matrix operations and by solving a finite-dimensional eigensystem.

Theorem 7 *Let $(\hat{\lambda}_i, \hat{\alpha}_i)_i$ be the eigensystem of $\frac{1}{s} \tilde{\mathbf{H}}_s^{1/2} K_s \tilde{\mathbf{H}}_s^{1/2}$ where $K_s := [K(Y_i^0, Y_j^0)]_{i,j \in [s]}$, $\mathbf{H}_s = \mathbf{I}_s - \frac{1}{s} \mathbf{1}_s \mathbf{1}_s^\top$, and $\tilde{\mathbf{H}}_s = \frac{s}{s-1} \mathbf{H}_s$. Define*

$$G := \sum_i \left(\frac{g_\lambda(\hat{\lambda}_i) - g_\lambda(0)}{\hat{\lambda}_i} \right) \hat{\alpha}_i \hat{\alpha}_i^\top.$$

Then

$$\hat{\eta}_\lambda^{TS} = \frac{1}{n(n-1)} (\textcircled{1} - \textcircled{2}) + \frac{1}{m(m-1)} (\textcircled{3} - \textcircled{4}) - \frac{2}{nm} \textcircled{5},$$

where

$$\begin{aligned} \textcircled{1} &= \mathbf{1}_n^\top \left(g_\lambda(0) K_n + \frac{1}{s} K_{ns} \tilde{\mathbf{H}}_s^{1/2} G \tilde{\mathbf{H}}_s^{1/2} K_{ns}^\top \right) \mathbf{1}_n, \\ \textcircled{2} &= \text{Tr} \left(g_\lambda(0) K_n + \frac{1}{s} K_{ns} \tilde{\mathbf{H}}_s^{1/2} G \tilde{\mathbf{H}}_s^{1/2} K_{ns}^\top \right), \\ \textcircled{3} &= \mathbf{1}_m^\top \left(g_\lambda(0) K_m + \frac{1}{s} K_{ms} \tilde{\mathbf{H}}_s^{1/2} G \tilde{\mathbf{H}}_s^{1/2} K_{ms}^\top \right) \mathbf{1}_m, \\ \textcircled{4} &= \text{Tr} \left(g_\lambda(0) K_m + \frac{1}{s} K_{ms} \tilde{\mathbf{H}}_s^{1/2} G \tilde{\mathbf{H}}_s^{1/2} K_{ms}^\top \right), \quad \text{and} \\ \textcircled{5} &= \mathbf{1}_m^\top \left(g_\lambda(0) K_{mn} + \frac{1}{s} K_{ms} \tilde{\mathbf{H}}_s^{1/2} G \tilde{\mathbf{H}}_s^{1/2} K_{ns}^\top \right) \mathbf{1}_n, \end{aligned}$$

with $K_n := [K(X_i, X_j)]_{i,j \in [n]}$, $K_m := [K(X_i^0, X_j^0)]_{i,j \in [m]}$, $K_{ns} := [K(X_i, Y_j^0)]_{i \in [n], j \in [s]}$, $K_{ms} := [K(X_i^0, Y_j^0)]_{i \in [m], j \in [s]}$, and $K_{mn} := [K(X_i^0, X_j)]_{i \in [m], j \in [n]}$.

Note that in the case of Tikhonov regularization, $G = \frac{-1}{\lambda} (\frac{1}{s} \tilde{\mathbf{H}}_s^{1/2} K_s \tilde{\mathbf{H}}_s^{1/2} + \lambda \mathbf{I}_s)^{-1}$. The complexity of computing $\hat{\eta}_\lambda^{TS}$ is given by $O(s^3 + m^2 + n^2 + ns^2 + ms^2)$. We would like to mention that since $\hat{\eta}_\lambda^{TS}$ is based on two sets of samples, a result very similar to Theorem 7 is presented in (Hagrass et al., 2024, Theorem 3) in the context of two-sample testing.

4.3 Spectral regularized concentration test (SRCT)

By applying Chebyshev inequality to $\hat{\eta}_\lambda^{TS}$ under H_0 , the following result provides an α -level test, which we refer to as SRCT. Define $\widehat{\mathcal{N}}_2(\lambda) := \left\| \hat{\Sigma}_{0,\lambda}^{-1/2} \hat{\Sigma}_0 \hat{\Sigma}_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}$.

Theorem 8 (Critical region–SRCT) *Let $n \geq 2$ and $m \geq 2$. Suppose (A_0) – (A_2) hold. Then for any $\alpha > 0$, $c_1 \geq 65$ and $\frac{4c_1\kappa}{s} \max\{\log \frac{96\kappa s}{\alpha}, \log \frac{12}{\alpha}\} \leq \lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$,*

$$P_{H_0} \{ \hat{\eta}_\lambda^{TS} \geq \gamma \} \leq \alpha,$$

where $\gamma = \frac{12(C_1+C_2)\widehat{\mathcal{N}}_2(\lambda)}{b_1\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right)$, $b_1 = \sqrt{\frac{4}{9} - \frac{16}{3\sqrt{3}c_1} - \frac{32}{9c_1}}$. Furthermore, if $C := \sup_i \|\phi_i\|_\infty < \infty$, the above bound holds for $4c_1C^2\mathcal{N}_1(\lambda) \log \frac{48\mathcal{N}_1(\lambda)}{\alpha} \leq s$.

Note that unlike in the Oracle test, the threshold γ and the test statistic $\hat{\eta}_\lambda^{TS}$ in the above result is completely data-driven and computable, with $\widehat{\mathcal{N}}_2(\lambda)$ being computed based on $(\hat{\lambda}_i)_i$ from Theorem 7. The following result provides sufficient conditions on the separation boundary to achieve the desired power.

Theorem 9 (Separation boundary–SRCT) *Suppose (A_0) – (A_4) and $m \geq n$. Let*

$$\sup_{P \in \mathcal{P}} \left\| \mathcal{T}^{-\theta} u \right\|_{L^2(P_0)} < \infty,$$

$\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} \geq \lambda = d_\theta \Delta_n^{\frac{1}{2\theta}}$, $d_\theta > 0$, where d_θ is a constant that depends on θ . For any $0 < \delta \leq 1$, if $s \geq 32c_1\kappa\lambda^{-1} \log(\max\{17920\kappa^2\lambda^{-1}, 6\}\delta^{-1})$ and Δ_n satisfies

$$\frac{\Delta_n^{\frac{2\theta+1}{2\theta}}}{\mathcal{N}_2 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)} \gtrsim \frac{d_\theta^{-1} \delta^{-2}}{n^2}, \quad \frac{\Delta_n}{\sqrt{\mathcal{N}_1 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)}} \gtrsim \frac{(\alpha^{-1/2} + \delta^{-1})}{n},$$

then

$$\inf_{P \in \mathcal{P}} P_{H_1} \{ \hat{\eta}_\lambda^{TS} \geq \gamma \} \geq 1 - 4\delta, \quad (12)$$

where $\gamma = \frac{12(C_1+C_2)\widehat{\mathcal{N}}_2(\lambda)}{b_1\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right)$, $b_1 = \sqrt{\frac{4}{9} - \frac{16}{3\sqrt{3}c_1} - \frac{32}{9c_1}}$, $c_1 \geq 65$ and $\tilde{\theta} = \min(\theta, \xi)$. Furthermore, suppose $C := \sup_i \|\phi_i\|_\infty < \infty$. Then (12) holds if $s \geq 32c_1C^2\mathcal{N}_1(\lambda) \log \frac{32\mathcal{N}_1(\lambda)}{\delta}$ and when the above conditions on Δ_n are replaced by

$$\frac{\Delta_n}{\mathcal{N}_1 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)} \gtrsim \frac{\delta^{-2}}{n^2}, \quad \frac{\Delta_n}{\sqrt{\mathcal{N}_1 \left(d_\theta \Delta_n^{\frac{1}{2\theta}} \right)}} \gtrsim \frac{(\alpha^{-1/2} + \delta^{-1})}{n}.$$

Remark 4 (i) Comparing the conditions on the separation boundary in Theorem 9 to those of Theorem 4, it is easy to verify that the claims in Corollaries 5 and 6 also hold for SRCT. Therefore, SRCT achieves minimax optimality in the same ranges of θ as the Oracle test.

(ii) In the case of polynomial decay, when $\tilde{\theta} > \frac{1}{2} - \frac{1}{4\beta}$, the condition on s —the number of samples needed to estimate the covariance operator Σ_0 —reduces to $s \gtrsim n^{\frac{2\beta}{4\theta\beta+1}} \log n$, which is of sub-linear order and is implied if $s \gtrsim n \log n$. When $\tilde{\theta} \leq \frac{1}{2} - \frac{1}{4\beta}$, the condition becomes $s \gtrsim n^{\frac{4\beta}{4\theta\beta+2\beta+1}} \log n$ which is implied for any θ and β if $s \gtrsim n^2 \log n$. Furthermore, under uniform boundedness, the condition on s becomes $s \gtrsim n^{\frac{2}{4\theta\beta+1}} \log n$ which is of sublinear order for $\theta > \frac{1}{4\beta}$. In case of exponential decay, for $\theta > \frac{1}{2}$, the condition is $s \gtrsim n^{\frac{1}{2\theta}} (\log n)^{1-\frac{1}{4\theta}}$, which is implied for any $\theta > \frac{1}{2}$, if $s \gtrsim n\sqrt{\log n}$. For $\theta < \frac{1}{2}$, the condition is $s \gtrsim n^{\frac{2}{2\theta+1}} (\log n)^{\frac{2\theta}{2\theta+1}}$ which is implied if $s \gtrsim n^2$. Furthermore, if $\sup_i \|\phi_i\|_\infty < \infty$ holds, then the condition is $s \gtrsim (\log n)(\log \log n)$.

4.4 Spectral regularized permutation test (SRPT)

Instead of using a concentration inequality-based test threshold as in SRCT, in this section, we study the permutation approach to compute the test threshold (Lehmann and Romano, 2006; Pesarin and Salmaso, 2010; Kim et al., 2022). We refer to the resulting test as SRPT. We show that SRPT achieves a minimax optimal separation boundary with a better constant compared to that of SRCT.

Recall that our test statistic defined in Section 4.2 involves three sets of independent samples, $(X_i)_{i=1}^n \stackrel{i.i.d.}{\sim} P$, $(X_j^0)_{j=1}^m \stackrel{i.i.d.}{\sim} P_0$, $(Y_i^0)_{i=1}^s \stackrel{i.i.d.}{\sim} P_0$. Define $(U_i)_{i=1}^n := (X_i)_{i=1}^n$, and $(U_{n+j})_{j=1}^m := (X_j^0)_{j=1}^m$. Let Π_{n+m} be the set of all possible permutations of $\{1, \dots, n+m\}$ with $\pi \in \Pi_{n+m}$ being a randomly selected permutation from the D possible permutations, where $D := |\Pi_{n+m}| = (n+m)!$. Define $(X_i^\pi)_{i=1}^n := (U_{\pi(i)})_{i=1}^n$ and $(X_j^{0\pi})_{j=1}^m := (U_{\pi(n+j)})_{j=1}^m$. Let $\hat{\eta}_\lambda^\pi := \hat{\eta}_\lambda^{TS}(X^\pi, X^{0\pi}, Y^0)$ be the statistic based on the permuted samples, and $(\pi^i)_{i=1}^B$ be B randomly selected permutations from Π_{n+m} . For simplicity, define $\hat{\eta}_\lambda^i := \hat{\eta}_\lambda^{\pi^i}$ to represent the statistic based on permuted samples w.r.t. the random permutation π^i . Given the samples $(X_i)_{i=1}^n$, $(X_j^0)_{j=1}^m$ and $(Y_i^0)_{i=1}^s$, define

$$F_\lambda(x) := \frac{1}{D} \sum_{\pi \in \Pi_{n+m}} \mathbb{1}(\hat{\eta}_\lambda^\pi \leq x)$$

to be the permutation distribution function, and define

$$q_{1-\alpha}^\lambda := \inf\{q \in \mathbb{R} : F_\lambda(q) \geq 1 - \alpha\}.$$

Furthermore, we define the empirical permutation distribution function based on B random permutations as

$$\hat{F}_\lambda^B(x) := \frac{1}{B+1} \sum_{i=0}^B \mathbb{1}(\hat{\eta}_\lambda^i \leq x),$$

where $\hat{\eta}_\lambda^0 = \hat{\eta}_\lambda^{TS}$ and define

$$\hat{q}_{1-\alpha}^{B,\lambda} := \inf\{q \in \mathbb{R} : \hat{F}_\lambda^B(q) \geq 1 - \alpha\}.$$

Based on these notations, the following result presents an α -level test with a completely data-driven critical level.

Theorem 10 (Critical region–SRPT) For any $0 < \alpha \leq 1$, $P_{H_0}\{\hat{\eta}_\lambda^{TS} \geq \hat{q}_{1-\alpha}^{B,\lambda}\} \leq \alpha$.

It is well known that the permutation approach exactly controls the type-I error. This follows from the exchangeability of samples under H_0 and the definition of $q_{1-\alpha}^\lambda$. Next, similar to Theorem 9, the following result provides general conditions under which the power can be controlled.

Theorem 11 (Separation boundary–SRPT) Suppose (A_0) – (A_4) hold. Let $m \geq n$,

$$\sup_{P \in \mathcal{P}} \|\mathcal{T}^{-\theta} u\|_{L^2(P_0)} < \infty,$$

$\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} \geq \lambda = d_\theta \Delta_n^{\frac{1}{2\theta}}$, for some $d_\theta > 0$, where d_θ is a constant that depends on θ . For any $0 < \delta \leq 1$, if $n \geq d_3 \delta^{-1/2} \log \frac{1}{\alpha}$ for some $d_3 > 0$, $B \geq \frac{3}{\alpha^2} (\log 2\delta^{-1} + \alpha(1-\alpha))$, $s \geq 280\kappa\lambda^{-1} \log(17920\kappa^2\lambda^{-1}\delta^{-1})$ and Δ_n satisfies

$$\frac{\Delta_n^{\frac{2\tilde{\theta}+1}{2\theta}}}{\mathcal{N}_2\left(d_\theta \Delta_n^{\frac{1}{2\theta}}\right)} \gtrsim \frac{d_\theta^{-1}(\delta^{-1} \log(1/\tilde{\alpha}))^2}{n^2}, \quad \frac{\Delta_n}{\mathcal{N}_2\left(d_\theta \Delta_n^{\frac{1}{2\theta}}\right)} \gtrsim \frac{\delta^{-1} \log(1/\tilde{\alpha})}{n},$$

then

$$\inf_{P \in \mathcal{P}} P_{H_1} \left\{ \hat{\eta}_\lambda^{TS} \geq \hat{q}_{1-\alpha}^{B,\lambda} \right\} \geq 1 - 5\delta, \quad (13)$$

where $\tilde{\theta} = \min(\theta, \xi)$. Furthermore, suppose $C := \sup_i \|\phi_i\|_\infty < \infty$. Then (13) holds if $s \geq 136C^2\mathcal{N}_1(\lambda) \log \frac{32\mathcal{N}_1(\lambda)}{\delta}$ and when the above conditions on Δ_n are replaced by

$$\frac{\Delta_n}{\mathcal{N}_1\left(d_\theta \Delta_n^{\frac{1}{2\theta}}\right)} \gtrsim \frac{(\delta^{-1} \log(1/\tilde{\alpha}))^2}{n^2}, \quad \frac{\Delta_n}{\mathcal{N}_2\left(d_\theta \Delta_n^{\frac{1}{2\theta}}\right)} \gtrsim \frac{\delta^{-1} \log(1/\tilde{\alpha})}{n}.$$

Corollary 12 Suppose $\lambda_i \asymp i^{-\beta}$, $\beta > 1$. Then for any $\delta > 0$,

$$\inf_{P \in \mathcal{P}} P_{H_1} \left\{ \hat{\eta}_\lambda^{TS} \geq \hat{q}_{1-\alpha}^{B,\lambda} \right\} \geq 1 - 5\delta,$$

when

$$\Delta_n = \begin{cases} c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\theta\beta+1}}, & \tilde{\theta} > \frac{1}{2} - \frac{1}{4\beta} \\ c(\alpha, \delta) n^{\frac{-8\tilde{\theta}\beta}{4\theta\beta+2\beta+1}}, & \tilde{\theta} \leq \frac{1}{2} - \frac{1}{4\beta} \end{cases},$$

with $c(\alpha, \delta) \gtrsim \delta^{-2}(\log \frac{1}{\alpha})^2$. Furthermore, if $\sup_i \|\phi_i\|_\infty < \infty$, then

$$\Delta_n = c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\theta\beta+1}}.$$

Corollary 13 Suppose $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$. Then for any $\delta > 0$, there exists $k_{\alpha, \delta}$ such that for all $n \geq k_{\alpha, \delta}$,

$$\inf_{(P, Q) \in \mathcal{P}} P_{H_1} \left\{ \hat{\eta}_\lambda^{TS} \geq \hat{q}_{1-\alpha}^{B,\lambda} \right\} \geq 1 - 2\delta,$$

when

$$\Delta_n = \begin{cases} c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n}, & \tilde{\theta} > \frac{1}{2} \\ c(\alpha, \delta, \theta) \left(\frac{\sqrt{\log n}}{n} \right)^{\frac{4\tilde{\theta}}{2\tilde{\theta}+1}}, & \tilde{\theta} \leq \frac{1}{2} \end{cases},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, 1 \right\} \delta^{-2} (\log \frac{1}{\alpha})^2$. Furthermore, if $\sup_i \|\phi_i\|_\infty < \infty$, then

$$\Delta_n = c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, \frac{1}{2\tilde{\theta}}, 1 \right\} \delta^{-2} (\log \frac{1}{\alpha})^2$.

The above results demonstrate the minimax optimality w.r.t. \mathcal{P} of the permutation-based test constructed in Theorem 10. Since the conditions on s in Theorem 11 match those of Theorem 9, the discussion in Remark 4(ii) also applies for SRPT.

Remark 5 Recently, Hagrass et al. (2024) proposed a spectral regularized two-sample test (SR2T) where the test statistic has a close resemblance to $\hat{\eta}_\lambda^{TS}$. Since we are solving a goodness-of-fit test question as a two-sample test, one could simply address it using SR2T, and therefore one may wonder about the need for the proposal of SRCT and SRPT, given their similarity to SR2T. While this is a valid question, comparing Corollaries 12 and 13 to that of (Hagrass et al., 2024, Corollaries 6, 7), we observe that while all these tests enjoy minimax separation rates over the same range of $\tilde{\theta}$, for the range of $\tilde{\theta}$ where the minimaxity of separation rate is not established, the proposed tests have faster convergence rate than that of SR2T, thereby demonstrating the advantage of the proposed tests over SR2T (see Section 1.2 for details).

4.5 Adaptation

In the previous sections, we have discussed two ways of constructing a test based on the statistic $\hat{\eta}_\lambda^{TS}$. In both these tests, the optimal λ to achieve the minimax separation boundary depends on unknown θ and β . In this section, we construct a test based on the union (aggregation) of multiple tests constructed for different values of λ taking values in a finite set, Λ . It turns out that the resultant test is guaranteed to be minimax optimal (up to log factors) for a wide range of θ (and β in the case of polynomially decaying eigenvalues). The aggregation method is quite classical and we employ the technique as used in Hagrass et al. (2024).

Define $\Lambda := \{\lambda_L, 2\lambda_L, \dots, \lambda_U\}$, where $\lambda_U = 2^b \lambda_L$, for $b \in \mathbb{N}$ so that $|\Lambda| = b + 1 = 1 + \log_2 \frac{\lambda_U}{\lambda_L}$, where $|\Lambda|$ is the cardinality of Λ . Let λ^* be the optimal λ that yields minimax optimality. The key idea is to choose λ_L and λ_U such that there is an element in Λ that is close to λ^* for any θ (and β in the case of polynomially decaying eigenvalues). Define $v^* := \sup\{x \in \Lambda : x \leq \lambda^*\}$. Then it is easy to verify that $v^* \asymp \lambda^*$, i.e., v^* is also an optimal choice for λ that belongs to Λ , since for $\lambda_L \leq \lambda^* \leq \lambda_U$, we have $\frac{\lambda^*}{2} \leq v^* \leq \lambda^*$. Motivated by this, in Theorems 14 and 16, we construct α -level tests that are adaptive versions of SRCT and SRPT, based on the union of corresponding tests over $\lambda \in \Lambda$ that rejects H_0 if one of the tests rejects H_0 . The separation boundary of these tests are analyzed in Theorems 15 and 17 under the polynomial and exponential decay rates of the eigenvalues of \mathcal{T} . These

results show that the adaptive versions achieve the same performance (up to log factors) as that of the Oracle test, i.e., minimax optimal w.r.t. \mathcal{P} over the range of θ mentioned in Theorem 2, without requiring the knowledge of λ^* . In contrast, (Balasubramanian et al., 2021, Theorem 5) considers an adaptive and asymptotic version of the Oracle test under $\mu_0 = 0$ and the uniform boundedness condition, where it only adapts over θ assuming β is known.

The following results relate to the adaptive version of SRCT.

Theorem 14 (Critical region–adaptation–SRCT) *Suppose (A_0) – (A_2) . Then for any $\alpha > 0$, $\frac{32c_1\kappa}{s} \max\{\log \frac{96\kappa s}{\tilde{\alpha}}, \log \frac{12}{\tilde{\alpha}}\} \leq \lambda_L \leq \lambda_U \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$, where $\tilde{\alpha} = \frac{\alpha}{|\Lambda|}$, $c_1 \geq 65$.*

$$P_{H_0} \left\{ \sup_{\lambda \in \Lambda} \frac{\hat{\eta}_\lambda^{TS}}{\hat{\mathcal{N}}_2(\lambda)} \geq \gamma \right\} \leq \alpha,$$

where $\gamma = \frac{12(C_1+C_2)}{b_1\sqrt{\tilde{\alpha}}} \left(\frac{1}{n} + \frac{1}{m} \right)$, $b_1 = \sqrt{\frac{4}{9} - \frac{16}{3\sqrt{3c_1}} - \frac{32}{9c_1}}$. Furthermore if $C := \sup_i \|\phi_i\|_\infty < \infty$, the above bound holds for $4c_1C^2\mathcal{N}_1(\lambda_L) \log \frac{48\mathcal{N}_1(\lambda_L)}{\delta} \leq s$.

Theorem 15 (Separation boundary–adaptation–SRCT) *Suppose (A_0) – (A_4) hold, $\tilde{\theta} = \min(\theta, \xi)$, $\tilde{\xi} = \max(\xi, \frac{1}{4})$, $\sup_{\theta > 0} \sup_{P \in \mathcal{P}} \|\mathcal{T}^{-\theta}u\|_{L^2(P_0)} < \infty$, $\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} \geq \lambda_U$, $\theta_l > 0$, $m \geq n$, and $0 < \alpha \leq 1$. Then for any $\delta > 0$ and γ defined as in Theorem 14,*

$$\inf_{\theta \geq \theta_l} \inf_{P \in \mathcal{P}} P_{H_1} \left\{ \sup_{\lambda \in \Lambda} \frac{\hat{\eta}_\lambda^{TS}}{\hat{\mathcal{N}}_2(\lambda)} \geq \gamma \right\} \geq 1 - 4\delta,$$

provided one of the following cases holds:

(i) $\lambda_i \asymp i^{-\beta}$, $1 < \beta \leq \beta_U$, $\lambda_L = r_1 n^{\frac{-4\beta_U}{1+2\beta_U}}$, $\lambda_U = r_2 \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-2}{4\tilde{\xi}+1}}$ for $r_1, r_2 > 0$, $s \geq 32c_1\kappa\lambda_L^{-1} \log(\max\{17920\kappa^2\lambda_L^{-1}, 6\}\delta^{-1})$, and

$$\Delta_n = c(\alpha, \delta) \max \left\{ n^{\frac{-8\tilde{\theta}\beta}{1+2\beta+4\tilde{\theta}\beta}}, \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}} \right\}.$$

Furthermore if $C := \sup_i \|\phi_i\|_\infty^2 < \infty$, then the above conditions on λ_L , λ_U and s can be replaced by $\lambda_L = r_3 \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-2\beta_U}{4\theta_l\beta_U+1}}$, $\lambda_U = r_4 \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-2}{4\tilde{\xi}+1}}$ for $r_3, r_4 > 0$, $s \geq 32c_1C^2\mathcal{N}_1(\lambda_L) \log \frac{32\mathcal{N}_1(\lambda_L)}{\delta}$, and

$$\Delta_n = c(\alpha, \delta) \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}},$$

where $c(\alpha, \delta) \gtrsim (\alpha^{-1/2} + \delta^{-2})$.

(ii) $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$, $\lambda_L = r_5 \left(\frac{n}{\sqrt{\log n}} \right)^{-2}$, $\lambda_U = r_6 \left(\frac{n}{\log n} \right)^{-1/2\tilde{\xi}}$ for some $r_5, r_6 > 0$, $s \geq 32c_1\kappa\lambda_L^{-1} \log(\max\{17920\kappa^2\lambda_L^{-1}, 6\}\delta^{-1})$, and

$$\Delta_n \gtrsim c(\alpha, \delta, \theta) \max \left\{ \left(\frac{n}{\log n} \right)^{-1}, \left(\frac{n}{\sqrt{\log n}} \right)^{-\frac{4\tilde{\theta}}{2\tilde{\theta}+1}} \right\},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$. Furthermore if $C := \sup_i \|\phi_i\|_\infty^2 < \infty$, then the above conditions on λ_L , λ_U and s can be replaced by $\lambda_L = r_7 \left(\frac{n}{\log \log n} \right)^{-1/2\theta_l}$, $\lambda_U = r_8 \left(\frac{n}{\log \log n} \right)^{-1/2\xi}$ for some $r_7, r_8 > 0$, $s \geq 32c_1C^2\mathcal{N}_1(\lambda_L) \log \frac{32\mathcal{N}_1(\lambda_L)}{\delta}$ and

$$\Delta_n = c(\alpha, \delta, \theta) \frac{\log n}{n},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, \frac{1}{2\tilde{\theta}}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$.

The following results handle the adapted version of SRPT, which show that the adapted test is minimax optimal w.r.t. \mathcal{P} up to a $\log \log n$ factor.

Theorem 16 (Critical region–adaptation–SRPT) *For any $0 < \alpha \leq 1$,*

$$P_{H_0} \left\{ \bigcup_{\lambda \in \Lambda} \hat{\eta}_\lambda^{TS} \geq \hat{q}_{1-\frac{\alpha}{|\Lambda|}}^{B,\lambda} \right\} \leq \alpha.$$

Theorem 17 (Separation boundary–adaptation–SRPT) *Suppose (A_0) – (A_4) hold, $\tilde{\theta} = \min(\theta, \xi)$, $\tilde{\xi} = \max(\xi, \frac{1}{4})$, $\sup_{\theta > 0} \sup_{P \in \mathcal{P}} \|\mathcal{T}^{-\theta} u\|_{L^2(P_0)} < \infty$, and $m \geq n$. Then for any $\delta > 0$, $B \geq \frac{3}{\tilde{\alpha}^2} (\log 2\delta^{-1} + \tilde{\alpha}(1 - \tilde{\alpha}))$, $0 < \alpha \leq e^{-1}$, $\tilde{\alpha} = \frac{\alpha}{|\Lambda|}$, $\theta_l > 0$, we have*

$$\inf_{\theta > \theta_l} \inf_{P \in \mathcal{P}} P_{H_1} \left\{ \bigcup_{\lambda \in \Lambda} \hat{\eta}_\lambda^{TS} \geq \hat{q}_{1-\frac{\alpha}{|\Lambda|}}^{B,\lambda} \right\} \geq 1 - 5\delta,$$

provided one of the following cases holds:

(i) $\lambda_i \asymp i^{-\beta}$, $1 < \beta \leq \beta_U$, $\lambda_L = r_1 \left(\frac{n}{\log \log n} \right)^{\frac{-4\beta_U}{1+2\beta_U}}$, $\lambda_U = r_2 \left(\frac{n}{\log \log n} \right)^{\frac{-2}{4\tilde{\xi}+1}}$ for $r_1, r_2 > 0$, $s \geq 32c_1\kappa\lambda_L^{-1} \log(\max\{17920\kappa^2\lambda_L^{-1}, 6\}\delta^{-1})$, and

$$\Delta_n = c(\alpha, \delta) \max \left\{ \left(\frac{n}{\log \log n} \right)^{\frac{-8\tilde{\theta}\beta}{1+2\beta+4\tilde{\theta}\beta}}, \left(\frac{n}{\log \log n} \right)^{\frac{-4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}} \right\}.$$

Furthermore if $C := \sup_i \|\phi_i\|_\infty^2 < \infty$, then the above conditions on λ_L , λ_U and s can be replaced by $\lambda_L = r_3 \left(\frac{n}{\log \log n} \right)^{\frac{-2\beta_U}{4\theta_l\beta_U+1}}$, $\lambda_U = r_4 \left(\frac{n}{\log \log n} \right)^{\frac{-2}{4\tilde{\xi}+1}}$ for $r_3, r_4 > 0$,

$s \geq 32c_1C^2\mathcal{N}_1(\lambda_L) \log \frac{32\mathcal{N}_1(\lambda_L)}{\delta}$, and

$$\Delta_n = c(\alpha, \delta) \left(\frac{n}{\log \log n} \right)^{\frac{-4\tilde{\theta}\beta}{4\theta\beta+1}},$$

where $c(\alpha, \delta) \gtrsim \delta^{-2}(\log \frac{1}{\alpha})^2$.

(ii) $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$, $\lambda_L = r_5 \left(\frac{n}{\sqrt{\log n \log \log n}} \right)^{-2}$, $\lambda_U = r_6 \left(\frac{n}{\sqrt{\log n \log \log n}} \right)^{-1/2\tilde{\xi}}$ for some $r_5, r_6 > 0$, $s \geq 32c_1\kappa\lambda_L^{-1} \log(\max\{17920\kappa^2\lambda_L^{-1}, 6\}\delta^{-1})$, and

$$\Delta_n \gtrsim c(\alpha, \delta, \theta) \max \left\{ \left(\frac{n}{\sqrt{\log n \log \log n}} \right)^{-1}, \left(\frac{n}{\sqrt{\log n \log \log n}} \right)^{-\frac{4\tilde{\theta}}{2\theta+1}} \right\},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, 1 \right\} \delta^{-2}(\log \frac{1}{\alpha})^2$. Furthermore if $C := \sup_i \|\phi_i\|_\infty^2 < \infty$, then the above conditions on λ_L , λ_U and s can be replaced by

$$\lambda_L = r_7 \left(\frac{n}{\sqrt{\log n \log \log n}} \right)^{-1/2\theta_l}, \quad \lambda_U = r_8 \left(\frac{n}{\sqrt{\log n \log \log n}} \right)^{-1/2\xi}$$

for some $r_7, r_8 > 0$, $s \geq 4c_1C^2\mathcal{N}_1(\lambda_L) \log \frac{8\mathcal{N}_1(\lambda_L)}{\delta}$ and

$$\Delta_n = c(\alpha, \delta, \theta) \frac{\sqrt{\log n \log \log n}}{n},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, \frac{1}{2\tilde{\theta}}, 1 \right\} \delta^{-2}(\log \frac{1}{\alpha})^2$.

The discussion so far has dealt with adapting to unknown θ and β associated with a given kernel. The natural question is how to choose the kernel, for example, suppose the kernel is a Gaussian kernel, then what is the right choice of bandwidth? This is an important question because it is not easy to characterize the class of kernels that satisfy the range space and eigenvalue decay conditions for a given P_0 . This question can be addressed by starting with a family of kernels, \mathcal{K} and constructing an adaptive test by taking the union of tests jointly over $\lambda \in \Lambda$ and $K \in \mathcal{K}$, so that the resulting test is jointly adaptive over λ and the kernel class \mathcal{K} . This idea has been explored recently in (Hagrass et al., 2024, Section 4.5) to construct a minimax optimal (up to a log factor) test that is jointly adaptive to both λ and \mathcal{K} (\mathcal{K} is assumed to be finite). Since the same idea can be explored for SRCT and SRPT to create kernel adaptive tests that yield results that are similar to those of Theorems 14–17 along with their proofs, we skip the details here and encourage the reader to refer to Hagrass et al. (2024).

5 Experiments

In this section, we investigate the empirical performance of the proposed regularized goodness-of-fit tests, SRCT and SRPT with adaptation to λ and the kernel. Note that SRCT and

SRPT are approximations to the Oracle test, since the latter is not easy to compute in general. In Section 5.1, using a periodic spline kernel, we compare the performance of SRPT to the moderated MMD (M3D) test (i.e., Oracle test) of Balasubramanian et al. (2021), which requires the knowledge of the eigenvalues and eigenfunctions of the kernel with respect to P_0 . Since SRCT and SRPT can be treated as two-sample tests, in Sections 5.2–5.4, we compare their performance to other popular two-sample tests in the literature such as adaptive MMD test (MMDAgg) (Schrab et al., 2021), Energy test (Szekely and Rizzo, 2004), Kolmogorov-Smirnov (KS) test (Puritz et al., 2022; Fasano and Franceschini, 1987) and the spectral regularized two sample test (SR2T) proposed in Hagrass et al. (2024) with Showalter regularization. For these experiments we used Gaussian kernel, defined as $K(x, y) = \exp\left(-\frac{\|x-y\|_2^2}{2h}\right)$, where h is the bandwidth. For our tests, we construct adaptive versions by taking the union of tests jointly over $\lambda \in \Lambda$ and $h \in W$. Let $\hat{\eta}_{\lambda, h}$ be the test statistic based on λ and bandwidth h . We reject H_0 if $\hat{\eta}_{\lambda, h} \geq \hat{q}_{1-\frac{\alpha}{|\Lambda||W|}}^{B, \lambda, h}$ for any $(\lambda, h) \in \Lambda \times W$. We performed such a test for $\Lambda := \{\lambda_L, 2\lambda_L, \dots, \lambda_U\}$, and $W := \{w_L h_m, 2w_L h_m, \dots, w_U h_m\}$, where $h_m := \text{median}\{\|q - q'\|_2^2 : q, q' \in X \cup X^0\}$, $X := (X_1, \dots, X_n)$ and $X^0 := (X_1^0, \dots, X_m^0)$. In our experiments, we set $\lambda_L = 10^{-6}$, $\lambda_U = 5$, $w_L = 0.01$ and $w_U = 100$. All tests are repeated 200 times and the average power is reported. For all experiments, we set $\alpha = 0.05$. For the tests SRPT and SR2T, we set the number of permutations to $B = 60$ and the number of samples used to estimate the covariance operator to $s = 100$.

5.1 Periodic spline kernel & perturbed uniform distribution: Oracle test

In this section, we compare the power of SRPT to that of M3D Balasubramanian et al. (2021). To be able to compute the M3D test, we use the periodic spline kernel, defined as $K(x, y) = \frac{(-1)^{r-1}}{(2r)!} B_{2r}([x-y])$, where B_r is the Bernoulli polynomial and $[t]$ is the fractional part of t . We set $r = 1$ and consider testing uniformity on the unit interval $\mathcal{X} = [0, 1]$. Under this setting, the eigenvalues and eigenfunctions of K are known explicitly (see Balasubramanian et al., 2021, Section 5 for details) so that the test statistic can be exactly computed. We examine testing the null hypothesis of uniform distribution against perturbed uniform distribution (see Hagrass et al., 2024, Section 5.1.1 for details), where the perturbed uniform distribution is indexed by a parameter P that characterizes the degree of perturbation. The larger the P is, the associated distribution is closer to uniform, implying that it becomes more difficult to distinguish between the null and the alternative. Figure 1 shows the power of SRPT in comparison to M3D for varying sample sizes n . SRPT($m = n$) and SRPT($m = 3n$) refer to our proposed permutation test while setting $m = n$, and $m = 3n$ respectively (recall that m is the number of samples from P_0 used to estimate the mean function μ_{P_0}). We can observe that SRPT with enough samples from P_0 can yield power almost as good as M3D (Oracle test), while not requiring the exact eigenvalues and eigenfunctions of \mathcal{T} . We also observe that s (the number of samples used to estimate the covariance operator, Σ_0) does not have much significance on the power and the choice of $s = 100$ seems to be good enough to accurately estimate Σ_0 .

Other than this experiment, unfortunately, we are not able to replicate any other experiment from (Balasubramanian et al., 2021) since no details about the parameter settings of the null and the alternative distributions are provided (i.e., if P_0 is normal, its mean

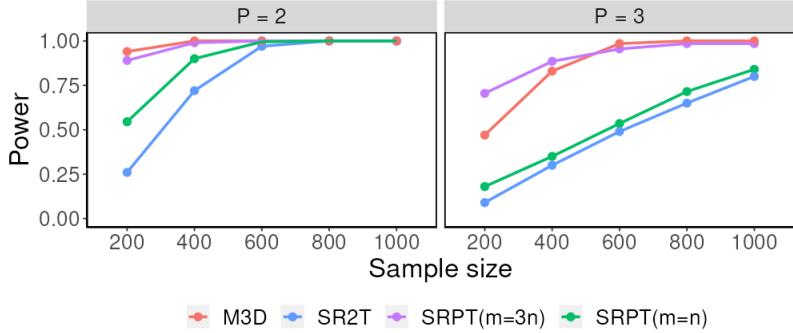


Figure 1: Oracle test (M3D) and SRPT to test for uniformity using periodic spline kernel on $[0, 1]$. P denotes the degree of perturbation where large P makes the alternative distribution (i.e., the perturbed uniform distribution) to be closer to the null (uniform distribution).

and variance are not mentioned). Moreover, the exact details about the computation of the eigenvalues and eigenfunctions of \mathcal{T} are not provided.

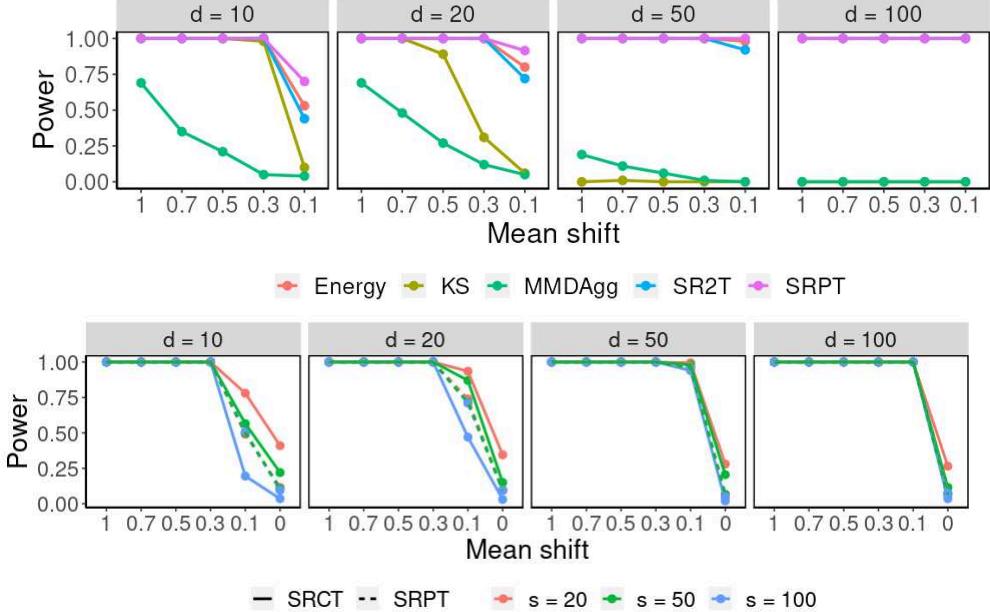
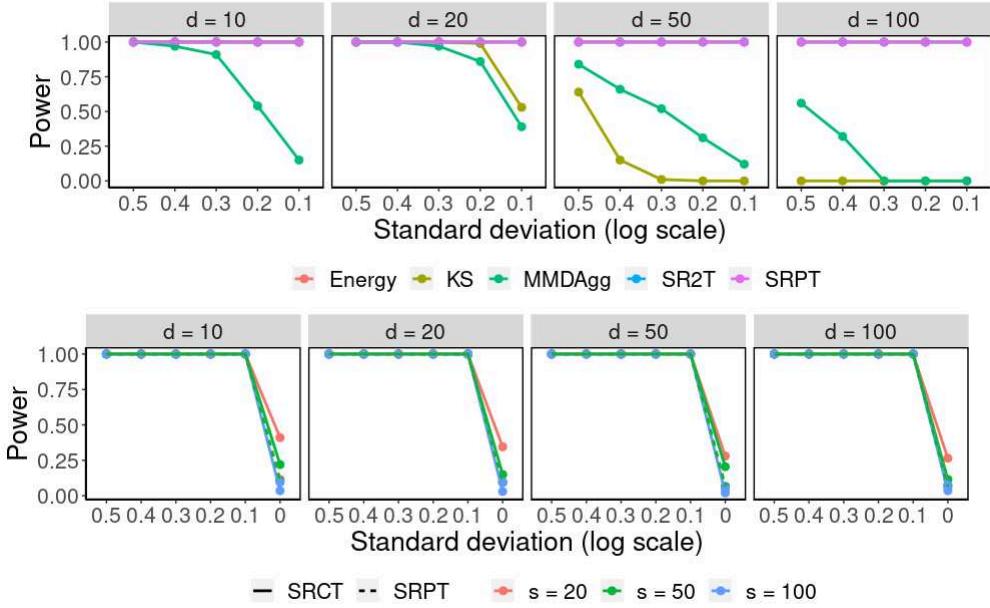
Remark 6 *Theorem 17 states that choosing any $m \geq n$ should be enough to achieve the same separation boundary up to constants as the Oracle test. Using $m = 3n$ as compared to $m = n$ will theoretically yield the same separation boundary in terms of n but with a better constant closer to that of the Oracle test. To demonstrate this point, for the rest of the experiments, we used $m = 3n$.*

5.2 Gaussian distribution

In this section, we examine the Gaussian location shift and covariance scale problems, where the observed samples are generated from a Gaussian distribution with a shifted mean or scaled covariance matrix (by scaling the diagonal elements of the identity matrix). The goal is to test the null hypothesis of standard Gaussian distribution. Figure 2(a) shows the power for different mean shifts and different dimensions from which we note that the Energy test gives the best power closely followed by the SRPT test. Figure 2(b) shows the result for different choices of s for both SRCT and SRPT tests with Showalter regularization. We can see that SRPT is not very sensitive to the choice of s as opposed to SRCT, which seems to give higher power for lower values of s , however with the cost of a worse Type-I error (shown at mean shift = 0). We can see that the choice of $s = 100$ controls both power and Type-I error for SRCT and for this choice of $s = 100$, the permutation test SRPT yields a higher power while still controlling the Type-I error. Similarly, Figure 3 shows the power for different scaling factors with different dimensions and different choices of s , demonstrating similar results.

5.3 Perturbed uniform distribution

In this part, we examine testing the null hypothesis of uniform distribution against perturbed uniform distribution for different values of perturbation, P (see Hagrass et al., 2024,

Figure 2: Power for Gaussian shift experiments with different d and s using $n = 200$.Figure 3: Power for Gaussian covariance scale experiments with different d and s using $n = 200$.

Section 5.1.1 for details). Figure 4(a) shows the result for $d \in \{1, 2\}$ for different perturbations, wherein we can see that the highest power is achieved by SRPT. Figure 4(b) shows the power for SRCT and SRPT for different choices of s , with $P = 0$ corresponding to no perturbations and thus showing Type-I error. Similar to the observation from the previous section, SRPT is not very sensitive to the choice of s , while SRCT is sensitive to s , with

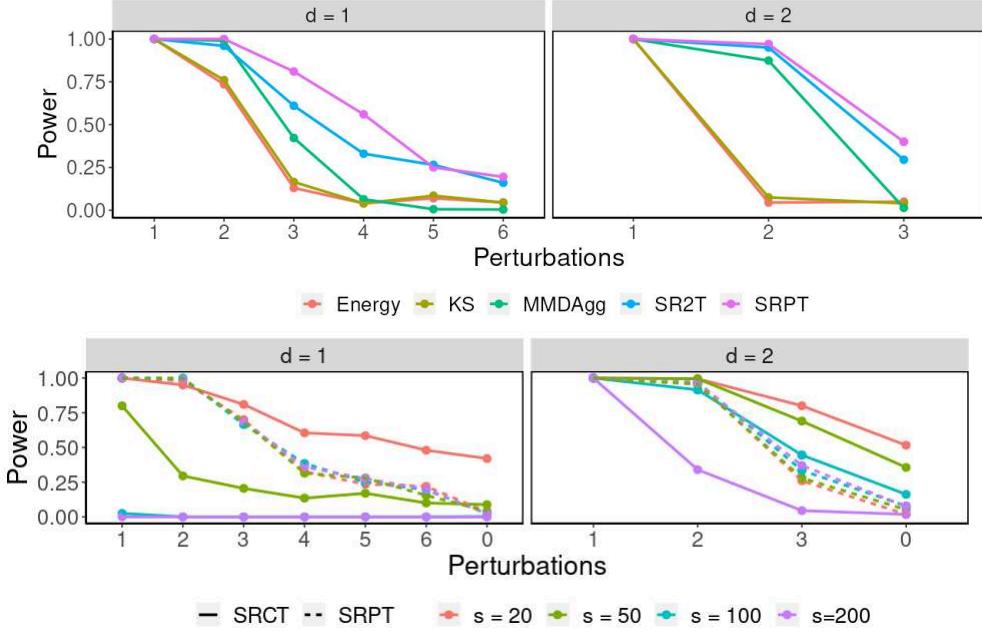


Figure 4: Power for perturbed uniform distributions for $d = 1$ ($n = 500$) and $d = 2$ ($n = 2000$).

$s = 50$ and $s = 200$ being the reasonable choices, respectively for $d = 1$ and $d = 2$ that controls both power and Type-I error.

5.4 Directional data

In this section, we investigate two experiments involving directional domains, where we focus on testing for a multivariate von Mises-Fisher distribution, which serves as the Gaussian analog on the unit sphere defined by the density $f(x, \mu, k) = \frac{k^{d/2-1}}{2\pi^{d/2} I_{d/2-1}(k)} \exp(k\mu^T x)$, $x \in \mathbb{S}^{d-1}$, with $k \geq 0$ being the concentration parameter, μ being the mean parameter and I being the modified Bessel function. Figure 5(a) shows the results for testing von Mises-Fisher distribution against spherical uniform distribution ($k = 0$) for different concentration parameters. We can see from Figure 5(a) that the best power is achieved by the Energy test followed closely by SRPT. Figure 5(b) shows that SRPT is less sensitive to the choice of s as opposed to SRCT which achieves its best power at $s = 100$ while still controlling the Type-I error. In the second experiment, we explore a mixture of two multivariate Watson distributions, representing axially symmetric distributions on a sphere, given by $f(x, \mu, k) = \frac{\Gamma(d/2)}{2\pi^{d/2} M(1/2, d/2, \kappa)} \exp(k(\mu^T x)^2)$, $x \in \mathbb{S}^{d-1}$, where $k \geq 0$ is the concentration parameter, μ is the mean parameter and M is Kummer's confluent hypergeometric function. Using equal weights we drew 500 samples from a mixture of two Watson distributions with similar concentration parameter k and mean parameter μ_1, μ_2 respectively, where $\mu_1 = (1, \dots, 1) \in \mathbb{R}^d$ and $\mu_2 = (-1, 1, \dots, 1) \in \mathbb{R}^d$. Figure 6(a) shows the power against spherical uniform distribution for different concentration parameters. We can see that SRPT outperforms all the other methods. Figure 6(b) shows how the power and Type-I error are affected by the choice of s , which similar to the previous sections shows that SRPT is not

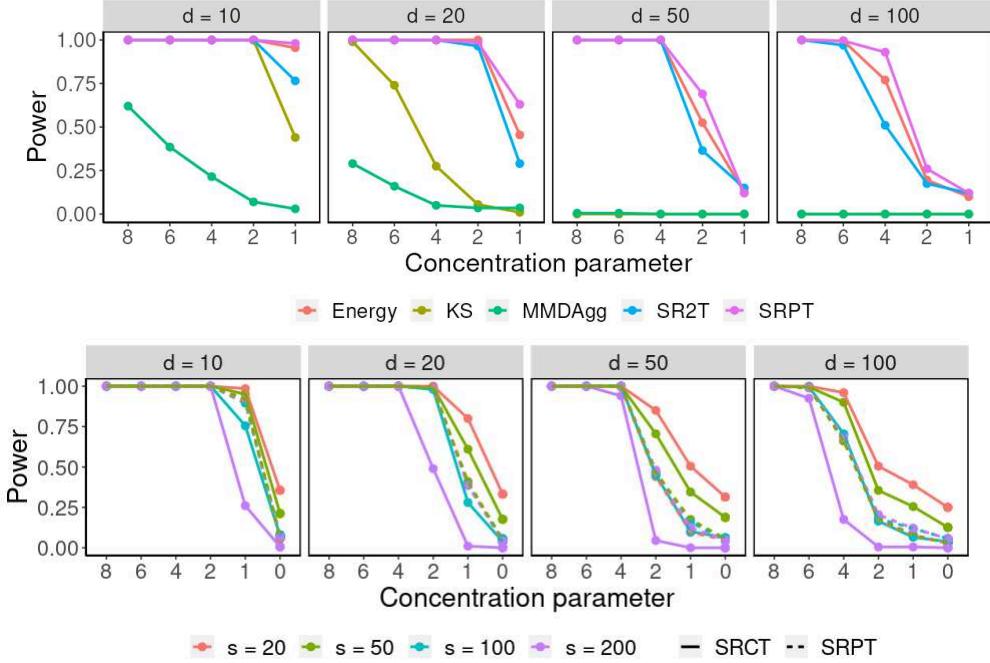


Figure 5: Power for von Mises-Fisher distribution with different concentration parameter k and s using $n = 500$.

very sensitive to s , while SRCT achieves its best power while still controlling for Type-I error at $s = 100$.

6 Discussion

To summarize, in this work, we have extended and generalized the theoretical properties of the Oracle test proposed by Balasubramanian et al. (2021) by employing a general spectral regularization approach, wherein we obtained sufficient conditions for the separation boundary under weaker assumptions and for a wider range of alternatives. Under the assumption that we have access to samples from P_0 , we addressed the problem of the practicality of the Oracle test by proposing two completely data-driven tests (SRCT and SRPT) that adapt to the choice of the kernel, the eigenvalue decay rate and the smoothness of the likelihood ratio, while still being minimax optimal (up to logarithmic factors) w.r.t. \mathcal{P} . Through numerical experiments, we established the superior performance of the proposed spectral regularized tests over the MMD-based test and the closely related two-sample test proposed in Haggrass et al. (2024). However, there are still some open questions for future consideration: (i) Improving the computational complexity of the proposed tests using approximation schemes like random Fourier features (Rahimi and Recht, 2008), Nyström method (e.g., Williams and Seeger 2001; Drineas and Mahoney 2005) or sketching (Yang et al., 2017), and studying the computational vs. statistical trade-off for the approximate test. (ii) The proposed test requires access to i.i.d. samples from P_0 , which might not be easy to generate or easily available. As an alternative, instead of regularizing w.r.t. $\hat{\Sigma}_{P_0}$ (i.e., the empirical covariance operator estimated based on the samples from P_0), one can regularize with respect to $\hat{\Sigma}_P$

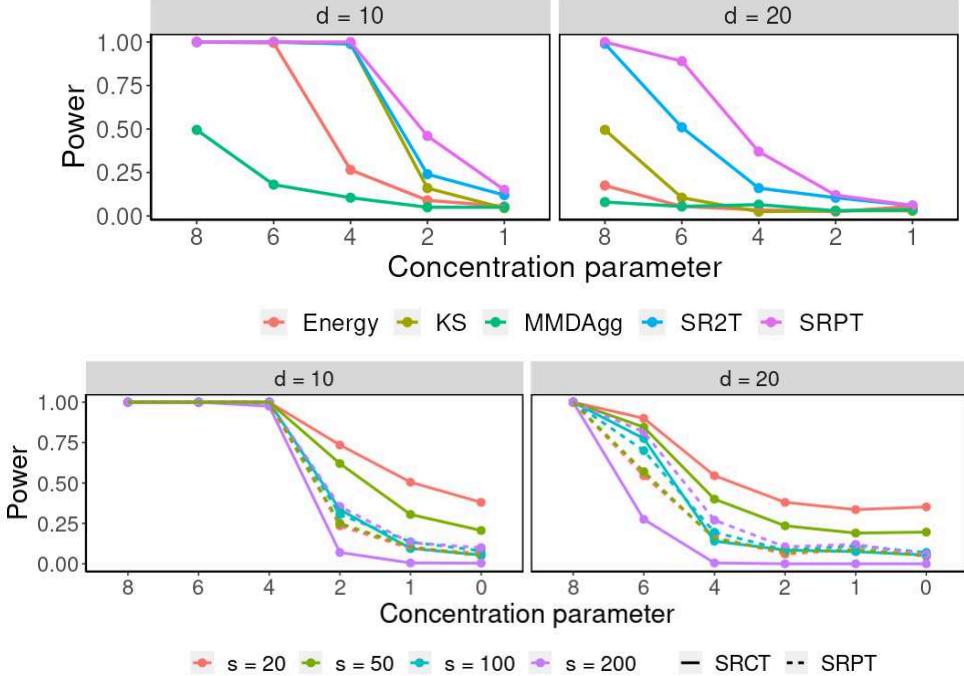


Figure 6: Power for mixture of Watson distributions with different concentration parameter k and s using $n = 500$.

(i.e., the empirical covariance operator estimated from the available samples drawn from P). However, even this approach is not practical unless the mean element μ_0 is computable, which need not be the case. To completely get around this issue, we can consider applying the idea of spectral regularization to Kernel Stein Discrepancy (KSD) (Chwialkowski et al., 2016; Liu et al., 2016) which does not require computing any integrals with respect to P_0 , and study its minimax optimality.

7 Proofs

In this section, we present the proofs of the main results of the paper.

7.1 Proof of Theorem 1

Define $b(x) = K(\cdot, x) - \mu_P$ and $a(x) = b(x) + (\mu_P - \mu_0) = K(\cdot, x) - \mu_0$. Thus we can write,

$$\begin{aligned}
 \hat{D}_{\text{MMD}}^2 &= \frac{1}{n(n-1)} \sum_{i \neq j} \langle a(X_i), a(X_j) \rangle_{\mathcal{H}} \\
 &= \frac{1}{n(n-1)} \sum_{i \neq j} \langle b(X_i), b(X_j) \rangle_{\mathcal{H}} + \frac{2}{n} \sum_{i=1}^n \langle b(X_i), (\mu_P - \mu_0) \rangle_{\mathcal{H}} + D_{\text{MMD}}^2 \\
 &= I_1 + I_2 + D_{\text{MMD}}^2,
 \end{aligned}$$

where

$$I_1 = \frac{1}{n(n-1)} \sum_{i \neq j} \langle b(X_i), b(X_j) \rangle_{\mathcal{H}}, \quad \text{and} \quad I_2 = \frac{2}{n} \sum_{i=1}^n \langle b(X_i), (\mu_P - \mu_0) \rangle_{\mathcal{H}}.$$

so that $\mathbb{E}[(\hat{D}_{\text{MMD}}^2 - D_{\text{MMD}}^2)^2] = \mathbb{E}_P(I_1 + I_2)^2 \leq 2\mathbb{E}_P(I_1^2) + 2\mathbb{E}_P(I_2^2)$. Next, following similar ideas as in the proofs of (Hagrass et al., 2024, Lemmas A.4, A.5), we can bound I_1 and I_2 as

$$\mathbb{E}(I_1^2) \leq \frac{4}{n^2} \|\Sigma_P\|_{\mathcal{L}^2(\mathcal{H})}^2, \quad \text{and} \quad \mathbb{E}(I_2^2) \leq \frac{4}{n} \|\Sigma_P\|_{\mathcal{L}^\infty(\mathcal{H})} \|\mu_P - \mu_0\|_{\mathcal{H}}^2,$$

respectively. Combining these bounds yields that

$$\mathbb{E}[(\hat{D}_{\text{MMD}}^2 - D_{\text{MMD}}^2)^2] \lesssim \frac{1}{n^2} + \frac{D_{\text{MMD}}^2}{n}. \quad (14)$$

When $P = P_0$, we have $D_{\text{MMD}}^2 = 0$. Therefore under H_0 ,

$$\mathbb{E}[(\hat{D}_{\text{MMD}}^2)^2] \leq \frac{4 \|\Sigma_0\|_{\mathcal{L}^2(\mathcal{H})}^2}{n^2} \stackrel{(*)}{\leq} \frac{16\kappa^2}{n^2}, \quad (15)$$

where in $(*)$ we used where in $(*)$ we used $\|\Sigma_0\|_{\mathcal{L}^2(\mathcal{H})}^2 \leq 4\kappa^2$. Thus using (15) and Chebyshev's inequality yields

$$P_{H_0}\{\hat{D}_{\text{MMD}}^2 \geq \gamma\} \leq \alpha,$$

where $\gamma = \frac{4\kappa}{\sqrt{\alpha n}}$.

Next, we use the bound in (14) to bound the power. Let $\gamma_1 = \frac{1}{\sqrt{\delta n}} + \frac{\sqrt{D_{\text{MMD}}^2}}{\sqrt{\delta} \sqrt{n}}$. Then

$$\begin{aligned} P_{H_1}\{\hat{D}_{\text{MMD}}^2 \geq \gamma\} &\stackrel{(*)}{\geq} P_{H_1}\{\hat{D}_{\text{MMD}}^2 > D_{\text{MMD}}^2 - \gamma_1\} \\ &\stackrel{(**)}{\geq} P_{H_1}\{|\hat{D}_{\text{MMD}}^2 - D_{\text{MMD}}^2| \leq \gamma_1\} \geq 1 - \delta, \end{aligned}$$

where $(*)$ holds when $D_{\text{MMD}}^2 \geq \gamma + \gamma_1$, which is implied if $D_{\text{MMD}}^2 \gtrsim \frac{1}{n}$, which in turn is implied if $\|u\|_{L^2(P_0)}^2 \gtrsim n^{\frac{-2\theta}{2\theta+1}}$, where the last implication follows from (Hagrass et al., 2024, Lemma A.19). $(**)$ follows from (14) and an application of Chebyshev's inequality. The desired result, therefore, holds by taking infimum over $P \in \mathcal{P}$.

Finally, we will show that we cannot achieve a rate better than $n^{\frac{-2\theta}{2\theta+1}}$ over \mathcal{P} . Recall that $\mathcal{T} = \sum_{i \in I} \lambda_i \tilde{\phi}_i \otimes_{L^2(P_0)} \tilde{\phi}_i$. Let $\bar{\phi}_i = \phi_i - \mathbb{E}_{P_0} \phi_i$, where $\phi_i = \frac{\mathfrak{I}^* \tilde{\phi}_i}{\lambda_i}$. Then $\mathfrak{I} \bar{\phi}_i = \mathfrak{I} \phi_i = \frac{\mathfrak{I} \tilde{\phi}_i}{\lambda_i} = \tilde{\phi}_i$. Assuming $\lambda_i = h(i)$, where h is an invertible, continuous function (for example $h = i^{-\beta}$ and $h = e^{-\tau i}$ correspond to polynomial and exponential decays respectively), let $k = \lfloor h^{-1}(n^{\frac{-1}{2\theta+1}}) \rfloor$, hence $\lambda_k = n^{\frac{-1}{2\theta+1}}$. Define

$$f := b \bar{\phi}_k,$$

where $b < \sqrt{\frac{4\kappa}{4\sqrt{\alpha}}}$. Then $\|f\|_{L^2(P_0)}^2 = b^2 \lesssim 1$, and thus $f \in L^2(P_0)$. Define

$$\tilde{u} := \mathcal{T}^\theta f = b \lambda_k^\theta \tilde{\phi}_k, \quad \text{and} \quad u := b \lambda_k^\theta \bar{\phi}_k.$$

Note that $\mathbb{E}_{P_0} u = b\lambda_k^\theta \mathbb{E}_{P_0} \bar{\phi}_k = 0$. Since $\mathfrak{I}u = \tilde{u}$, we have $u \in [\tilde{u}]_\sim \in \text{Ran}(\mathcal{T}^\theta)$, $\|u\|_{L^2(P_0)}^2 = b^2 \lambda_k^{2\theta} > \Delta_n$. Next we bound $|u(x)|$ in the following two cases.

Case I: $\theta \geq \frac{1}{2}$ and $\sup_i \|\phi_i\|_\infty$ is not finite.

Note that

$$|u(x)| = b\lambda_k^\theta |\langle k(\cdot, x) - \mu_0, \phi_k \rangle_{\mathcal{H}}| \leq b\lambda_k^\theta \|k(\cdot, x) - \mu_0\|_{\mathcal{H}} \|\phi_k\|_{\mathcal{H}} \stackrel{(*)}{\leq} 2b\sqrt{\kappa} \lambda_k^{\theta - \frac{1}{2}} \stackrel{(\dagger)}{\leq} 1,$$

where in $(*)$ we used $\|\phi_k\|_{\mathcal{H}}^2 = \lambda_k^{-2} \langle \mathfrak{I}^* \tilde{\phi}_k, \mathfrak{I}^* \phi_k \rangle = \lambda_k^{-1}$. In (\dagger) we used $\theta > \frac{1}{2}$.

Case II: $\sup_i \|\phi_i\|_\infty < \infty$.

In this case,

$$|u(x)| \leq 2b \sup_k \|\phi_k\|_\infty \lambda_k^\theta \leq 1,$$

for n large enough. This implies that we can find $P \in \mathcal{P}$ such that $\frac{dP}{dP_0} = u + 1$. Then for such P , we have $D_{\text{MMD}}^2 = \|\mathcal{T}^{1/2}u\|_{L^2(P_0)}^2 = b^2 \lambda_k^{2\theta+1} = \frac{b^2}{n} < \frac{4\kappa}{\sqrt{\alpha n}} = \gamma$. Therefore there exists some $\varepsilon > 0$ such that $D_{\text{MMD}}^2 < \gamma - \varepsilon$. Hence, we have

$$\begin{aligned} P_{H_1}\{\hat{D}_{\text{MMD}}^2 \geq \gamma\} &< P_{H_1}\{\hat{D}_{\text{MMD}}^2 \geq D_{\text{MMD}}^2 + \varepsilon\} < P_{H_1}\{|\hat{D}_{\text{MMD}}^2 - D_{\text{MMD}}^2| \geq \varepsilon\} \\ &\stackrel{(*)}{\lesssim} \frac{1}{\varepsilon^2 n^2} \leq \delta, \end{aligned}$$

where we used (14) along with Chebyshev's inequality in $(*)$, and the last inequality holds for $n > \frac{1}{\varepsilon\sqrt{\delta}}$.

7.2 Proof of Theorem 2

As shown in (Kim et al., 2022, Lemma G.1), in order to show that a separation boundary Δ_n will imply $R_{\Delta_n}^* \geq \delta$, it is sufficient to find a set of distributions $\{P_k\}_{k=1}^L \subset \mathcal{P}$, such that

$$\mathbb{E}_{P_0^n} \left[\left(\frac{1}{L} \sum_{k=1}^L \frac{dP_k^n}{dP_0^n} \right)^2 \right] \leq 1 + 4(1 - \alpha - \delta)^2. \quad (16)$$

Then the proof follows the same ideas as used in the proof of (Hagras et al., 2024, Theorem 2) as shown briefly below.

Recall $\mathcal{T} = \sum_{i \in I} \lambda_i \tilde{\phi}_i \otimes_{L^2(P_0)} \tilde{\phi}_i$. Let $\bar{\phi}_i = \phi_i - \mathbb{E}_{P_0} \phi_i$, where $\phi_i = \frac{\mathfrak{I}^* \tilde{\phi}_i}{\lambda_i}$. Then $\mathfrak{I} \bar{\phi}_i = \mathfrak{I} \phi_i = \frac{\mathcal{T} \tilde{\phi}_i}{\lambda_i} = \tilde{\phi}_i$.

Polynomial decay (Case I): $\lambda_i \asymp i^{-\beta}$, $\beta > 1$, $\sup_i \|\phi_i\|_\infty < \infty$ and $\theta \geq \frac{1}{4\beta}$.

Let

$$B_n = \left\lfloor \frac{\Delta_n^{-1/2\theta\beta}}{16(\sup_i \|\phi_i\|_\infty)^2} \right\rfloor,$$

$C_n = \lfloor \sqrt{B_n} \rfloor$ and $a_n = \sqrt{\frac{\Delta_n}{C_n}}$. For $k \in \{1, \dots, L\}$, define

$$u_{n,k} := a_N \sum_{i=1}^{B_n} \varepsilon_{ki} \bar{\phi}_i,$$

where $\varepsilon_k := \{\varepsilon_{k1}, \varepsilon_{k2}, \dots, \varepsilon_{kB_n}\} \in \{0, 1\}^{B_n}$ such that $\sum_{i=1}^{B_n} \varepsilon_{ki} = C_n$, thus $L = \binom{B_n}{C_n}$. Then it can be shown (see the proof of Hagrass et al. 2024, Theorem 2) that we can find $P_k \in \mathcal{P}$ such that $\frac{dP_k}{dP_0} = u_{n,k} + 1$, and that (16) holds for $\theta > \frac{1}{4\beta}$ when $\Delta_n \leq c(\alpha, \delta) n^{\frac{-4\theta\beta}{4\theta\beta+1}}$ for some $c(\alpha, \delta)$.

Polynomial decay (Case II): $\lambda_i \asymp i^{-\beta}$, $\beta > 1$, $\theta \geq \frac{1}{2}$ and $\sup_i \|\phi_i\|_\infty$ is not finite.

Since $\lambda_i \asymp i^{-\beta}$, $\beta > 1$, there exists constants $\underline{A} > 0$ and $\bar{A} > 0$ such that $\underline{A}i^{-\beta} \leq \lambda_i \leq \bar{A}i^{-\beta}$. Let $B_n = \lfloor \left(\frac{\underline{A}\Delta_n^{-1}}{4\kappa} \right)^{\frac{1}{2\theta\beta}} \rfloor$, $C_n = \lfloor \sqrt{B_n} \rfloor$ and $a_n = \sqrt{\frac{\Delta_n}{C_n}}$. Then similar to Case I, it can be shown that (see the proof of Hagrass et al. 2024, Theorem 2) we can find $P_k \in \mathcal{P}$ such that $\frac{dP_k}{dP_0} = u_{n,k} + 1$, and that (16) holds for $\theta > \frac{1}{2}$ when $\Delta_n \leq c(\alpha, \delta) n^{\frac{-4\theta\beta}{4\theta\beta+1}}$ for some $c(\alpha, \delta)$.

Exponential decay: $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$.

Since $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$, there exists constants $\underline{A} > 0$ and $\bar{A} > 0$ such that $\underline{A}e^{-\tau i} \leq \lambda_i \leq \bar{A}e^{-\tau i}$. Let $B_n = \lfloor (2\tau \max\{\theta, \frac{1}{2}\})^{-1} \log(\frac{\underline{A}}{4\kappa\Delta_n}) \rfloor$, $C_n = \lfloor \sqrt{B_n} \rfloor$ and $a_n = \sqrt{\frac{\Delta_n}{C_n}}$, where $\theta > 0$. Then similar to the previous cases it can be shown that (16) holds when $\Delta_n \leq c(\alpha, \delta, \theta) \frac{(\log n)^b}{n}$ for any $b < \frac{1}{2}$. Thus the desired bound holds by taking supremum over $b < \frac{1}{2}$.

7.3 Proof of Theorem 3

By defining $\mathcal{B} := g_\lambda^{1/2}(\Sigma_0)\Sigma_{0,\lambda}^{1/2}$ and $a(x) = \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(K(\cdot, x) - \mu_0)$, we have

$$\hat{\eta}_\lambda = \frac{1}{n(n-1)} \sum_{i \neq j} \langle a(X_i), a(X_j) \rangle_{\mathcal{H}}.$$

By replacing Σ_{PQ} with Σ_0 in the proof of (Hagrass et al., 2024, Lemma A.4), we have

$$\mathbb{E}_{P_0}(\hat{\eta}_\lambda^2) \leq \frac{4}{n^2} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \mathcal{N}_2^2(\lambda).$$

The result therefore follows by applying Chebyshev's inequality and noting from (Hagrass et al., 2024, Lemma A.8(ii)) that $\|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \leq (C_1 + C_2)$.

7.4 Proof of Theorem 4

Define $\mathcal{B} := g_\lambda^{1/2}(\Sigma_0)\Sigma_{0,\lambda}^{1/2}$, $b(x) = \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(K(\cdot, x) - \mu_P)$, and $a(x) = b(x) + \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(\mu_P - \mu_0) = \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(K(\cdot, x) - \mu_0)$. Thus we can write,

$$\begin{aligned} \hat{\eta}_\lambda &= \frac{1}{n(n-1)} \sum_{i \neq j} \langle b(X_i), b(X_j) \rangle_{\mathcal{H}} + \frac{2}{n} \sum_{i=1}^n \left\langle b(X_i), \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(\mu_P - \mu_0) \right\rangle_{\mathcal{H}} + \eta_\lambda \\ &= I_1 + I_2 + \eta_\lambda, \end{aligned}$$

where

$$I_1 = \frac{1}{n(n-1)} \sum_{i \neq j} \langle b(X_i), b(X_j) \rangle_{\mathcal{H}}, \quad \text{and} \quad I_2 = \frac{2}{n} \sum_{i=1}^n \left\langle b(X_i), \mathcal{B} \Sigma_{0,\lambda}^{-1/2} (\mu_P - \mu_0) \right\rangle_{\mathcal{H}}.$$

Thus, $\text{Var}_P[\hat{\eta}_\lambda] = \mathbb{E}_P(I_1 + I_2)^2 \leq 2\mathbb{E}_P(I_1^2) + 2\mathbb{E}_P(I_2^2)$. Next we bound $\mathbb{E}_P(I_1^2)$ and $\mathbb{E}_P(I_2^2)$ using (Haggrass et al., 2024, Lemmas A.4 and A.5) by replacing Σ_{PQ} with Σ_0 , which yields

$$\begin{aligned} \text{Var}_P[\hat{\eta}_\lambda] &\leq \frac{8}{n^2} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}^2 \\ &\quad + \frac{8}{n} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^\infty(\mathcal{H})} \left\| \Sigma_{0,\lambda}^{-1/2} (\mu_P - \mu_0) \right\|_{\mathcal{H}}^2 \\ &\stackrel{(*)}{\leq} \frac{8}{n^2} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 (4C_\lambda \|u\|_{L^2(P_0)}^2 + 2\mathcal{N}_2^2(\lambda)) \\ &\quad + \frac{8}{n} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 (2\sqrt{C_\lambda} \|u\|_{L^2(P_0)} + 1) \|u\|_{L^2(P_0)}^2, \end{aligned}$$

where C_λ as defined in Lemma A.3 and in (*), we used Lemmas A.2 and A.3. Then we can easily deduce by using Chebyshev's inequality that $P_{H_1}\{\hat{\eta}_\lambda \geq \gamma\} \geq 1 - \delta$ holds if $\eta_\lambda \geq \gamma + \sqrt{\frac{\text{Var}_P[\hat{\eta}_\lambda]}{\delta}}$. Using Lemma A.2, we have $\eta_\lambda \geq \frac{B_3}{4} \|u\|_{L^2(P_0)}^2$ under the assumptions $u \in \text{Ran}(\mathcal{T}^\theta)$, and

$$\|u\|_{L^2(P_0)}^2 \geq \frac{4C_3}{3B_3} \|\mathcal{T}\|_{\mathcal{L}^\infty(\mathcal{H})}^{2\max(\theta-\xi,0)} \lambda^{2\tilde{\theta}} \left\| \mathcal{T}^{-\theta} u \right\|_{L^2(P_0)}^2. \quad (17)$$

Note that $u \in \text{Ran}(\mathcal{T}^\theta)$ is guaranteed since $P \in \mathcal{P}$ and

$$\|u\|_{L^2(P_0)}^2 \geq c_4 \lambda^{2\tilde{\theta}} \quad (18)$$

guarantees (17) since $\|\mathcal{T}\|_{\mathcal{L}^\infty(\mathcal{H})} = \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} \leq 2\kappa$ and $c_1 := \sup_{P \in \mathcal{P}} \|\mathcal{T}^{-\theta} u\|_{L^2(P_0)} < \infty$, where $c_4 = \frac{4c_1^2 C_3 (2\kappa)^{2\max(\theta-\xi,0)}}{3B_3}$. Thus,

$$\frac{B_3}{2} \|u\|_{L^2(P_0)}^2 \geq \gamma + \sqrt{\frac{\text{Var}_P[\hat{\eta}_\lambda]}{\delta}} \quad (19)$$

guarantees $\eta_\lambda \geq \gamma + \sqrt{\frac{\text{Var}_P[\hat{\eta}_\lambda]}{\delta}}$. Hence it remains to verify (18) and (19). Using $\|u\|_{L^2(P_0)}^2 \geq \Delta_n$, it is easy to see that (18) is implied when $\lambda = (c_4^{-1} \Delta_n)^{1/2\tilde{\theta}}$. Using $\|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \leq (C_1 + C_2)^2$, which follows from (Haggrass et al., 2024, Lemma A.8 (ii)) by replacing Σ_{PQ} with Σ_0 , and substituting the expressions of γ and $\text{Var}_P[\hat{\eta}_\lambda]$ in (19), we can verify that (19) is implied if $\Delta_n \geq \frac{r_1 \mathcal{N}_2(\lambda)}{n\sqrt{\alpha}}$, $\Delta_n \geq \frac{r_2 C_\lambda}{\delta^2 n^2}$ and $\Delta_n \geq \frac{r_3 \mathcal{N}_2(\lambda)}{\delta n}$ for some constants $r_1, r_2, r_3 > 0$. Hence the desired result follows by taking infimum over $P \in \mathcal{P}$.

7.5 Proof of Corollary 5

When $\lambda_i \asymp i^{-\beta}$, we have $\mathcal{N}_2(\lambda) \leq \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^\infty(\mathcal{H})}^{1/2} \mathcal{N}_1^{1/2}(\lambda) \lesssim \lambda^{-1/2\beta}$ (see Sriperumbudur and Sterge 2022, Lemma B.9). Using this bound in the conditions mentioned in

Theorem 4, ensures that these conditions on the separation boundary hold if

$$\Delta_n \gtrsim \max \left\{ \left(\frac{n}{\alpha^{-1/2} + \delta^{-1}} \right)^{-\frac{4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, (\delta n)^{-\frac{8\tilde{\theta}\beta}{4\tilde{\theta}\beta+2\beta+1}} \right\}, \quad (20)$$

which in turn is implied if

$$\Delta_n = \begin{cases} c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, & \tilde{\theta} > \frac{1}{2} - \frac{1}{4\beta} \\ c(\alpha, \delta) n^{-\frac{8\tilde{\theta}\beta}{4\tilde{\theta}\beta+2\beta+1}}, & \tilde{\theta} \leq \frac{1}{2} - \frac{1}{4\beta} \end{cases},$$

where $c(\alpha, \delta) \gtrsim (\alpha^{-1/2} + \delta^{-2})$ and we used that $\tilde{\theta} > \frac{1}{2} - \frac{1}{4\beta} \Leftrightarrow \frac{4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1} < \frac{8\tilde{\theta}\beta}{4\tilde{\theta}\beta+2\beta+1}$. On the other hand when $C := \sup_i \|\phi_i\|_\infty < \infty$, we obtain the corresponding condition as

$$\Delta_n \gtrsim \max \left\{ \left(\frac{n}{\alpha^{-1/2} + \delta^{-1}} \right)^{-\frac{4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, (\delta n)^{-\frac{4\tilde{\theta}\beta}{2\tilde{\theta}\beta+1}} \right\}, \quad (21)$$

which is implied if

$$\Delta_n = c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}.$$

7.6 Proof of Corollary 6

When $\lambda_i \asymp e^{-\tau i}$, we have $\mathcal{N}_2(\lambda) \leq \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^\infty(\mathcal{H})}^{1/2} \mathcal{N}_1^{1/2}(\lambda) \lesssim \sqrt{\log \frac{1}{\lambda}}$ (see Sriperumbudur and Sterge 2022, Lemma B.9). Thus, substituting this in the conditions from Theorem 4 and using $n \geq \max\{e^2, \alpha^{-1/2} + \delta^{-1}\}$, we can write the separation boundary as

$$\Delta_n \gtrsim \max \left\{ \left(\frac{\sqrt{2\tilde{\theta}}(\alpha^{-1/2} + \delta^{-1})^{-1}n}{\sqrt{\log n}} \right)^{-1}, \left(\frac{\delta n}{\sqrt{\log n}} \right)^{-\frac{4\tilde{\theta}}{2\tilde{\theta}+1}} \right\}, \quad (22)$$

which is implied if

$$\Delta_n = \begin{cases} c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n}, & \tilde{\theta} > \frac{1}{2} \\ c(\alpha, \delta, \theta) \left(\frac{\sqrt{\log n}}{n} \right)^{\frac{4\tilde{\theta}}{2\tilde{\theta}+1}}, & \tilde{\theta} \leq \frac{1}{2} \end{cases},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$ and we used that $\tilde{\theta} > \frac{1}{2} \Leftrightarrow 1 < \frac{4\tilde{\theta}}{2\tilde{\theta}+1}$.

On the other hand when $C := \sup_i \|\phi_i\|_\infty < \infty$, we obtain

$$\Delta_n \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}} \left(\frac{(\alpha^{-1/2} + \delta^{-1})^{-1}n}{\sqrt{\log n}} \right)^{-1}, \frac{1}{2\tilde{\theta}} \left(\frac{\delta n}{\sqrt{\log n}} \right)^{-2} \right\},$$

which in turn is implied if

$$\Delta_n = c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, \frac{1}{2\tilde{\theta}}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$.

7.7 Proof of Theorem 7

The proof is exactly similar to that of (Hagras et al., 2024, Theorem 3) by replacing Σ_{PQ} with Σ_0 .

7.8 Proof of Theorem 8

Since $\mathbb{E}(\hat{\eta}_\lambda^{TS}|(Y_i^0)_{i=1}^s) = 0$, an application of Chebyshev's inequality via Lemma A.8 yields,

$$P_{H_0} \left\{ |\hat{\eta}_\lambda^{TS}| \geq \frac{\sqrt{6}(C_1 + C_2) \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \mathcal{N}_2(\lambda)}{\sqrt{\delta}} \left(\frac{1}{n} + \frac{1}{m} \right) \middle| (Y_i^0)_{i=1}^s \right\} \leq \delta.$$

Let $\gamma_1 := \frac{2\sqrt{6}(C_1 + C_2)\mathcal{N}_2(\lambda)}{\sqrt{\delta}} \left(\frac{1}{n} + \frac{1}{m} \right)$, and $\gamma_2 := \frac{\sqrt{6}(C_1 + C_2)\|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \mathcal{N}_2(\lambda)}{\sqrt{\delta}} \left(\frac{1}{n} + \frac{1}{m} \right)$. Then

$$\begin{aligned} P_{H_0}\{\hat{\eta}_\lambda^{TS} \leq \gamma_1\} &\geq P_{H_0}\{\hat{\eta}_\lambda^{TS} \leq \gamma_2\} \cap \{\gamma_2 \leq \gamma_1\} \\ &\geq 1 - P_{H_0}\{\hat{\eta}_\lambda^{TS} \geq \gamma_2\} - P_{H_0}\{\gamma_2 \geq \gamma_1\} \stackrel{(*)}{\geq} 1 - 3\delta, \end{aligned}$$

where $(*)$ follows using

$$P_{H_0}\{\hat{\eta}_\lambda^{TS} \geq \gamma_2\} \leq P_{H_0}\{|\hat{\eta}_\lambda^{TS}| \geq \gamma_2\} = \mathbb{E}_{P_0^m} [P_{H_0}\{|\hat{\eta}_\lambda^{TS}| \geq \gamma_2 | (Y_i^0)_{i=1}^s\}] \leq \delta,$$

and

$$P_{H_0}\{\gamma_2 \geq \gamma_1\} = P_{H_0}\{\|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \geq 2\} \stackrel{(\dagger)}{\leq} 2\delta,$$

where (\dagger) follows from (Sriperumbudur and Sterge, 2022, Lemma B.2(ii)), under the condition that $\frac{140\kappa}{s} \log \frac{16\kappa s}{\delta} \leq \lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$. When $C := \sup_i \|\phi_i\|_\infty < \infty$, using (Hagras et al., 2024, Lemma A.17), we can obtain an improved condition on λ satisfying $136C^2\mathcal{N}_1(\lambda) \log \frac{8\mathcal{N}_1(\lambda)}{\delta} \leq s$ and $\lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$. Thus setting $\delta = \frac{\alpha}{6}$, yields that

$$P_{H_0} \left\{ \hat{\eta}_\lambda^{TS} \geq \frac{12(C_1 + C_2)\mathcal{N}_2(\lambda)}{\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right) \right\} \leq \frac{\alpha}{2}. \quad (23)$$

Finally, the desired result follows by writing

$$\begin{aligned} P_{H_0} \left\{ \hat{\eta}_\lambda^{TS} \leq \frac{12(C_1 + C_2)\widehat{\mathcal{N}}_2(\lambda)}{b_1\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right) \right\} \\ \geq P_{H_0} \left\{ \left\{ \hat{\eta}_\lambda^{TS} \leq \frac{12(C_1 + C_2)\mathcal{N}_2(\lambda)}{\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right) \right\} \cap \{\widehat{\mathcal{N}}_2(\lambda) \geq b_1\mathcal{N}_2(\lambda)\} \right\} \\ \geq 1 - P_{H_0} \left\{ \hat{\eta}_\lambda^{TS} \geq \frac{12(C_1 + C_2)\mathcal{N}_2(\lambda)}{\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right) \right\} - P_{H_0}\{\widehat{\mathcal{N}}_2(\lambda) \leq b_1\mathcal{N}_2(\lambda)\} \\ \stackrel{(\ddagger)}{\geq} 1 - \alpha, \end{aligned}$$

where (\ddagger) follows using (23) and Lemma A.7 under the condition that

$$\frac{4c_1\kappa}{s} \max\{\log \frac{96\kappa s}{\alpha}, \log \frac{12}{\alpha}\} \leq \lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}.$$

The above condition can be replaced with $4c_1C^2\mathcal{N}_1(\lambda) \log \frac{48\mathcal{N}_1(\lambda)}{\alpha} \leq s$ if $C := \sup_i \|\phi_i\|_\infty < \infty$.

7.9 Proof of Theorem 9

Let $\mathcal{M} = \hat{\Sigma}_{0,\lambda}^{-1/2} \Sigma_{0,\lambda}^{1/2}$, $\gamma_1 = \frac{1}{\sqrt{\delta}} \left(\frac{\sqrt{C_\lambda} \|u\|_{L^2(P_0)} + \mathcal{N}_2(\lambda)}{n} + \frac{C_\lambda^{1/4} \|u\|_{L^2(P_0)}^{3/2} + \|u\|_{L^2(P_0)}}{\sqrt{n}} \right)$, where C_λ is defined in Lemma A.3. Then Lemma A.8 implies $\tilde{C} \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \gamma_1 \geq \sqrt{\frac{\text{Var}(\hat{\eta}_\lambda^{TS} | (Y_i^0)_{i=1}^s)}{\delta}}$ for some constant $\tilde{C} > 0$. By (Haggrass et al., 2024, Lemma A.1), if

$$P \left\{ \gamma \geq \zeta - \tilde{C} \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \gamma_1 \right\} \leq \delta, \quad (24)$$

for any $P \in \mathcal{P}$, then we obtain $P\{\hat{\eta}_\lambda^{TS} \geq \gamma\} \geq 1 - 2\delta$. The result follows by taking the infimum over $P \in \mathcal{P}$. Therefore, it remains to verify (24), which we do below. Define $d_\lambda := \mathcal{N}_2(\lambda) + \left(\frac{\sqrt{2}}{c_1} + \frac{1}{\sqrt{2c_1}} \right) \sqrt{\mathcal{N}_1(\lambda)}$, $\gamma_3 := \frac{12d_\lambda(C_1+C_2)}{b_1\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right)$ and $c_2 := B_3 C_4 (C_1+C_2)^{-1}$. Consider

$$\begin{aligned} & P_{H_1} \left\{ \gamma \leq \zeta - \tilde{C} \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \gamma_1 \right\} \\ & \stackrel{(**)}{\geq} P_{H_1} \left\{ \left\{ \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \gamma_3 \leq c_2 \|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^{-2} \|u\|_{L^2(P_0)}^2 - \tilde{C} \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \gamma_1 \right\} \right. \\ & \quad \left. \cap \left\{ \gamma \leq \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \gamma_3 \right\} \right\} \\ & \geq 1 - P_{H_1} \left\{ \frac{\|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 (\tilde{C} \gamma_1 + \gamma_3)}{c_2 \|u\|_{L^2(P_0)}^2} \geq 1 \right\} - P_{H_1} \left\{ \gamma \geq \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \gamma_3 \right\} \\ & \stackrel{(*)}{\geq} P_{H_1} \left\{ \frac{\|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 (\tilde{C} \gamma_1 + \gamma_3)}{c_2 \|u\|_{L^2(P_0)}^2} \leq 1 \right\} - \delta \\ & \stackrel{(\dagger)}{\geq} 1 - P \left\{ \left\{ \|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \leq \frac{3}{2} \right\} \cap \left\{ \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \leq 2 \right\} \right\} - \delta \\ & \geq 1 - P \left\{ \|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \geq \frac{3}{2} \right\} - P \left\{ \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \geq 2 \right\} - \delta \\ & \stackrel{(\ddagger)}{\geq} 1 - 2\delta, \end{aligned}$$

where $(**)$ follows by using $\zeta \geq c_2 \|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^{-2} \|u\|_{L^2(P_0)}^2$, which is obtained by combining (Haggrass et al. 2024, Lemma A.11) by replacing Σ_{PQ} and μ_Q with Σ_0 and μ_0 , respectively) with Lemma A.2 under the assumptions of $u \in \text{Ran}(\mathcal{T}^\theta)$, and (17). Note that $u \in \text{Ran}(\mathcal{T}^\theta)$ is guaranteed since $P \in \mathcal{P}$ and (18) guarantees (17) as discussed in the proof of Theorem 4. $(*)$ follows by Lemma A.6 under the condition $s \geq 32c_1\kappa\lambda^{-1} \log(\max\{17920\kappa^2\lambda^{-1}, 6\}\delta^{-1})$ and $\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} \geq \lambda$ (when $C := \sup_i \|\phi_i\|_\infty < \infty$, the condition can be replaced by $s \geq 32c_1C^2\mathcal{N}_1(\lambda) \log \frac{6}{\delta}$). (\dagger) follows when

$$\|u\|_{L^2(P_0)}^2 \geq \frac{3(\tilde{C}\gamma_1 + \gamma_3)}{c_2}. \quad (25)$$

(\ddagger) follows from (Sriperumbudur and Sterge, 2022, Lemma B.2(ii)) under the assumption that

$$\frac{140\kappa}{s} \log \frac{64\kappa s}{\delta} \leq \lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}, \quad (26)$$

which is implied by $s \geq 280\kappa\lambda^{-1} \log(17920\kappa^2\lambda^{-1}\delta^{-1})$. When $C := \sup_i \|\phi_i\|_\infty < \infty$, (‡) follows from (Haggrass et al., 2024, Lemma A.17) by replacing (26) with

$$136C^2\mathcal{N}_1(\lambda) \log \frac{32\mathcal{N}_1(\lambda)}{\delta} \leq s, \quad \text{and} \quad \lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}. \quad (27)$$

Thus it remains only to verify (25). Using $m \geq n$ and

$$\mathcal{N}_2(\lambda) \leq \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^\infty(\mathcal{H})}^{1/2} \mathcal{N}_1^{1/2}(\lambda),$$

it can be checked that (25) is implied by $\Delta_n \geq \frac{r_1\mathcal{N}_1^{1/2}(\lambda)}{n\sqrt{\alpha}}$, $\Delta_n \geq \frac{r_2C_\lambda}{\delta^2n^2}$ and $\Delta_n \geq \frac{r_3\mathcal{N}_1^{1/2}(\lambda)}{\delta n}$ for some constants $r_1, r_2, r_3 > 0$.

7.10 Proof of Theorem 10

The proof is exactly similar to that of (Haggrass et al., 2024, Theorem 8).

7.11 Proof of Theorem 11

Let $\mathcal{M} = \hat{\Sigma}_{0,\lambda}^{-1/2} \Sigma_0^{1/2}$ and $\gamma_1 = \frac{1}{\sqrt{\delta}} \left(\frac{\sqrt{C_\lambda} \|u\|_{L^2(P_0)} + \mathcal{N}_2(\lambda)}{n} + \frac{C_\lambda^{1/4} \|u\|_{L^2(P_0)}^{3/2} + \|u\|_{L^2(P_0)}}{\sqrt{n}} \right)$. Then following the proof of Theorem 9 in (Haggrass et al., 2024), Lemma A.9 along with $m \geq n$ yields that the power will be controlled to the desired level when $\Delta_n \geq \frac{r_1C_\lambda(\log(1/\alpha))^2}{\delta^2n^2}$ and $\Delta_n \geq \frac{r_2\mathcal{N}_2(\lambda)\log(1/\alpha)}{\delta n}$ for some constants $r_1, r_2 > 0$, and under the condition (26) which can be replaced by (27) when $C := \sup_i \|\phi_i\|_\infty < \infty$.

7.12 Proof of Corollary 12

The proof is similar to that of Corollary 5. Since $\lambda_i \asymp i^{-\beta}$, we have $\mathcal{N}_2(\lambda) \lesssim \lambda^{-1/2\beta}$. By using this bound in the conditions of Theorem 11, we obtain that the conditions on Δ_n hold if

$$\Delta_n \gtrsim \max \left\{ \left(\frac{\delta n}{\log(1/\alpha)} \right)^{-\frac{4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, \left(\frac{\delta n}{\log(1/\alpha)} \right)^{-\frac{8\tilde{\theta}\beta}{4\tilde{\theta}\beta+2\beta+1}} \right\}. \quad (28)$$

By exactly using the same arguments as in the proof of Corollary 5, it is easy to verify that the above condition on Δ_n is implied if

$$\Delta_n = \begin{cases} c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, & \tilde{\theta} > \frac{1}{2} - \frac{1}{4\beta}, \\ c(\alpha, \delta) n^{\frac{-8\tilde{\theta}\beta}{4\tilde{\theta}\beta+2\beta+1}}, & \tilde{\theta} \leq \frac{1}{2} - \frac{1}{4\beta} \end{cases},$$

where $c(\alpha, \delta) \gtrsim \delta^{-2}(\log \frac{1}{\alpha})^2$. On the other hand when $C := \sup_i \|\phi_i\|_\infty < \infty$, we obtain the corresponding condition as

$$\Delta_n \gtrsim \max \left\{ \left(\frac{\delta n}{\log(1/\alpha)} \right)^{-\frac{4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, \left(\frac{\delta n}{\log(1/\alpha)} \right)^{-\frac{4\tilde{\theta}\beta}{2\tilde{\theta}\beta+1}} \right\}, \quad (29)$$

which is implied if

$$\Delta_n = c(\alpha, \delta) n^{\frac{-4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}.$$

7.13 Proof of Corollary 13

The proof is similar to that of Corollary 6. When $\lambda_i \asymp e^{-\tau i}$, we have $\mathcal{N}_2(\lambda) \lesssim \sqrt{\log \frac{1}{\lambda}}$. Thus substituting this in the conditions from Theorem 11 and assuming that

$$n \geq \max\{e^2, \delta^{-1}(\log 1/\alpha)\},$$

we can write the separation boundary as

$$\Delta_n \gtrsim \max \left\{ \left(\frac{\sqrt{2\tilde{\theta}}\delta n}{\log(1/\alpha)\sqrt{\log n}} \right)^{-1}, \left(\frac{\delta n \log(1/\alpha)^{-1}}{\sqrt{\log n}} \right)^{-\frac{4\tilde{\theta}}{2\tilde{\theta}+1}} \right\}, \quad (30)$$

which is implied if

$$\Delta_n = \begin{cases} c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n}, & \tilde{\theta} > \frac{1}{2} \\ c(\alpha, \delta, \theta) \left(\frac{\sqrt{\log n}}{n} \right)^{\frac{4\tilde{\theta}}{2\tilde{\theta}+1}}, & \tilde{\theta} \leq \frac{1}{2} \end{cases},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, 1 \right\} \delta^{-2} (\log \frac{1}{\alpha})^2$.

On the other hand when $C := \sup_i \|\phi_i\|_\infty < \infty$, we obtain

$$\Delta_n \gtrsim \max \left\{ \left(\frac{\sqrt{2\tilde{\theta}}\delta n}{\log(1/\alpha)\sqrt{\log n}} \right)^{-1}, \frac{1}{2\tilde{\theta}} \left(\frac{\delta n}{\log(1/\alpha)\sqrt{\log n}} \right)^{-2} \right\}, \quad (31)$$

which in turn is implied if

$$\Delta_n = c(\alpha, \delta, \theta) \frac{\sqrt{\log n}}{n},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, \frac{1}{2\tilde{\theta}}, 1 \right\} \delta^{-2} (\log \frac{1}{\alpha})^2$.

7.14 Proof of Theorem 14

First note that the following two events,

$$A := \bigcup_{\lambda \in \Lambda} \{ \hat{\eta}_\lambda^{TS} \geq \gamma(\tilde{\alpha}, \lambda) \},$$

and

$$B := \sup_{\lambda \in \Lambda} \frac{\hat{\eta}_\lambda^{TS}}{\hat{\mathcal{N}}_2(\lambda)} \geq \frac{12(C_1 + C_2)}{b_1 \sqrt{\tilde{\alpha}}} \left(\frac{1}{n} + \frac{1}{m} \right)$$

are equivalent, where $\gamma(\alpha, \lambda) = \frac{12(C_1 + C_2)\hat{\mathcal{N}}_2(\lambda)}{b_1 \sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right)$. The proof therefore follows from Theorem 8 and (Hagrass et al., 2024, Lemma A.16).

7.15 Proof Theorem 15

The same steps as in the proof of Theorem 9 will follow, with the only difference being α is replaced by $\frac{\alpha}{|\Lambda|}$, where $|\Lambda| = 1 + \log_2 \frac{\lambda_U}{\lambda_L} \lesssim \log(n)$.

For the case of $\lambda_i \asymp i^{-\beta}$, we can deduce from the proof of Corollary 5 (see (20)) that when $\lambda = d_3^{-1/2\tilde{\theta}} \Delta_{N,M}^{1/2\tilde{\theta}}$ for some $d_3 > 0$, then

$$P_{H_1} \{ \hat{\eta}_\lambda^{TS} \geq \tilde{\gamma} \} \geq 1 - 4\delta,$$

where $\tilde{\gamma} = \frac{12\hat{\mathcal{N}}_2(\lambda)\sqrt{|\Lambda|}(C_1+C_2)}{b_1\sqrt{\alpha}} \left(\frac{1}{n} + \frac{1}{m} \right)$ and the condition on the separation boundary becomes

$$\Delta_n \gtrsim \max \left\{ \left(\frac{n}{\tilde{\alpha}^{-1/2} + \delta^{-1}} \right)^{-\frac{4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, (\delta n)^{-\frac{8\tilde{\theta}\beta}{4\tilde{\theta}\beta+2\beta+1}} \right\},$$

where $\tilde{\alpha} = \frac{\alpha}{|\Lambda|}$. In turn, this is implied if

$$\Delta_n = c(\alpha, \delta) \max \left\{ \left(\frac{n}{\sqrt{\log n}} \right)^{-\frac{4\tilde{\theta}\beta}{4\tilde{\theta}\beta+1}}, n^{-\frac{8\tilde{\theta}\beta}{4\tilde{\theta}\beta+2\beta+1}} \right\},$$

where $c(\alpha, \delta) \gtrsim (\alpha^{-1/2} + \delta^{-2})$. Note that the optimal choice of λ is given by

$$\lambda = \lambda^* := d_3^{-1/2\tilde{\theta}} c(\alpha, \delta, \theta)^{1/2\tilde{\theta}} \max \left\{ \left(\frac{n}{\sqrt{\log n}} \right)^{-\frac{2\beta}{4\theta\beta+1}}, n^{-\frac{4\beta}{4\theta\beta+2\beta+1}} \right\}.$$

Thus it can be verified that for any θ and β , the optimal lambda can be bounded as

$$r_1 n^{\frac{-4\beta_U}{1+2\beta_U}} \leq \lambda \leq r_2 \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-2}{4\tilde{\theta}\beta+1}}$$

for some constants $r_1, r_2 > 0$.

♦ Define $v^* := \sup\{x \in \Lambda : x \leq \lambda^*\}$. From the definition of Λ , it is easy to see that $\lambda_L \leq \lambda^* \leq \lambda_U$ and $\frac{\lambda^*}{2} \leq v^* \leq \lambda^*$. Thus $v^* \in \Lambda$ is an optimal choice of λ that will yield the same form of the separation boundary up to constants. Therefore, by (Hagras et al., 2024, Lemma A.16), for any θ and any P in \mathcal{P} , we have

$$P_{H_1} \left\{ \sup_{\lambda \in \Lambda} \frac{\hat{\eta}_\lambda^{TS}}{\hat{\mathcal{N}}_2(\lambda)} \geq \gamma \right\} \geq 1 - 4\delta.$$

Thus the desired result holds by taking the infimum over $P \in \mathcal{P}$ and θ . ♦

When $\lambda_i \asymp i^{-\beta}$ and $C := \sup_i \|\phi_i\|_\infty < \infty$, then using (21), the conditions on the separation boundary becomes

$$\Delta_n = c(\alpha, \delta) \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-4\tilde{\theta}\beta}{4\theta\beta+1}},$$

where $c(\alpha, \delta) \gtrsim (\alpha^{-1/2} + \delta^{-2})$. This yields the optimal λ to be

$$\lambda^* := d_3^{-1/2\tilde{\theta}} c(\alpha, \delta, \theta)^{1/2\tilde{\theta}} \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-2\beta}{4\theta\beta+1}}.$$

Then as in the previous case we deduce that for any θ and β ,

$$r_3 \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-2\beta_U}{4\theta_l\beta_U+1}} \leq \lambda \leq r_4 \left(\frac{n}{\sqrt{\log n}} \right)^{\frac{-2}{4\xi+1}}$$

for some constants $r_3, r_4 > 0$. The claim therefore follows by using the argument mentioned between \blacklozenge and \blacklozenge .

For the case $\lambda_i \asymp e^{-\tau i}$, $\tau > 0$, the condition on the separation boundary from (22) becomes

$$\Delta_n \gtrsim c(\alpha, \delta, \theta) \max \left\{ \left(\frac{n}{\log n} \right)^{-1}, \left(\frac{n}{\sqrt{\log n}} \right)^{-\frac{4\tilde{\theta}}{2\tilde{\theta}+1}} \right\},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$. Thus

$$\lambda^* = d_3^{-1/2\tilde{\theta}} c(\alpha, \delta, \theta)^{1/2\tilde{\theta}} \max \left\{ \left(\frac{n}{\log n} \right)^{-1/2\tilde{\theta}}, \left(\frac{n}{\sqrt{\log n}} \right)^{-\frac{2}{2\tilde{\theta}+1}} \right\},$$

which can be bounded as $r_5 \left(\frac{n}{\sqrt{\log n}} \right)^{-2} \leq \lambda \leq r_6 \left(\frac{n}{\log n} \right)^{-1/2\xi}$ for some $r_5, r_6 > 0$. Furthermore when $C := \sup_i \|\phi_i\|_\infty < \infty$, the condition on the separation boundary becomes

$$\Delta_n = c(\alpha, \delta, \theta) \frac{\log n}{n},$$

where $c(\alpha, \delta, \theta) \gtrsim \max \left\{ \sqrt{\frac{1}{2\tilde{\theta}}}, \frac{1}{2\tilde{\theta}}, 1 \right\} (\alpha^{-1/2} + \delta^{-2})$. Thus

$$\lambda^* = d_3^{-1/2\tilde{\theta}} c(\alpha, \delta, \theta)^{1/2\tilde{\theta}} \left(\frac{n}{\log n} \right)^{-1/2\tilde{\theta}},$$

which can be bounded by $r_7 \left(\frac{n}{\log n} \right)^{-1/2\theta_l} \leq \lambda \leq r_8 \left(\frac{n}{\log n} \right)^{-1/2\xi}$. The claim, therefore, follows by using the same argument as mentioned in the polynomial decay case.

7.16 Proof of Theorem 16

The proof follows from Theorem 10 and (Hagrass et al., 2024, Lemma A.16) by using $\frac{\alpha}{|\Lambda|}$ in the place of α .

7.17 Proof of Theorem 17

The proof follows from Theorem 11 using $\frac{\alpha}{|\Lambda|}$ instead of α in the expressions (28), (29), (30) and (31), and then bounding the expressions for the resulting optimal λ^* using the ideas similar to that used in the proof of Theorem 15.

Acknowledgments

OH and BKS are partially supported by National Science Foundation (NSF) CAREER award DMS-1945396. BL is supported by NSF grant DMS-2210775.

References

N. Aronszajn. Theory of reproducing kernels. *Trans. Amer. Math. Soc.*, pages 68:337–404, 1950.

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.

F. Bauer, S. Pereverzev, and L. Rosasco. On regularization algorithms in learning theory. *Journal of Complexity*, 23(1):52–72, 2007.

M. V. Burnashev. On the minimax detection of an inaccurately known signal in a white Gaussian noise background. *Theory of Probability & Its Applications*, 24(1):107–119, 1979.

A. Caponnetto and E. D. Vito. Optimal rates for regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7:331–368, 2007.

K. Chwialkowski, H. Strathmann, and A. Gretton. A kernel test of goodness of fit. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2606–2615, New York, New York, USA, 2016. PMLR.

F. Cucker and D. X. Zhou. *Learning Theory: An Approximation Theory Viewpoint*. Cambridge University Press, Cambridge, UK, 2007.

V. de la Peña and E. Giné. *Decoupling: From Dependence to Independence*. Springer Science & Business Media, 2012.

N. Dinculeanu. *Vector Integration and Stochastic Integration in Banach Spaces*. John-Wiley & Sons, Inc., 2000.

P. Drineas and M. W. Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6:2153–2175, December 2005.

H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1996.

G. Fasano and A. Franceschini. A multidimensional version of the Kolmogorov–Smirnov test. *Monthly Notices of the Royal Astronomical Society*, 225(1):155–170, 03 1987. ISSN 0035-8711.

A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A kernel method for the two-sample problem. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems*, volume 19, pages 513–520. MIT Press, 2006.

A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.

O. Hagrass, B. K. Sriperumbudur, and B. Li. Spectral regularized kernel two-sample tests. *The Annals of Statistics*, 52(3):1076–1101, 2024.

W. Hoeffding. A class of statistics with asymptotically normal distribution. In *Breakthroughs in Statistics*, pages 308–334, 1992.

Y. I. Ingster. Minimax testing of nonparametric hypotheses on a distribution density in the L_p metrics. *Theory of Probability & Its Applications*, 31(2):333–337, 1987.

Y. I. Ingster. Asymptotically minimax hypothesis testing for nonparametric alternatives i, ii, iii. *Mathematical Methods of Statistics*, 2(2):85–114, 1993.

I. Kim, S. Balakrishnan, and L. Wasserman. Minimax optimality of permutation tests. *The Annals of Statistics*, 50(1):225–251, 2022.

E. L. Lehmann and J. P. Romano. *Testing Statistical Hypotheses*. Springer Science & Business Media, 2006.

T. Li and M. Yuan. On the optimality of Gaussian kernel based nonparametric tests against smooth alternatives. 2019. <https://arxiv.org/pdf/1909.03302.pdf>.

Q. Liu, J. Lee, and M. Jordan. A kernelized Stein discrepancy for goodness-of-fit tests. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 276–284, New York, New York, USA, 20–22 Jun 2016. PMLR.

S. Mendelson and J. Neeman. Regularization in kernel learning. *The Annals of Statistics*, 38(1):526–565, 2010.

H. Q. Minh, P. Niyogi, and Y. Yao. Mercer’s theorem, feature maps, and smoothing. In G. Lugosi and H. U. Simon, editors, *Learning Theory*, pages 154–168, Berlin, 2006. Springer.

K. Muandet, K. Fukumizu, B. Sriperumbudur, and B. Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2):1–141, 2017.

F. Pesarin and L. Salmaso. *Permutation Tests for Complex Data: Theory, Applications and Software*. John Wiley & Sons, 2010.

C. Puritz, E. Ness-Cohn, and R. Braun. fasano.franceschini.test: An implementation of a multidimensional KS test in R, 2022.

A. Rahimi and B. Recht. Random features for large-scale kernel machines. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1177–1184. Curran Associates, Inc., 2008.

M. Reed and B. Simon. *Methods of Modern Mathematical Physics: Functional Analysis I*. Academic Press, New York, 1980.

A. Schrab, I. Kim, M. Albert, B. Laurent, B. Guedj, and A. Gretton. MMD aggregated two-sample test. 2021. <https://arxiv.org/pdf/2110.15073.pdf>.

R. J. Serfling. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, New York, 2009.

C. Simon-Gabriel and B. Schölkopf. Kernel distribution embeddings: Universal kernels, characteristic kernels and kernel metrics on distributions. *Journal of Machine Learning Research*, 19(44):1–29, 2018.

A. J. Smola, A. Gretton, L. Song, and B. Schölkopf. A Hilbert space embedding for distributions. In M. Hutter, R. A. Servedio, and E. Takimoto, editors, *Algorithmic Learning Theory*, pages 13–31. Springer-Verlag, Berlin, Germany, 2007.

B. K. Sriperumbudur. On the optimal estimation of probability measures in weak and strong topologies. *Bernoulli*, 22(3):1839 – 1893, 2016.

B. K. Sriperumbudur and N. Sterge. Approximate kernel PCA using random features: Computational vs. statistical trade-off. *The Annals of Statistics*, 50(5):2713–2736, 2022.

B. K. Sriperumbudur, K. Fukumizu, A. Gretton, G. R. G. Lanckriet, and B. Schölkopf. Kernel choice and classifiability for RKHS embeddings of probability distributions. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 1750–1758, Cambridge, MA, 2009. MIT Press.

B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.

B. K. Sriperumbudur, K. Fukumizu, and G. R. Lanckriet. Universality, characteristic kernels and RKHS embedding of measures. *Journal of Machine Learning Research*, 12:2389–2410, 2011.

I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, New York, 2008.

I. Steinwart and C. Scovel. Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constructive Approximation*, 35:363–417, 2012.

I. Steinwart, D. R. Hush, and C. Scovel. An explicit description of the reproducing kernel Hilbert spaces of Gaussian RBF kernels. *IEEE Transactions on Information Theory*, 52: 4635–4643, 2006.

G. Szekely and M. Rizzo. Testing for equal distributions in high dimension. *InterStat*, 5, 11 2004.

G. Wahba. *Spline Models for Observational Data*, volume 59. Siam, 1990.

C. K. I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In V. T. T. K. Leen, T. G. Diettrich, editor, *Advances in Neural Information Processing Systems 13*, pages 682–688, Cambridge, MA, 2001. MIT Press.

Y. Yang, M. Pilanci, and M. J. Wainwright. Randomized sketches for kernels: Fast and optimal non-parametric regression. *The Annals of Statistics*, 45(3):991–1023, 2017.

A Technical results

In the following, we present technical results that are used to prove the main results of the paper. Unless specified otherwise, the notation used in this section matches that of the main paper.

Lemma A.1 *Let $(X_i)_{i=1}^n \stackrel{i.i.d.}{\sim} Q$, $(Y_i)_{i=1}^m \stackrel{i.i.d.}{\sim} P$ and $\mathcal{B} : \mathcal{H} \rightarrow \mathcal{H}$ be a bounded operator. Define*

$$I = \frac{2}{nm} \sum_{i,j} \langle a(X_i), b(Y_j) \rangle_{\mathcal{H}},$$

where $a(x) = \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(K(\cdot, x) - \mu_Q)$, $b(x) = \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(K(\cdot, x) - \mu_P)$, $\mu_Q = \int_{\mathcal{X}} K(\cdot, x) dQ(x)$ and $\mu_P = \int_{\mathcal{Y}} K(\cdot, y) dP(y)$. Then

$$(i) \quad \mathbb{E} \langle a(X_i), b(Y_j) \rangle_{\mathcal{H}}^2 \leq \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_Q \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})};$$

$$(ii) \quad \mathbb{E} (I^2) \leq \frac{4}{nm} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_Q \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}.$$

Proof (i) Note that

$$\begin{aligned} \mathbb{E} \langle a(X_i), b(Y_j) \rangle_{\mathcal{H}}^2 &= \mathbb{E} \langle a(X_i) \otimes_{\mathcal{H}} a(X_i), b(Y_j) \otimes_{\mathcal{H}} b(Y_j) \rangle_{\mathcal{L}^2(\mathcal{H})} \\ &= \left\langle \mathcal{B}\Sigma_{0,\lambda}^{-1/2} \Sigma_Q \Sigma_{0,\lambda}^{-1/2} \mathcal{B}^*, \mathcal{B}\Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \mathcal{B}^* \right\rangle_{\mathcal{L}^2(\mathcal{H})} \\ &\leq \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_Q \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}. \end{aligned}$$

(ii) follows by noting that

$$\mathbb{E} (I^2) \stackrel{(\dagger)}{=} \frac{4}{n^2 m^2} \sum_{i,j} \mathbb{E} \langle a(X_i), b(Y_j) \rangle_{\mathcal{H}}^2$$

where (\dagger) follows from (Haggrass et al., 2024, Lemma A.3 (i)), and the result follows from (i). \blacksquare

Lemma A.2 *Let $u = \frac{dP}{dP_0} - 1 \in L^2(P_0)$ and $\eta = \left\| g_{\lambda}^{1/2}(\Sigma_0)(\mu_0 - \mu_P) \right\|_{\mathcal{H}}^2$, where g_{λ} satisfies (A₁)–(A₄). Then*

$$\eta \leq C_1 \|u\|_{L^2(P_0)}^2.$$

Furthermore, if $u \in \text{Ran}(\mathcal{T}^\theta)$, $\theta > 0$ and

$$\|u\|_{L^2(P_0)}^2 \geq \frac{4C_3}{3B_3} \|\mathcal{T}\|_{\mathcal{L}^\infty(L^2(P_0))}^{2\max(\theta-\xi,0)} \lambda^{2\tilde{\theta}} \left\| \mathcal{T}^{-\theta} u \right\|_{L^2(P_0)}^2,$$

where $\tilde{\theta} = \min(\theta, \xi)$, then,

$$\eta \geq \frac{B_3}{4} \|u\|_{L^2(P_0)}^2.$$

Proof The proof uses the same approach as in the proof of (Hagras et al., 2024, Lemma A.7) by noting that $\eta = \langle \mathcal{T}g_\lambda(\mathcal{T})u, u \rangle_{L^2(P_0)}$ and involves replacing μ_Q , R and Σ_{PQ} with μ_0 , P_0 and Σ_0 , respectively. \blacksquare

Lemma A.3 Define $\mathcal{N}_1(\lambda) := \text{Tr}(\Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2})$, $\mathcal{N}_2(\lambda) := \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}$, and $u := \frac{dP}{dP_0} - 1 \in L^2(P_0)$. Then the following hold:

- (i) $\left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}^2 \leq 4C_\lambda \|u\|_{L^2(P_0)}^2 + 2\mathcal{N}_2^2(\lambda);$
- (ii) $\left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^\infty(\mathcal{H})} \leq 2\sqrt{C_\lambda} \|u\|_{L^2(P_0)} + 1,$

where

$$C_\lambda = \begin{cases} \mathcal{N}_1(\lambda) \sup_i \|\phi_i\|_\infty^2, & \sup_i \|\phi_i\|_\infty^2 < \infty \\ \frac{2\mathcal{N}_2(\lambda)}{\lambda} \sup_x \|K(\cdot, x)\|_{\mathcal{H}}^2, & \text{otherwise} \end{cases}.$$

Proof The proof is similar to that of (Hagras et al., 2024, Lemma A.9) and involves replacing R with P_0 and Σ_{PQ} with Σ_0 . \blacksquare

Lemma A.4 For any $0 < \delta < \frac{1}{2}$,

$$P_0^s \left\{ (Y_i)_{i=1}^s : \left\| \Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} \leq \frac{32\kappa \log \frac{3}{\delta}}{\lambda s} + \sqrt{\frac{16\kappa \mathcal{N}_1(\lambda) \log \frac{2}{\delta}}{\lambda s}} \right\} \geq 1 - 2\delta.$$

Furthermore, suppose $C := \sup_i \|\phi_i\|_\infty < \infty$. Then

$$\begin{aligned} P_0^s \left\{ (Y_i)_{i=1}^s : \left\| \Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} \leq \frac{32C^2 \mathcal{N}_1(\lambda) \log \frac{3}{\delta}}{s} + \sqrt{\frac{16C^2 \mathcal{N}_1^2(\lambda) \log \frac{2}{\delta}}{s}} \right\} \\ \geq 1 - 2\delta. \end{aligned}$$

Proof Define $s(x) := K(\cdot, x)$, $A(x, y) := \frac{1}{\sqrt{2}}(s(x) - s(y))$, $U(x, y) := \Sigma_{0,\lambda}^{-1/2} A(x, y)$, and $Z(x, y) = U(x, y) \otimes_{\mathcal{H}} U(x, y)$. Then

$$\Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} = \frac{1}{s(s-1)} \sum_{i \neq j} Z(Y_i, Y_j) - \mathbb{E}(Z(X, Y)).$$

Also,

$$\sup_{x,y} \|Z(x,y)\|_{\mathcal{L}^2(\mathcal{H})} = \sup_{x,y} \|U(x,y)\|_{\mathcal{H}}^2 = \frac{1}{2} \sup_{x,y} \left\| \Sigma_{0,\lambda}^{-1/2} (s(x) - s(y)) \right\|_{\mathcal{H}}^2 \leq \frac{2\kappa}{\lambda}.$$

Define $\zeta(x) := \mathbb{E}_Y[Z(x, Y)]$. Then

$$\begin{aligned} \mathbb{E} \|\zeta(X) - \Sigma_0\|_{\mathcal{L}^2(\mathcal{H})}^2 &\leq \mathbb{E} \|\zeta(X)\|_{\mathcal{L}^2(\mathcal{H})}^2 = \mathbb{E} \left\| \Sigma_{0,\lambda}^{-1/2} \mathbb{E}_Y[A(X, Y) \otimes_{\mathcal{H}} A(X, Y)] \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}^2 \\ &= \mathbb{E} \text{Tr} \left(\Sigma_{0,\lambda}^{-1/2} \mathbb{E}_Y[A(X, Y) \otimes_{\mathcal{H}} A(X, Y)] \Sigma_{0,\lambda}^{-1} \mathbb{E}_Y[A(X, Y) \otimes_{\mathcal{H}} A(X, Y)] \Sigma_{0,\lambda}^{-1/2} \right) \\ &\leq \sup_x \|\zeta(x)\|_{\mathcal{L}^\infty(\mathcal{H})} \text{Tr}(\Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}) \\ &\leq \sup_{x,y} \|U(x,y)\|_{\mathcal{H}}^2 \mathcal{N}_1(\lambda) \leq \frac{2\kappa \mathcal{N}_1(\lambda)}{\lambda}. \end{aligned}$$

When $C := \sup_i \|\phi_i\|_\infty < \infty$, we can use the same approach as in the proof of (Haggrass et al., 2024, Lemma A.17) to show that $\sup_{x,y} \|U(x,y)\|_{\mathcal{H}}^2 \leq 2C^2 \mathcal{N}_1(\lambda)$ which in turn yields that

$$\sup_{x,y} \|Z(x,y)\|_{\mathcal{L}^2(\mathcal{H})} \leq 2C^2 \mathcal{N}_1(\lambda),$$

and

$$\mathbb{E} \|\zeta(X) - \Sigma_0\|_{\mathcal{L}^2(\mathcal{H})}^2 \leq 2C^2 \mathcal{N}_1^2(\lambda).$$

Then the result follows from (Sriperumbudur and Sterge, 2022, Theorem D.3(ii)). \blacksquare

Lemma A.5 *Let $I = \left\langle \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}, \Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} \right\rangle_{\mathcal{L}^2(\mathcal{H})}$. Then for any $\delta > 0$,*

$$P_0^s \left\{ (Y_i)_{i=1}^s : |I| \leq \frac{4\kappa \log \frac{2}{\delta}}{\lambda s} + \sqrt{\frac{12\kappa \mathcal{N}_2^2(\lambda) \log \frac{2}{\delta}}{\lambda s}} \right\} \geq 1 - \delta.$$

Furthermore, suppose $C := \sup_i \|\phi_i\|_\infty < \infty$. Then

$$P_0^s \left\{ (Y_i)_{i=1}^s : |I| \leq \frac{4C^2 \mathcal{N}_1(\lambda) \log \frac{2}{\delta}}{s} + \sqrt{\frac{12C^2 \mathcal{N}_1(\lambda) \mathcal{N}_2^2(\lambda) \log \frac{2}{\delta}}{s}} \right\} \geq 1 - \delta.$$

Proof Define $s(x) := K(\cdot, x)$, $A(x, y) := \frac{1}{\sqrt{2}}(s(x) - s(y))$, $U(x, y) := \Sigma_{0,\lambda}^{-1/2} A(x, y)$, $Z(x, y) := U(x, y) \otimes_{\mathcal{H}} U(x, y)$, $B := \mathbb{E} Z(X, Y) = \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}$, and

$$\tilde{Z}(x, y) := \langle B, Z(x, y) \rangle_{\mathcal{L}^2(\mathcal{H})}.$$

Then

$$I = \frac{1}{s(s-1)} \sum_{i \neq j} \tilde{Z}(Y_i, Y_j) - \mathbb{E} \tilde{Z}(X, Y).$$

Moreover,

$$\begin{aligned} \sup_{x,y} |\tilde{Z}(x,y)| &= \sup_{x,y} |\langle B, Z(x,y) \rangle_{\mathcal{L}^2(\mathcal{H})}| \\ &= \sup_{x,y} |\text{Tr}(B(U(x,y) \otimes_{\mathcal{H}} U(x,y)))| \leq \|B\|_{\mathcal{L}^\infty(\mathcal{H})} \sup_{x,y} \|U(x,y)\|_{\mathcal{H}}^2 \leq \frac{2\kappa}{\lambda}, \end{aligned}$$

and

$$\mathbb{E} \tilde{Z}^2(X, Y) \stackrel{(*)}{\leq} \sup_{x,y} |\tilde{Z}(x,y)| \mathbb{E} \langle B, Z(X, Y) \rangle_{\mathcal{L}^2(\mathcal{H})} \leq \frac{2\kappa \mathcal{N}_2^2(\lambda)}{\lambda},$$

where $(*)$ follows by using $\tilde{Z}(x,y) \geq 0$, which can be shown by writing,

$$\begin{aligned} \tilde{Z}(x,y) &= \langle B, Z(x,y) \rangle = \text{Tr}(B(U(x,y) \otimes_{\mathcal{H}} U(x,y))) \\ &= \text{Tr} \left(\Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2} [U(x,y) \otimes_{\mathcal{H}} U(x,y)] \right) \\ &= \text{Tr}(\Sigma_0^{1/2} \Sigma_{0,\lambda}^{-1/2} (U(x,y) \otimes_{\mathcal{H}} U(x,y)) \Sigma_{0,\lambda}^{-1/2} \Sigma_0^{1/2}) \\ &= \left\| \Sigma_0^{1/2} \Sigma_{0,\lambda}^{-1/2} U(x,y) \right\|_{\mathcal{H}}^2 \geq 0. \end{aligned}$$

When $C := \sup_i \|\phi_i\|_\infty < \infty$, we can use the same approach as in (Hagras et al., 2024, Lemma A.17) to show that $\sup_{x,y} \|U(x,y)\|_{\mathcal{H}}^2 \leq 2C^2 \mathcal{N}_1(\lambda)$ which in turn yields that

$$\sup_{x,y} |\tilde{Z}(x,y)| \leq 2C^2 \mathcal{N}_1(\lambda),$$

and

$$\mathbb{E} \tilde{Z}^2(X, Y) \leq 2C^2 \mathcal{N}_1(\lambda) \mathcal{N}_2^2(\lambda).$$

Thus the result follows by using Hoeffding's inequality as stated in (de la Peña and Giné, 2012, Theorem 4.1.8). \blacksquare

Lemma A.6 For any $c_1 > 0$, $\delta > 0$ and $\frac{32c_1\kappa}{s} \log \frac{3}{\delta} \leq \lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$, we have

$$P_0^s \left\{ (Y_i)_{i=0}^s : \widehat{\mathcal{N}}_2(\lambda) \leq \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \left(\mathcal{N}_2(\lambda) + \left(\frac{\sqrt{2}}{c_1} + \frac{1}{\sqrt{2c_1}} \right) \sqrt{\mathcal{N}_1(\lambda)} \right) \right\} \geq 1 - 2\delta.$$

Furthermore if $C := \sup_i \|\phi_i\|_\infty < \infty$, the above bound holds for $32c_1C^2 \mathcal{N}_1(\lambda) \log \frac{3}{\delta} \leq s$ and $\lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$.

Proof Let $\mathcal{M} = \hat{\Sigma}_{0,\lambda}^{-1/2} \Sigma_{0,\lambda}^{1/2}$. Then

$$\begin{aligned} \widehat{\mathcal{N}}_2(\lambda) &\leq \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \left\| \Sigma_{0,\lambda}^{-1/2} \hat{\Sigma}_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} \\ &\leq \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \left(\mathcal{N}_2(\lambda) + \left\| \Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} \right). \end{aligned}$$

From Lemma A.4 and the assumption that $\frac{32c_1\kappa}{s} \log \frac{3}{\delta} \leq \lambda$, we have with probability at least $1 - 2\delta$ that,

$$\begin{aligned} \left\| \Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})} &\leq \frac{1}{c_1} + \frac{1}{\sqrt{2c_1}} \sqrt{\mathcal{N}_1(\lambda)} \\ &\stackrel{(*)}{\leq} \left(\frac{\sqrt{2}}{c_1} + \frac{1}{\sqrt{2c_1}} \right) \sqrt{\mathcal{N}_1(\lambda)}, \end{aligned}$$

where in $(*)$ we used $\sqrt{\mathcal{N}_1(\lambda)} = \sqrt{\sum_i \frac{\lambda_i}{\lambda_i + \lambda}} \geq \sqrt{\frac{\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}}{\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} + \lambda}} \stackrel{(\dagger)}{\geq} \frac{1}{\sqrt{2}}$, and (\dagger) follows from $\lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$. Similarly when $C := \sup_i \|\phi_i\|_\infty < \infty$, the same bound holds by Lemma A.4 for $32c_1C^2\mathcal{N}_1(\lambda) \log \frac{3}{\delta} \leq s$. \blacksquare

Lemma A.7 *For any $c_1 > 0$, $\delta > 0$, and $\max\{\frac{140\kappa}{s} \log \frac{16\kappa s}{\delta}, \frac{4c_1\kappa}{s} \log \frac{2}{\delta}\} \leq \lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$, we have*

$$P_0^s \left\{ (Y_i)_{i=0}^s : \hat{\mathcal{N}}_2^2(\lambda) \geq \left(\frac{4}{9} - \frac{16}{3\sqrt{3c_1}} - \frac{32}{9c_1} \right) \mathcal{N}_2^2(\lambda) \right\} \geq 1 - 3\delta.$$

Furthermore if $C := \sup_i \|\phi_i\|_\infty < \infty$, the above bound holds for $\lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$, and $C^2\mathcal{N}_1(\lambda) \max\{4c_1 \log \frac{2}{\delta}, 136 \log \frac{8\mathcal{N}_1(\lambda)}{\delta}\} \leq s$.

Proof Let $\mathcal{M} := \hat{\Sigma}_{0,\lambda}^{-1/2} \Sigma_{0,\lambda}^{1/2}$ and $I := \left\langle \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}, \Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} \right\rangle_{\mathcal{L}^2(\mathcal{H})}$. Then

$$\begin{aligned} \hat{\mathcal{N}}_2^2(\lambda) &\geq \frac{1}{\|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^4} \left\| \Sigma_{0,\lambda}^{-1/2} \hat{\Sigma}_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}^2 \\ &= \frac{1}{\|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^4} \left(\left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}^2 + 2I + \left\| \Sigma_{0,\lambda}^{-1/2} (\hat{\Sigma}_0 - \Sigma_0) \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}^2 \right) \\ &\geq \frac{1}{\|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^4} (\mathcal{N}_2^2(\lambda) - 2|I|). \end{aligned}$$

Then from Lemma A.5 and the assumption $\frac{4\kappa c_1}{s} \log \frac{2}{\delta} \leq \lambda$ (similarly when $C := \sup_i \|\phi_i\|_\infty < \infty$ the same bound holds by Lemma A.5 for $32c_1C^2\mathcal{N}_1(\lambda) \log \frac{3}{\delta} \leq s$), we have with probability at least $1 - \delta$,

$$|I| \leq \frac{1}{c_1} + \sqrt{\frac{3}{c_1}} \mathcal{N}_2(\lambda) \leq \left(\sqrt{\frac{3}{c_1}} + \frac{2}{c_1} \right) \mathcal{N}_2(\lambda) \leq 2 \left(\sqrt{\frac{3}{c_1}} + \frac{2}{c_1} \right) \mathcal{N}_2^2(\lambda),$$

where in the last two inequalities we used $\mathcal{N}_2(\lambda) \geq \frac{\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}}{\|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})} + \lambda} \geq \frac{1}{2}$ when $\lambda \leq \|\Sigma_0\|_{\mathcal{L}^\infty(\mathcal{H})}$.

Define

$$S_1 := \left\{ (Y_i)_{i=1}^s : |I| \leq 2 \left(\sqrt{\frac{3}{c_1}} + \frac{2}{c_1} \right) \mathcal{N}_2^2(\lambda) \right\}$$

and

$$S_2 := \left\{ (Y_i)_{i=1}^s : \|\mathcal{M}^{-1}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \leq \frac{9}{4} \right\}.$$

Then,

$$\begin{aligned} P_0^s \left\{ (Y_i)_{i=0}^s : \widehat{\mathcal{N}}_2^2(\lambda) \geq \left(\frac{4}{9} - \frac{16}{3\sqrt{3c_1}} - \frac{32}{9c_1} \right) \mathcal{N}_2^2(\lambda) \right\} &\geq P(S_1 \cap S_2) \\ &\geq 1 - P(S'_1) - P(S'_2) \geq 1 - 3\delta, \end{aligned}$$

where we used (Sriperumbudur and Sterge, 2022, Lemma B.2 (iii)) in the last inequality, with S' being the complement of set S . Similarly, when $C := \sup_i \|\phi_i\|_\infty < \infty$ the same bound holds using (Hagras et al., 2024, Lemma A.17 (iii)). \blacksquare

Lemma A.8 *Let $\zeta = \left\| g_\lambda^{1/2}(\hat{\Sigma}_0)(\mu_P - \mu_0) \right\|_{\mathcal{H}}^2$, $\mathcal{M} = \hat{\Sigma}_{0,\lambda}^{-1/2} \Sigma_{0,\lambda}^{1/2}$, and $m \geq n$. Then*

$$\begin{aligned} \mathbb{E} [(\hat{\eta}_\lambda^{TS} - \zeta)^2 | (Y_i^0)_{i=1}^s] &\leq \tilde{C} \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left\{ \frac{C_\lambda \|u\|_{L^2(P_0)}^2 + \mathcal{N}_2^2(\lambda)}{n^2} + \frac{\sqrt{C_\lambda} \|u\|_{L^2(P_0)}^3 + \|u\|_{L^2(P_0)}^2}{n} \right\}, \end{aligned}$$

where C_λ is defined in Lemma A.3 and \tilde{C} is a constant that depends only on C_1 and C_2 . Furthermore, if $P = P_0$, then

$$\mathbb{E} [(\hat{\eta}_\lambda^{TS})^2 | (Y_i^0)_{i=1}^s] \leq 6(C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \mathcal{N}_2^2(\lambda) \left(\frac{1}{m^2} + \frac{1}{n^2} \right).$$

Proof Define $a(x) = \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(K(\cdot, x) - \mu_P)$, and $b(x) = \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(K(\cdot, x) - \mu_0)$, where $\mathcal{B} = g_\lambda^{1/2}(\hat{\Sigma}_0)\Sigma_{0,\lambda}^{1/2}$. Then replacing Σ_{PQ} by Σ_0 and μ_Q by μ_0 in the proof of (Hagras et al., 2024, Lemma A.12), it can be shown that

$$\begin{aligned} \hat{\eta}_\lambda^{TS} - \zeta &= \underbrace{\frac{1}{n(n-1)} \sum_{i \neq j} \langle a(X_i), a(X_j) \rangle_{\mathcal{H}}}_{(1)} + \underbrace{\frac{1}{m(m-1)} \sum_{i \neq j} \langle b(X_i^0), b(X_j^0) \rangle_{\mathcal{H}}}_{(2)} \\ &\quad + \underbrace{\frac{2}{m} \sum_{i=1}^m \langle b(X_i^0), \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(\mu_0 - \mu_P) \rangle_{\mathcal{H}}}_{(3)} - \underbrace{\frac{2}{nm} \sum_{i,j} \langle a(X_i), b(X_j^0) \rangle_{\mathcal{H}}}_{(4)} \\ &\quad - \underbrace{\frac{2}{n} \sum_{i=1}^n \langle a(X_i), \mathcal{B}\Sigma_{0,\lambda}^{-1/2}(\mu_0 - \mu_P) \rangle_{\mathcal{H}}}_{(5)} \end{aligned}$$

and

$$\|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})} \leq (C_1 + C_2)^{1/2} \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}.$$

Next, we bound each of these terms using Lemmas A.1, A.3 and (Hagrass et al., 2024, Lemma A.4, Lemma A.5). It follows from Lemma A.3(i) and (Hagrass et al., 2024, Lemma A.4(ii)) that

$$\begin{aligned}\mathbb{E} \left(\bigcircled{1}^2 | (Y_i^0)_{i=1}^s \right) &\leq \frac{4}{n^2} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left\| \Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2} \right\|_{\mathcal{L}^2(\mathcal{H})}^2 \\ &\leq \frac{4}{n^2} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left(4C_\lambda \|u\|_{L^2(P_0)}^2 + 2\mathcal{N}_2^2(\lambda) \right),\end{aligned}$$

and

$$\mathbb{E} \left(\bigcircled{2}^2 | (Y_i^0)_{i=1}^s \right) \leq \frac{4}{m^2} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \mathcal{N}_2^2(\lambda).$$

Using Lemma A.3(ii) and (Hagrass et al., 2024, Lemma A.5), we obtain

$$\begin{aligned}\mathbb{E} \left(\bigcircled{3}^2 | (Y_i^0)_{i=1}^s \right) &\leq \frac{4}{m} \|\Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}\|_{\mathcal{L}^\infty(\mathcal{H})} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \|\Sigma_{0,\lambda}^{-1/2} (\mu_P - \mu_0)\|_{\mathcal{H}}^2 \\ &\leq \frac{4}{m} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \|\Sigma_{0,\lambda}^{-1/2} (\mu_P - \mu_0)\|_{\mathcal{H}}^2 \\ &\stackrel{(*)}{\leq} \frac{4}{m} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \|u\|_{L^2(P_0)}^2,\end{aligned}$$

and

$$\begin{aligned}\mathbb{E} \left(\bigcircled{5}^2 | (Y_i^0)_{i=1}^s \right) &\leq \frac{4}{n} \|\Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2}\|_{\mathcal{L}^\infty(\mathcal{H})} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \|\Sigma_{0,\lambda}^{-1/2} (\mu_0 - \mu_P)\|_{\mathcal{H}}^2 \\ &\leq \frac{4}{n} (1 + 2\sqrt{C_\lambda} \|u\|_{L^2(P_0)}) (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \|\Sigma_{0,\lambda}^{-1/2} (\mu_0 - \mu_P)\|_{\mathcal{H}}^2 \\ &\stackrel{(*)}{\leq} \frac{4}{n} (1 + 2\sqrt{C_\lambda} \|u\|_{L^2(P_0)}) (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \|u\|_{L^2(P_0)}^2,\end{aligned}$$

where (*) follows from using $g_\lambda(x) = (x + \lambda)^{-1}$ with $C_1 = 1$ in Lemma A.2. For term (4), using Lemmas A.1 and A.3, we have

$$\begin{aligned}\mathbb{E} \left(\bigcircled{4}^2 | (Y_i^0)_{i=1}^s \right) &\leq \frac{4}{nm} \|\mathcal{B}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \|\Sigma_{0,\lambda}^{-1/2} \Sigma_P \Sigma_{0,\lambda}^{-1/2}\|_{\mathcal{L}^2(\mathcal{H})} \|\Sigma_{0,\lambda}^{-1/2} \Sigma_0 \Sigma_{0,\lambda}^{-1/2}\|_{\mathcal{L}^2(\mathcal{H})} \\ &\leq \frac{4}{nm} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 (2\sqrt{C_\lambda} \|u\|_{L^2(P_0)} + \sqrt{2}\mathcal{N}_2(\lambda)) \mathcal{N}_2(\lambda) \\ &\leq \frac{4}{nm} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 (2\sqrt{C_\lambda} \mathcal{N}_2(\lambda) \|u\|_{L^2(P_0)} + \sqrt{2}\mathcal{N}_2^2(\lambda)).\end{aligned}$$

Combining these bounds with the fact that $\sqrt{ab} \leq \frac{a}{2} + \frac{b}{2}$, and that $(\sum_{i=1}^k a_k)^2 \leq k \sum_{i=1}^k a_k^2$ for any $a, b, a_k \in \mathbb{R}$, $k \in \mathbb{N}$ yields that

$$\begin{aligned}\mathbb{E} \left[(\hat{\eta}_\lambda^{TS} - \zeta)^2 | (Y_i^0)_{i=1}^s \right] &\lesssim \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left(\frac{C_\lambda \|u\|_{L^2(P_0)}^2 + \mathcal{N}_2^2(\lambda)}{n^2} + \frac{\sqrt{C_\lambda} \|u\|_{L^2(P_0)}^3 + \|u\|_{L^2(P_0)}^2}{n} \right. \\ &\quad \left. + \frac{\mathcal{N}_2^2(\lambda)}{m^2} + \frac{\|u\|_{L^2(P_0)}^2}{m} \right) \\ &\lesssim \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \left(\frac{C_\lambda \|u\|_{L^2(P_0)}^2 + \mathcal{N}_2^2(\lambda)}{n^2} + \frac{\sqrt{C_\lambda} \|u\|_{L^2(P_0)}^3 + \|u\|_{L^2(P_0)}^2}{n} \right),\end{aligned}$$

where in the last inequality we used $m \geq n$.

When $P = P_0$, and using the same lemmas as above, we have

$$\begin{aligned}\mathbb{E} \left(\bigcircled{1}^2 | (Y_i^0)_{i=1}^s \right) &\leq \frac{4}{n^2} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \mathcal{N}_2^2(\lambda), \\ \mathbb{E} \left(\bigcircled{2}^2 | (Y_i^0)_{i=1}^s \right) &\leq \frac{4}{m^2} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \mathcal{N}_2^2(\lambda), \\ \mathbb{E} \left(\bigcircled{4}^2 | (Y_i^0)_{i=1}^s \right) &\leq \frac{4}{nm} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \mathcal{N}_2^2(\lambda),\end{aligned}$$

and $\bigcircled{3} = \bigcircled{5} = 0$. Therefore,

$$\begin{aligned}\mathbb{E}[(\hat{\eta}_\lambda^{TS})^2 | (Y_i^0)_{i=1}^s] &= \mathbb{E} \left[(\bigcircled{1} + \bigcircled{2} + \bigcircled{4})^2 | (Y_i^0)_{i=1}^s \right] \stackrel{(*)}{=} \mathbb{E} \left(\bigcircled{1}^2 + \bigcircled{2}^2 + \bigcircled{4}^2 | (Y_i^0)_{i=1}^s \right) \\ &\stackrel{(\dagger)}{\leq} (C_1 + C_2)^2 \|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^4 \mathcal{N}_2^2(\lambda) \left(\frac{6}{m^2} + \frac{6}{n^2} \right),\end{aligned}$$

where $(*)$ follows by noting that $\mathbb{E}(\bigcircled{1} \cdot \bigcircled{2}) = \mathbb{E}(\bigcircled{1} \cdot \bigcircled{4}) = \mathbb{E}(\bigcircled{2} \cdot \bigcircled{4}) = 0$ under the assumption $P = P_0$, and (\dagger) follows using $\sqrt{ab} \leq \frac{a}{2} + \frac{b}{2}$. \blacksquare

Lemma A.9 *For $0 < \alpha \leq e^{-1}$, $\delta > 0$ and $m \geq n$, there exists a constant $C_5 > 0$ such that*

$$P_{H_1}(q_{1-\alpha}^\lambda \leq C_5 \gamma) \geq 1 - \delta,$$

where

$$\gamma = \frac{\|\mathcal{M}\|_{\mathcal{L}^\infty(\mathcal{H})}^2 \log \frac{1}{\alpha}}{\sqrt{\delta n}} \left(\sqrt{C_\lambda} \|u\|_{L^2(P_0)} + \mathcal{N}_2(\lambda) + C_\lambda^{1/4} \|u\|_{L^2(P_0)}^{3/2} + \|u\|_{L^2(P_0)} \right) + \frac{\zeta \log \frac{1}{\alpha}}{\sqrt{\delta n}},$$

$$\zeta = \left\| g_\lambda^{1/2}(\hat{\Sigma}_0)(\mu_0 - \mu_P) \right\|_{\mathcal{H}}^2, \text{ and } C_\lambda \text{ is defined in Lemma A.3.}$$

Proof The proof is similar to that of (Hagras et al., 2024, Lemma A.15) and involves replacing Σ_{PQ} with Σ_0 , R with P_0 , and μ_Q with μ_0 . Then the desired result follows by using $m \geq n$. \blacksquare