
Interpretable Distribution Features with Maximum Testing Power

Wittawat Jitkrittum, Zoltán Szabó, Kacper Chwialkowski, Arthur Gretton

wittawatj@gmail.com

zoltan.szabo.m@gmail.com

kacper.chwialkowski@gmail.com

arthur.gretton@gmail.com

Gatsby Unit, University College London

Abstract

Two semimetrics on probability distributions are proposed, given as the sum of differences of expectations of analytic functions evaluated at spatial or frequency locations (i.e., features). The features are chosen so as to maximize the distinguishability of the distributions, by optimizing a lower bound on test power for a statistical test using these features. The result is a parsimonious and interpretable indication of how and where two distributions differ locally. We show that the empirical estimate of the test power criterion converges with increasing sample size, ensuring the quality of the returned features. In real-world benchmarks on high-dimensional text and image data, linear-time tests using the proposed semimetrics achieve comparable performance to the state-of-the-art quadratic-time maximum mean discrepancy test, while returning human-interpretable features that explain the test results.

1 Introduction

We address the problem of discovering features of distinct probability distributions, with which they can most easily be distinguished. The distributions may be in high dimensions, can differ in non-trivial ways (i.e., not simply in their means), and are observed only through i.i.d. samples. One application for such divergence measures is to model criticism, where samples from a trained model are compared with a validation sample: in the univariate case, through the KL divergence (Cinzia Carota and Polson, 1996), or in the multivariate case, by use of the maximum mean discrepancy (MMD) (Lloyd and Ghahramani, 2015). An alternative, interpretable analysis of a multivariate difference in distributions may be obtained by projecting onto a discriminative direction, such that the Wasserstein distance on this projection is maximized (Mueller and Jaakkola, 2015). Note that both recent works require low dimensionality, either explicitly (in the case of Lloyd and Ghahramani, the function becomes difficult to plot in more than two dimensions), or implicitly in the case of Mueller and Jaakkola, in that a large difference in distributions must occur in projection along a particular one-dimensional axis. Distances between distributions in high dimensions may be more subtle, however, and it is of interest to find interpretable, distinguishing features of these distributions.

In the present paper, we take a hypothesis testing approach to discovering features which best distinguish two multivariate probability measures P and Q , as observed by samples $X := \{\mathbf{x}_i\}_{i=1}^n$ drawn independently and identically (i.i.d.) from P , and $Y := \{\mathbf{y}_i\}_{i=1}^n \subset \mathbb{R}^d$ from Q . Non-parametric two-sample tests based on RKHS distances (Eric et al., 2008; Fromont et al., 2012; Gretton et al., 2012a) or energy distances (Székely and Rizzo, 2004; Baringhaus and Franz, 2004) have as their test statistic an integral probability metric, the Maximum Mean Discrepancy (Gretton et al., 2012a; Sejdinovic et al., 2013). For this metric, a smooth witness function is computed, such that the amplitude is largest where the probability mass differs most (e.g. Gretton et al., 2012a,

Figure 1). Lloyd and Ghahramani (2015) used this witness function to compare the model output of the Automated Statistician (Lloyd et al., 2014) with a reference sample, yielding a visual indication of where the model fails. In high dimensions, however, the witness function cannot be plotted, and is less helpful. Furthermore, the witness function does not give an easily interpretable result for distributions with local differences in their characteristic functions. A more subtle shortcoming is that it does not provide a direct indication of the distribution features which, when compared, would maximize test power - rather, it is the witness function *norm*, and (broadly speaking) its *variance* under the null, that determine test power.

Our approach builds on the analytic representations of probability distributions of Chwialkowski et al. (2015), where differences in expectations of analytic functions at particular spatial or frequency locations are used to construct a two-sample test statistic, which can be computed in linear time. Despite the differences in these analytic functions being evaluated at random locations, the analytic tests have greater power than linear time tests based on subsampled estimates of the MMD (Gretton et al., 2012b; Zaremba et al., 2013). Our first theoretical contribution, in Sec. 3, is to derive a lower bound on the test power, which can be maximized over the choice of test locations. We propose two novel tests, both of which significantly outperform the random feature choice of Chwialkowski et al.. The (ME) test evaluates the difference of mean embeddings at locations chosen to maximize the test power lower bound (i.e., spatial features); unlike the maxima of the MMD witness function, these features are directly chosen to maximize the distinguishability of the distributions, and take variance into account. The Smooth Characteristic Function (SCF) test uses as its statistic the difference of the two smoothed empirical characteristic functions, evaluated at points in the frequency domain so as to maximize the same criterion (i.e., frequency features). Optimization of the mean embedding kernels/frequency smoothing functions themselves is achieved on a held-out data set with the same consistent objective.

As our second theoretical contribution in Sec. 3, we prove that the empirical estimate of the test power criterion asymptotically converges to its population quantity uniformly over the class of Gaussian kernels. Two important consequences follow: first, in testing, we obtain a more powerful test with fewer features. Second, we obtain a parsimonious and interpretable set of features that best distinguish the probability distributions. In Sec. 4, we provide experiments demonstrating that the proposed linear-time tests greatly outperform all previous linear time tests, and achieve performance that compares to or exceeds the more expensive quadratic-time MMD test (Gretton et al., 2012a). Moreover, the new tests discover features of text data (NIPS proceedings) and image data (distinct facial expressions) which have a clear human interpretation, thus validating our feature elicitation procedure in these challenging high-dimensional testing scenarios.

2 ME and SCF tests

In this section, we review the ME and SCF tests (Chwialkowski et al., 2015) for two-sample testing. In Sec. 3, we will extend these approaches to learn features that optimize the power of these tests. Given two samples $X := \{\mathbf{x}_i\}_{i=1}^n, Y := \{\mathbf{y}_i\}_{i=1}^n \subset \mathbb{R}^d$ independently and identically distributed (i.i.d.) according to P and Q , respectively, the goal of a two-sample test is to decide whether P is different from Q on the basis of the samples. The task is formulated as a statistical hypothesis test proposing a null hypothesis $H_0 : P = Q$ (samples are drawn from the same distribution) against an alternative hypothesis $H_1 : P \neq Q$ (the sample generating distributions are different). A test calculates a test statistic $\hat{\lambda}_n$ from X and Y , and rejects H_0 if $\hat{\lambda}_n$ exceeds a predetermined test threshold (critical value). The threshold is given by the $(1 - \alpha)$ -quantile of the (asymptotic) distribution of $\hat{\lambda}_n$ under H_0 i.e., the null distribution, and α is the significance level of the test.

ME test The ME test uses as its test statistic $\hat{\lambda}_n$, a form of Hotelling’s T-squared statistic, defined as $\hat{\lambda}_n := n\bar{\mathbf{z}}_n^\top \mathbf{S}_n^{-1} \bar{\mathbf{z}}_n$, where $\bar{\mathbf{z}}_n := \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i$, $\mathbf{S}_n := \frac{1}{n-1} \sum_{i=1}^n (\mathbf{z}_i - \bar{\mathbf{z}}_n)(\mathbf{z}_i - \bar{\mathbf{z}}_n)^\top$, and $\mathbf{z}_i := (k(\mathbf{x}_i, \mathbf{v}_j) - k(\mathbf{y}_i, \mathbf{v}_j))_{j=1}^J \in \mathbb{R}^J$. The statistic depends on a positive definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ (with $\mathcal{X} \subseteq \mathbb{R}^d$), and a set of J test locations $\mathcal{V} = \{\mathbf{v}_j\}_{j=1}^J \subset \mathbb{R}^d$. Under H_0 , asymptotically $\hat{\lambda}_n$ follows $\chi^2(J)$, a chi-squared distribution with J degrees of freedom. The ME test rejects H_0 if $\hat{\lambda}_n > T_\alpha$, where the test threshold T_α is given by the $(1 - \alpha)$ -quantile of the asymptotic null distribution $\chi^2(J)$. Although the distribution of $\hat{\lambda}_n$ under H_1 was not derived, Chwialkowski et al. (2015) showed that if k is analytic, integrable and characteristic (in the sense of Sriperumbudur et al. (2011)), under H_1 , $\hat{\lambda}_n$ can be arbitrarily large as $n \rightarrow \infty$, allowing the test to correctly reject H_0 .

One can intuitively think of the ME test statistic as a squared normalized (by the inverse covariance \mathbf{S}_n^{-1}) $L^2(\mathcal{X}, V_J)$ distance of the mean embeddings (Smola et al., 2007) of the empirical measures $P_n := \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{x}_i}$, and $Q_n := \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{y}_i}$ where $V_J := \frac{1}{J} \sum_{i=1}^J \delta_{\mathbf{v}_i}$, and $\delta_{\mathbf{x}}$ is the Dirac measure concentrated at \mathbf{x} . The unnormalized counterpart (i.e., without \mathbf{S}_n^{-1}) was shown by Chwialkowski et al. (2015) to be a metric on the space of probability measures for any \mathcal{V} . Both variants behave similarly for two-sample testing, with the normalized version being a semimetric having a more computationally tractable null distribution, i.e., $\chi^2(J)$.

SCF test The SCF uses the test statistic which has the same form as the ME test statistic with a modified $\mathbf{z}_i := [\hat{l}(\mathbf{x}_i) \sin(\mathbf{x}_i^\top \mathbf{v}_j) - \hat{l}(\mathbf{y}_i) \sin(\mathbf{y}_i^\top \mathbf{v}_j), \hat{l}(\mathbf{x}_i) \cos(\mathbf{x}_i^\top \mathbf{v}_j) - \hat{l}(\mathbf{y}_i) \cos(\mathbf{y}_i^\top \mathbf{v}_j)]_{j=1}^J \in \mathbb{R}^{2J}$, where $\hat{l}(\mathbf{x}) = \int_{\mathbb{R}^d} \exp(-i\mathbf{u}^\top \mathbf{x}) l(\mathbf{u}) d\mathbf{u}$ is the Fourier transform of $l(\mathbf{x})$, and $l : \mathbb{R}^d \rightarrow \mathbb{R}$ is an analytic translation-invariant kernel i.e., $l(\mathbf{x} - \mathbf{y})$ defines a positive definite kernel for \mathbf{x} and \mathbf{y} . In contrast to the ME test defining the statistic in terms of spatial locations, the locations $\mathcal{V} = \{\mathbf{v}_j\}_{j=1}^J \subset \mathbb{R}^d$ in the SCF test are in the frequency domain. As a brief description, let $\varphi_P(\mathbf{w}) := \mathbb{E}_{\mathbf{x} \sim P} \exp(i\mathbf{w}^\top \mathbf{x})$ be the characteristic function of P . Define a smooth characteristic function as $\phi_P(\mathbf{v}) = \int_{\mathbb{R}^d} \varphi_P(\mathbf{w}) l(\mathbf{v} - \mathbf{w}) d\mathbf{w}$ (Chwialkowski et al., 2015, Definition 2). Then, similar to the ME test, the statistic defined by the SCF test can be seen as a normalized (by \mathbf{S}_n^{-1}) version of $L^2(\mathcal{X}, V_J)$ distance of empirical $\phi_P(\mathbf{v})$ and $\phi_Q(\mathbf{v})$. The SCF test statistic has asymptotic distribution $\chi^2(2J)$ under H_0 . We will use J' to refer to the degrees of freedom of the chi-squared distribution i.e., $J' = J$ for the ME test, and $J' = 2J$ for the SCF test.

In this work, we modify the statistic with a regularization parameter $\gamma_n > 0$, giving $\hat{\lambda}_n := n\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n$, for stability of the matrix inverse. Using multivariate Slutsky's theorem, under H_0 , $\hat{\lambda}_n$ still asymptotically follows $\chi^2(J')$ provided that $\gamma_n \rightarrow 0$ as $n \rightarrow \infty$.

3 Lower bound on test power, consistency of empirical power statistic

This section contains our main results. We propose to optimize the test locations \mathcal{V} and kernel parameters (jointly referred to as θ) by maximizing a lower bound on the test power in Proposition 1. This criterion offers a simple objective function for fast parameter tuning. The bound may be of independent interest in other Hotelling's T-squared statistics, since apart from the Gaussian case (e.g. Bilodeau and Brenner, 2008, Ch. 8), the characterization of such statistics under the alternative distribution is challenging. The optimization procedure is given in Sec. 4. We use $\mathbb{E}_{\mathbf{x}\mathbf{y}}$ as a shorthand for $\mathbb{E}_{\mathbf{x} \sim P} \mathbb{E}_{\mathbf{y} \sim Q}$ and let $\|\cdot\|_F$ be the Frobenius norm.

Proposition 1 (Lower bound on ME test power). *Let \mathcal{K} be a uniformly bounded (i.e., $\exists B < \infty$ such that $\sup_{k \in \mathcal{K}} \sup_{(\mathbf{x}, \mathbf{y}) \in \mathcal{X}^2} |k(\mathbf{x}, \mathbf{y})| \leq B$) family of $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ measurable kernels. Let \mathbb{V} be a collection in which each element is a set of J test locations. Assume that $\tilde{c} := \sup_{\mathcal{V} \in \mathbb{V}, k \in \mathcal{K}} \|\Sigma^{-1}\|_F < \infty$. Then, the test power $\mathbb{P}(\hat{\lambda}_n \geq T_\alpha)$ of the ME test satisfies $\mathbb{P}(\hat{\lambda}_n \geq T_\alpha) \geq L(\lambda_n)$ where*

$$L(\lambda_n) := 1 - 2e^{-\xi_1(\lambda_n - T_\alpha)^2/n} - 2e^{-\frac{[\gamma_n(\lambda_n - T_\alpha)(n-1) - \xi_2 n]^2}{\xi_3 n(2n-1)^2}} - 2e^{-[(\lambda_n - T_\alpha)/3 - \bar{c}_3 n \gamma_n]^2 \gamma_n^2 / \xi_4},$$

and $\bar{c}_3, \xi_1, \dots, \xi_4$ are positive constants depending on only B, J and \tilde{c} . The parameter $\lambda_n := n\boldsymbol{\mu}^\top \Sigma^{-1} \boldsymbol{\mu}$ is the population counterpart of $\hat{\lambda}_n := n\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n$ where $\boldsymbol{\mu} = \mathbb{E}_{\mathbf{x}\mathbf{y}}[\mathbf{z}_1]$ and $\Sigma = \mathbb{E}_{\mathbf{x}\mathbf{y}}[(\mathbf{z}_1 - \boldsymbol{\mu})(\mathbf{z}_1 - \boldsymbol{\mu})^\top]$. For large n , $L(\lambda_n)$ is increasing in λ_n .

Proof (sketch). The idea is to construct a bound for $|\hat{\lambda}_n - \lambda_n|$ which involves bounding $\|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2$ and $\|\mathbf{S}_n - \Sigma\|_F$ separately using Hoeffding's inequality. The result follows after a reparameterization of the bound on $\mathbb{P}(|\hat{\lambda}_n - \lambda_n| \geq t)$ to have $\mathbb{P}(\hat{\lambda}_n \geq T_\alpha)$. See Sec. F for details. \square

Proposition 1 suggests that for large n it is sufficient to maximize λ_n to maximize a lower bound on the ME test power. The same conclusion holds for the SCF test (result omitted due to space constraints). Assume that k is characteristic (Sriperumbudur et al., 2011). It can be shown that $\lambda_n = 0$ if and only if $P = Q$ i.e., λ_n is a semimetric for P and Q . In this sense, one can see λ_n as encoding the ease of rejecting H_0 . The higher λ_n , the easier for the test to correctly reject H_0 when H_1 holds. This observation justifies the use of λ_n as a maximization objective for parameter tuning.

Contributions The statistic $\hat{\lambda}_n$ for both ME and SCF tests depends on a set of test locations \mathcal{V} and a kernel parameter σ . We propose to set $\theta := \{\mathcal{V}, \sigma\} = \arg \max_{\theta} \lambda_n = \arg \max_{\theta} \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$. The optimization of θ brings two benefits: first, it significantly increases the probability of rejecting H_0 when H_1 holds; second, the learned test locations act as discriminative features allowing an interpretation of how the two distributions differ. We note that optimizing parameters by maximizing a test power proxy (Gretton et al., 2012b) is valid under both H_0 and H_1 as long as the data used for parameter tuning and for testing are disjoint. If H_0 holds, then $\theta = \arg \max 0$ is arbitrary. Since the test statistic asymptotically follows $\chi^2(J')$ for any θ , the optimization does not change the null distribution. Also, the rejection threshold T_α depends on only J' and is independent of θ .

To avoid creating a dependency between θ and the data used for testing (which would affect the null distribution), we split the data into two disjoint sets. Let $D := (X, Y)$ and $D^{tr}, D^{te} \subset D$ such that $D^{tr} \cap D^{te} = \emptyset$ and $D^{tr} \cup D^{te} = D$. In practice, since $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are unknown, we use $\hat{\lambda}_{n/2}^{tr}$ in place of λ_n , where $\hat{\lambda}_{n/2}^{tr}$ is the test statistic computed on the training set D^{tr} . For simplicity, we assume that each of D^{tr} and D^{te} has half of the samples in D . We perform an optimization of θ with gradient ascent algorithm on $\hat{\lambda}_{n/2}^{tr}(\theta)$. The actual two-sample test is performed using the test statistic $\hat{\lambda}_{n/2}^{te}(\theta)$ computed on D^{te} . The full procedure from tuning the parameters to the actual two-sample test is summarized in Sec. A.

Since we use an empirical estimate $\hat{\lambda}_{n/2}^{tr}$ in place of λ_n for parameter optimization, we give a finite-sample bound in Theorem 2 guaranteeing the convergence of $\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n$ to $\boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$ as n increases, uniformly over all kernels $k \in \mathcal{K}$ (a family of uniformly bounded kernels) and all test locations in an appropriate class. Kernel classes satisfying conditions of Theorem 2 include the widely used isotropic Gaussian kernel class $\mathcal{K}_g = \{k_\sigma : (\mathbf{x}, \mathbf{y}) \mapsto \exp(-(2\sigma^2)^{-1} \|\mathbf{x} - \mathbf{y}\|^2) \mid \sigma > 0\}$, and the more general full Gaussian kernel class $\mathcal{K}_{\text{full}} = \{k : (\mathbf{x}, \mathbf{y}) \mapsto \exp(-(\mathbf{x} - \mathbf{y})^\top \mathbf{A}(\mathbf{x} - \mathbf{y})) \mid \mathbf{A} \text{ is positive definite}\}$ (see Lemma 5 and Lemma 6).

Theorem 2 (Consistency of $\hat{\lambda}_n$ in the ME test). *Let $\mathcal{X} \subseteq \mathbb{R}^d$ be a measurable set, and \mathbb{V} be a collection in which each element is a set of J test locations. All suprema over \mathcal{V} and k are to be understood as $\sup_{\mathcal{V} \in \mathbb{V}}$ and $\sup_{k \in \mathcal{K}}$ respectively. For a class of kernels \mathcal{K} on $\mathcal{X} \subseteq \mathbb{R}^d$, define*

$$\mathcal{F}_1 := \{\mathbf{x} \mapsto k(\mathbf{x}, \mathbf{v}) \mid k \in \mathcal{K}, \mathbf{v} \in \mathcal{X}\}, \quad \mathcal{F}_2 := \{\mathbf{x} \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{x}, \mathbf{v}') \mid k \in \mathcal{K}, \mathbf{v}, \mathbf{v}' \in \mathcal{X}\}, \quad (1)$$

$$\mathcal{F}_3 := \{(\mathbf{x}, \mathbf{y}) \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{y}, \mathbf{v}') \mid k \in \mathcal{K}, \mathbf{v}, \mathbf{v}' \in \mathcal{X}\}. \quad (2)$$

Assume that (1) \mathcal{K} is a uniformly bounded (by B) family of $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ measurable kernels, (2) $\bar{c} := \sup_{\mathcal{V}, k} \|\boldsymbol{\Sigma}^{-1}\|_F < \infty$, and (3) $\mathcal{F}_i = \{f_{\theta_i} \mid \theta_i \in \Theta_i\}$ is VC-subgraph with VC-index $VC(\mathcal{F}_i)$, and $\theta \mapsto f_{\theta_i}(m)$ is continuous ($\forall m, i = 1, 2, 3$). Let $\bar{c}_1 := 4B^2 J \sqrt{J} \bar{c}$, $\bar{c}_2 := 4B \sqrt{J} \bar{c}$, and $\bar{c}_3 := 4B^2 J \bar{c}^2$. Let C_i -s ($i = 1, 2, 3$) be the universal constants associated to \mathcal{F}_i -s according to Theorem 2.6.7 in van der Vaart and Wellner (2000). Then for any $\delta \in (0, 1)$ with probability at least $1 - \delta$,

$$\begin{aligned} & \sup_{\mathcal{V}, k} \left| \bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} \right| \\ & \leq 2T_{\mathcal{F}_1} \left(\frac{2}{\gamma_n} \bar{c}_1 B J \frac{2n-1}{n-1} + \bar{c}_2 \sqrt{J} \right) + \frac{2}{\gamma_n} \bar{c}_1 J (T_{\mathcal{F}_2} + T_{\mathcal{F}_3}) + \frac{8}{\gamma_n} \frac{\bar{c}_1 B^2 J}{n-1} + \bar{c}_3 \gamma_n, \text{ where} \\ T_{\mathcal{F}_j} &= \frac{16\sqrt{2}B^{\zeta_j}}{\sqrt{n}} \left(2\sqrt{\log[C_j \times VC(\mathcal{F}_j)(16e)^{VC(\mathcal{F}_j)}]} + \frac{\sqrt{2\pi[VC(\mathcal{F}_j) - 1]}}{2} \right) + B^{\zeta_j} \sqrt{\frac{2\log(5/\delta)}{n}}, \end{aligned}$$

for $j = 1, 2, 3$ and $\zeta_1 = 1, \zeta_2 = \zeta_3 = 2$.

Proof (sketch). The idea is to lower bound the difference with an expression involving $\sup_{\mathcal{V}, k} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2$ and $\sup_{\mathcal{V}, k} \|\mathbf{S}_n - \boldsymbol{\Sigma}\|_F$. These two quantities can be seen as suprema of empirical processes, and can be bounded by Rademacher complexities of their respective function classes (i.e., $\mathcal{F}_1, \mathcal{F}_2$, and \mathcal{F}_3). Finally, the Rademacher complexities can be upper bounded using Dudley entropy bound and VC subgraph properties of the function classes. Proof details are given in Sec. D. \square

Theorem 2 implies that if we set $\gamma_n = \mathcal{O}(n^{-1/4})$, then we have $\sup_{\mathcal{V}, k} \left| \bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} \right| = \mathcal{O}_p(n^{-1/4})$ as the rate of convergence. Both

Proposition 1 and Theorem 2 require $\tilde{c} := \sup_{\mathcal{V} \in \mathbb{V}, k \in \mathcal{K}} \|\Sigma^{-1}\|_F < \infty$ as a precondition. To guarantee that $\tilde{c} < \infty$, a concrete construction of \mathcal{K} is the isotropic Gaussian kernel class \mathcal{K}_g , where σ is constrained to be in a compact set. Also, consider $\mathbb{V} := \{\mathcal{V} \mid \text{any two locations are at least } \epsilon \text{ distance apart, and all test locations have their norms bounded by } \zeta\}$ for some $\epsilon, \zeta > 0$. Then, for any non-degenerate P, Q , we have $\tilde{c} < \infty$ since $(\sigma, \mathcal{V}) \mapsto \lambda_n$ is continuous, and thus attains its supremum over compact sets \mathcal{K} and \mathbb{V} .

4 Experiments

In this section, we demonstrate the effectiveness of the proposed methods on both toy and real problems. We consider the isotropic Gaussian kernel class \mathcal{K}_g in all kernel-based tests. We study seven two-sample test algorithms. For the SCF test, we set $\hat{l}(\mathbf{x}) = k(\mathbf{x}, \mathbf{0})$. Denote by ME-full and SCF-full the ME and SCF tests whose test locations and the Gaussian width σ are fully optimized using gradient ascent on a separate training sample (D^{tr}) of the same size as the test set (D^{te}). ME-grid and SCF-grid are as in Chwialkowski et al. (2015) where only the Gaussian width is optimized by a grid search,¹ and the test locations are randomly drawn from a multivariate normal distribution. MMD-quad (quadratic-time) and MMD-lin (linear-time) refer to the nonparametric tests based on maximum mean discrepancy of Gretton et al. (2012a), where to ensure a fair comparison, the Gaussian kernel width is also chosen so as to maximize a criterion for the test power on training data, following the same principle as (Gretton et al., 2012b). For MMD-quad, since its null distribution is given by an infinite sum of weighted chi-squared variables (no closed-form quantiles), in each trial we randomly permute the two samples 400 times to approximate the null distribution. Finally, T^2 is the standard two-sample Hotelling’s T-squared test, which serves as a baseline with Gaussian assumptions on P and Q .


In all the following experiments, each problem is repeated for 500 trials. For toy problems, new samples are generated from the specified P, Q distributions in each trial. For real problems, samples are partitioned randomly into training and test sets in each trial. In all of the simulations, we report an empirical estimate of $\mathbb{P}(\hat{\lambda}_{n/2}^{te} \geq T_\alpha)$ which is the proportion of the number of times the statistic $\hat{\lambda}_{n/2}^{te}$ is above T_α . This quantity is an estimate of type-I error under H_0 , and corresponds to test power when H_1 is true. We set $\alpha = 0.01$ in all the experiments. All the code and preprocessed data are available at <https://github.com/wittawatj/interpretable-test>.

Optimization The parameter tuning objective $\hat{\lambda}_{n/2}^{tr}(\theta)$ is a function of θ consisting of one real-valued σ and J test locations each of d dimensions. The parameters θ can thus be regarded as a $Jd + 1$ Euclidean vector. We take the derivative of $\hat{\lambda}_{n/2}^{tr}(\theta)$ with respect to θ , and use gradient ascent to maximize it. J is pre-specified and fixed. For the ME test, we initialize the test locations with realizations from two multivariate normal distributions fitted to samples from P and Q ; this ensures that the initial locations are well supported by the data. For the SCF test, initialization using the standard normal distribution is found to be sufficient. The parameter γ_n is not optimized; we set the regularization parameter γ_n to be as small as possible while being large enough to ensure that $(S_n + \gamma_n I)^{-1}$ can be stably computed. We emphasize that both the optimization and testing are linear in n . The testing cost $\mathcal{O}(J^3 + J^2n + dJn)$ and the optimization costs $\mathcal{O}(J^3 + dJ^2n)$ per gradient ascent iteration. Runtimes of all methods are reported in Sec. C in the appendix.

1. Informative features: simple demonstration We begin with a demonstration that the proxy $\hat{\lambda}_{n/2}^{tr}(\theta)$ for the test power is informative for revealing the difference of the two samples in the ME

Table 1: Four toy problems. H_0 holds only in SG.

Data	P	Q
SG	$\mathcal{N}(\mathbf{0}_d, I_d)$	$\mathcal{N}(\mathbf{0}_d, I_d)$
GMD	$\mathcal{N}(\mathbf{0}_d, I_d)$	$\mathcal{N}((1, 0, \dots, 0)^\top, I_d)$
GVD	$\mathcal{N}(\mathbf{0}_d, I_d)$	$\mathcal{N}(\mathbf{0}_d, \text{diag}(2, 1, \dots, 1))$
Blobs	Gaussian mixtures in \mathbb{R}^2 as studied in Chwialkowski et al. (2015); Gretton et al. (2012b).	



¹Chwialkowski et al. (2015) chooses the Gaussian width that minimizes the median of the p-values, a heuristic that does not directly address test power. Here, we perform a grid search to choose the best Gaussian width by maximizing $\hat{\lambda}_{n/2}^{tr}$ as done in ME-full and SCF-full.




Figure 2: Plots of type-I error/test power against the test sample size n^{te} in the four toy problems.

test. We consider the Gaussian Mean Difference (GMD) problem (see Table 1), where both P and Q are two-dimensional normal distributions with the difference in means. We use $J = 2$ test locations \mathbf{v}_1 and \mathbf{v}_2 , where \mathbf{v}_1 is fixed to the location indicated by the black triangle in Fig. 1. The contour plot shows $\mathbf{v}_2 \mapsto \hat{\lambda}_{n/2}^{tr}(\mathbf{v}_1, \mathbf{v}_2)$.

Fig. 1 (top) suggests that $\hat{\lambda}_{n/2}^{tr}$ is maximized when \mathbf{v}_2 is placed in either of the two regions that captures the difference of the two samples i.e., the region in which the probability masses of P and Q have less overlap. Fig. 1 (bottom), we consider placing \mathbf{v}_1 in one of the two key regions. In this case, the contour plot shows that \mathbf{v}_2 should be placed in the other region to maximize $\hat{\lambda}_{n/2}^{tr}$, implying that placing multiple test locations in the same neighborhood will not increase the discriminability. The two modes on the left and right suggest two ways to place the test location in a region that reveals the difference. The non-convexity of the $\hat{\lambda}_{n/2}^{tr}$ is an indication of many informative ways to detect differences of P and Q , rather than a drawback. A convex objective would not capture this multimodality.

2. Test power vs. sample size n We now demonstrate the rate of increase of test power with sample size. When the null hypothesis holds, the type-I error stays at the specified level α . We consider the following four toy problems: Same Gaussian (SG), Gaussian mean difference (GMD), Gaussian variance difference (GVD), and Blobs. The specifications of P and Q are summarized in Table. 1. In the Blobs problem, P and Q are defined as a mixture of Gaussian distributions arranged on a 4×4 grid in \mathbb{R}^2 . This problem is challenging as the difference of P and Q is encoded at a much smaller length scale compared to the global structure (Gretton et al., 2012b). Specifically, the eigenvalue ratio for the covariance of each Gaussian distribution is 2.0 in P , and 1.0 in Q . We set $J = 5$ in this experiment.

The results are shown in Fig. 2 where type-I error (for SG problem), and test power (for GMD, GVD and Blobs problems) are plotted against test sample size. A number of observations are worth noting. In the SG problem, we see that the type-I error roughly stays at the specified level: the rate of rejection of H_0 when it is true is roughly at the specified level $\alpha = 0.01$.

GMD with 100 dimensions turns out to be an easy problem for all the tests except MMD-lin. In the GVD and Blobs cases, ME-full and SCF-full achieve substantially higher test power than ME-grid and SCF-grid, respectively, suggesting a clear advantage from optimizing the test locations. Remarkably, ME-full consistently outperforms the quadratic-time MMD across all test sample sizes in the GVD case. When the difference of P and Q is subtle as in the Blobs problem, ME-grid, which uses randomly drawn test locations, can perform poorly (see Fig. 2d) since it is unlikely that randomly drawn locations will be placed in the key regions that reveal the difference. In this case, optimization of the test locations can considerably boost the test power (see ME-full in Fig. 2d). Note also that SCF variants perform significantly better than ME variants on the Blobs problem, as the difference in P and Q is localized in the frequency domain; ME-full and ME-grid would require many more test locations in the spatial domain to match the test powers of the SCF variants. For the same reason, SCF-full does much better than the quadratic-time MMD across most sample sizes, as the latter represents a weighted distance between characteristic functions integrated across the entire frequency domain (Sriperumbudur et al., 2010, Corollary 4).




Figure 1: A contour plot of $\hat{\lambda}_{n/2}^{tr}$ as a function of \mathbf{v}_2 when $J = 2$ and \mathbf{v}_1 is fixed (black triangle). The objective $\hat{\lambda}_{n/2}^{tr}$ is high in the regions that reveal the difference of the two samples.




Figure 3: Plots of type-I error/test power against the dimensions d in the four toy problems in Table 1.

Table 2: Type-I errors and powers of various tests in the problem of distinguishing NIPS papers from two categories. $\alpha = 0.01$. $J = 1$. n_{te} denotes the test sample size of each of the two samples.

Problem	n_{te}	ME-full	ME-grid	SCF-full	SCF-grid	MMD-quad	MMD-lin
Bayes-Bayes	215	.012	.018	.012	.004	.022	.008
Bayes-Deep	216	.954	.034	.688	.180	.906	.262
Bayes-Learn	138	.990	.774	.836	.534	1.00	.238
Bayes-Neuro	394	1.00	.300	.828	.500	.952	.972
Learn-Deep	149	.956	.052	.656	.138	.876	.500
Learn-Neuro	146	.960	.572	.590	.360	1.00	.538

3. Test power vs. dimension d We next investigate how the dimension (d) of the problem can affect type-I errors and test powers of ME and SCF tests. We consider the same artificial problems: SG, GMD and GVD. This time, we fix the test sample size to 10000, set $J = 5$, and vary the dimension. The results are shown in Fig. 3. Due to the large dimensions and sample size, it is computationally infeasible to run MMD-quad.

We observe that all the tests except the T-test can maintain type-I error at roughly the specified significance level $\alpha = 0.01$ as dimension increases. The type-I performance of the T-test is incorrect at large d because of the difficulty in accurately estimating the covariance matrix in high dimensions. It is interesting to note the high performance of ME-full in the GMD problem in Fig. 3b. ME-full achieves the maximum test power of 1.0 throughout and matches the power T-test, in spite of being nonparametric and making no assumption on P and Q (the T-test is further advantaged by its excessive Type-I error). However, this is true only with optimization of the test locations. This is reflected in the test power of ME-grid in Fig. 3b which drops monotonically as dimension increases, highlighting the importance of test location optimization. The performance of MMD-lin degrades quickly with increasing dimension, as expected from Ramdas et al. (2015).

4. Distinguishing articles from two categories We now turn to performance on real data. We first consider the problem of distinguishing two categories of publications at the conference on Neural Information Processing Systems (NIPS). Out of 5903 papers published in NIPS from 1988 to 2015, we manually select disjoint subsets related to Bayesian inference (Bayes), neuroscience (Neuro), deep learning (Deep), and statistical learning theory (Learn) (see Sec. B). Each paper is represented as a bag of words using TF-IDF (Manning et al., 2008) as features. We perform stemming, remove all stop words, and retain only nouns. A further filtering of document-frequency (DF) of words that satisfies $5 \leq DF \leq 2000$ yields approximately 5000 words from which 2000 words (i.e., $d = 2000$ dimensions) are randomly selected. See Sec. B for more details on the preprocessing. For ME and SCF tests, we use only one test location i.e., set $J = 1$. We perform 1000 permutations to approximate the null distribution of MMD-quad in this and the following experiments.

Type-I errors and test powers are summarized in Table. 2. The first column indicates the categories of the papers in the two samples. In Bayes-Bayes problem, papers on Bayesian inference are randomly partitioned into two samples in each trial. This task represents a case in which H_0 holds. Among all the linear-time tests, we observe that ME-full has the highest test power in all the tasks, attaining a maximum test power of 1.0 in the Bayes-Neuro problem. This high performance assures that although different test locations \mathcal{V} may be selected in different trials, these locations are each informative. It is interesting to observe that ME-full has performance close to or better than MMD-quad, which requires $O(n^2)$ runtime complexity. Besides clear advantages of interpretability and linear runtime of the proposed tests, these results suggest that evaluating the differences in expectations of analytic functions at particular locations can yield an equally powerful test at a much lower cost, as opposed to

Table 3: Type-I errors and powers in the problem of distinguishing positive (+) and negative (-) facial expressions. $\alpha = 0.01$. $J = 1$.

Problem	n^{te}	ME-full	ME-grid	SCF-full	SCF-grid	MMD-quad	MMD-lin
\pm vs. \pm	201	.010	.012	.014	.002	.018	.008
$+$ vs. $-$	201	.998	.656	1.00	.750	1.00	.578

computing the RKHS norm of the witness function as done in MMD. Unlike Blobs, however, Fourier features are less powerful in this setting.

We further investigate the interpretability of the ME test by the following procedure. For the learned test location $\mathbf{v}^t \in \mathbb{R}^d$ ($d = 2000$) in trial t , we construct $\tilde{\mathbf{v}}^t = (\tilde{v}_1^t, \dots, \tilde{v}_d^t)$ such that $\tilde{v}_j^t = |v_j^t|$. Let $\eta_j^t \in \{0, 1\}$ be an indicator variable taking value 1 if \tilde{v}_j^t is among the top five largest for all $j \in \{1, \dots, d\}$, and 0 otherwise. Define $\eta_j := \sum_t \eta_j^t$ as a proxy indicating the significance of word j i.e., η_j is high if word j is frequently among the top five largest as measured by \tilde{v}_j^t . The top seven words as sorted in descending order by η_j in the Bayes-Neuro problem are *spike*, *markov*, *cortex*, *dropout*, *recur*, *iii*, *gibb*, showing that the learned test locations are highly interpretable. Indeed, “markov” and “gibb” (i.e., stemmed from Gibbs) are discriminative terms in Bayesian inference category, and “spike” and “cortex” are key terms in neuroscience. We give full lists of discriminative terms learned in all the problems in Sec. B.1. To show that not all the randomly selected 2000 terms are informative, if the definition of η_j^t is modified to consider the least important words (i.e., η_j is high if word j is frequently among the top five smallest as measured by \tilde{v}_j^t), we instead obtain *circumfer*, *bra*, *dominiqu*, *rhino*, *mitra*, *kid*, *impostor*, which are not discriminative.

5. Distinguishing positive and negative emotions In the final experiment, we study how well ME and SCF tests can distinguish two samples of photos of people showing positive and negative facial expressions. Our emphasis is on the discriminative features of the faces identified by ME test showing how the two groups differ. For this purpose, we use Karolinska Directed Emotional Faces (KDEF) dataset (Lundqvist et al., 1998) containing 5040 aligned face images of 70 amateur actors, 35 females and 35 males. We use only photos showing front views of the faces. In the dataset, each actor displays seven expressions: happy (HA), neutral (NE), surprised (SU), sad (SA), afraid (AF), angry (AN), and disgusted (DI). We assign HA, NE, and SU faces into the positive emotion group (i.e., samples from P), and AF, AN and DI faces into the negative emotion group (samples from Q). We denote this problem as “+ vs. -”. Examples of six facial expressions from one actor are shown in Fig. 4. Photos of the SA group are unused to keep the sizes of the two samples the same. Each image of size 562×762 pixels is cropped to exclude the background, resized to $48 \times 34 = 1632$ pixels (d), and converted to grayscale.




Figure 4: (a)-(f): Six facial expressions of actor AM05 in the KDEF data. (g): Average across trials of the learned test locations \mathbf{v}_1 .

We run the tests 500 times with the same setting used previously i.e., Gaussian kernels, and $J = 1$. The type-I errors and test powers are shown in Table 3. In the table, “ \pm vs. \pm ” is a problem in which all faces expressing the six emotions are randomly split into two samples of equal sizes i.e., H_0 is true. Both ME-full and SCF-full achieve high test powers while maintaining the correct type-I errors.

As a way to interpret how positive and negative emotions differ, we take an average across trials of the learned test locations of ME-full in the “+ vs. -” problem. This average is shown in Fig. 4g. We see that the test locations faithfully capture the difference of positive and negative emotions by giving more weights to the regions of nose, upper lip, and nasolabial folds (smile lines), confirming the interpretability of the test in a high-dimensional setting.

Acknowledgement

We thank the Gatsby Charitable Foundation for the financial support.

References

- L. Baringhaus and C. Franz. On a new multivariate two-sample test. *Journal of Multivariate Analysis*, 88: 190–206, 2004.
- M. Bilodeau and D. Brenner. *Theory of multivariate statistics*. Springer Science & Business Media, 2008.
- S. Bird, E. Klein, and E. Loper. *Natural Language Processing with Python*. O’Reilly Media, 1st edition, 2009.
- O. Bousquet. New approaches to statistical learning theory. *Annals of the Institute of Statistical Mathematics*, 55:371–389, 2003.
- K. Chwialkowski, A. Ramdas, D. Sejdinovic, and A. Gretton. Fast two-sample testing with analytic representations of probability measures. In *NIPS*, pages 1972–1980, 2015.
- G. P. Cinzia Carota and N. G. Polson. Diagnostic measures for model criticism. *Journal of the American Statistical Association*, 91(434):753–762, 1996.
- M. Eric, F. R. Bach, and Z. Harchaoui. Testing for homogeneity with kernel Fisher discriminant analysis. In *NIPS*, pages 609–616, 2008.
- M. Fromont, B. Laurent, M. Lerasle, and P. Reynaud-Bouret. Kernels based tests with non-asymptotic bootstrap approaches for two-sample problems. In *COLT*, pages 23.1–23.22, 2012.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012a.
- A. Gretton, D. Sejdinovic, H. Strathmann, S. Balakrishnan, M. Pontil, K. Fukumizu, and B. K. Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In *NIPS*, pages 1205–1213, 2012b.
- M. R. Kosorok. *Introduction to Empirical Processes and Semiparametric Inference*. Springer, 2008.
- J. R. Lloyd and Z. Ghahramani. Statistical model criticism using kernel two sample tests. In *NIPS*, pages 829–837, 2015.
- J. R. Lloyd, D. Duvenaud, R. Grosse, J. B. Tenenbaum, and Z. Ghahramani. Automatic construction and Natural-Language description of nonparametric regression models. In *AAAI*, pages 1242–1250, 2014.
- D. Lundqvist, A. Flykt, and A. Öhman. The Karolinska directed emotional faces-KDEF. Technical report, ISBN 91-630-7164-9, 1998.
- C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to information retrieval*. Cambridge University Press, 2008.
- J. Mueller and T. Jaakkola. Principal differences analysis: Interpretable characterization of differences between distributions. In *NIPS*, pages 1693–1701, 2015.
- A. Ramdas, S. Jakkam Reddi, B. Póczos, A. Singh, and L. Wasserman. On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions. In *AAAI*, pages 3571–3577, 2015.
- D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Annals of Statistics*, 41(5):2263–2291, 2013.
- A. Smola, A. Gretton, L. Song, and B. Schölkopf. A Hilbert space embedding for distributions. In *ALT*, pages 13–31, 2007.
- N. Srebro and S. Ben-David. Learning bounds for support vector machines with learned kernels. In *COLT*, pages 169–183, 2006.
- B. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schoelkopf, and 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. *The Journal of Machine Learning Research*, 12:2389–2410, 2011.
- I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, 2008.
- G. Székely and M. Rizzo. Testing for equal distributions in high dimension. *InterStat*, (5), 2004.
- A. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics (Springer Series in Statistics)*. Springer, 2000.
- W. Zaremba, A. Gretton, and M. Blaschko. B-test: A non-parametric, low variance kernel two-sample test. In *NIPS*, pages 755–763, 2013.

Interpretable Distribution Features with Maximum Testing Power

Supplementary Material

A Algorithm

The full algorithm for the proposed tests from parameter tuning to the actual two-sample testing is given in Algorithm 1.

Algorithm 1 Optimizing parameters and testing

Input: Two samples X, Y , significance level α , and number of test locations J

- 1: Split $D := (X, Y)$ into disjoint training and test sets, D^{tr} and D^{te} , of the same size n^{te} .
 - 2: Optimize parameters $\theta = \arg \max_{\theta} \hat{\lambda}_{n/2}^{tr}(\theta)$ where $\hat{\lambda}_{n/2}^{tr}(\theta)$ is computed with the training set D^{tr} .
 - 3: Set T_{α} to the $(1 - \alpha)$ -quantile of $\chi^2(J')$.
 - 4: Compute the test statistic $\hat{\lambda}_{n/2}^{te}(\theta)$ using D^{te} .
 - 5: Reject H_0 if $\hat{\lambda}_{n/2}^{te}(\theta) > T_{\alpha}$.
-

B Experiments on NIPS text collection

The full procedure for processing the NIPS text collection is summarized as following.

1. Download all 5903 papers from 1988 to 2015 from <https://papers.nips.cc/> as PDF files.
2. Convert each PDF file to text with `pdftotext`².
3. Remove all stop words. We use the list of stop words from <http://www.ranks.nl/stopwords>.
4. Keep only nouns. We use the list of nouns as available in WordNet-3.0³.
5. Keep only words which contain only English alphabets i.e., does not contain punctuations or numbers. Also, word length must be between 3 and 20 characters (inclusive).
6. Keep only words which occur in at least 5 documents, and in no more than 2000 documents.
7. Convert all characters to small case. Stem all words with SnowballStemmer in NLTK (Bird et al., 2009). For example, “recognize” and “recognizer” become “recogn” after stemming.
8. Categorize papers into disjoint collections. A paper is treated as belonging to a group if its title has at least one word from the list of keywords for the category. Papers that match the criteria of both categories are not considered. The lists of keywords are as follows.
 - (a) **Bayesian inference** (Bayes): graphical model, bayesian, inference, mcmc, monte carlo, posterior, prior, variational, markov, latent, probabilistic, exponential family.
 - (b) **Deep learning** (Deep): deep, drop out, auto-encod, convolutional, neural net, belief net, boltzmann.
 - (c) **Learning theory** (Learn): learning theory, consistency, theoretical guarantee, complexity, pac-bayes, pac-learning, generalization, uniform converg, bound, deviation, inequality, risk min, minimax, structural risk, VC, rademacher, asymptotic.
 - (d) **Neuroscience** (Neuro): motor control, neural, neuron, spiking, spike, cortex, plasticity, neural decod, neural encod, brain imag, biolog, perception, cognitive, emotion, synap, neural population, cortical, firing rate, firing-rate, sensor.
9. Randomly select 2000 words from the remaining words.
10. Treat each paper as a bag of words and construct a feature vector with TF-IDF (Manning et al., 2008).

²`pdftotext` is available at <http://poppler.freedesktop.org>.

³WordNet is available online at <https://wordnet.princeton.edu/wordnet/citing-wordnet/>.

B.1 Discriminative terms identified by ME test

In this section, we provide full lists of discriminative terms following the procedure described in Sec. 4. The top ten words in each problem are as follows.

- **Bayes-Bayes:** collabor, traffic, bay, permut, net, central, occlus, mask, draw, joint.
- **Bayes-Deep:** infer, bay, mont, adaptor, motif, haplotyp, ecg, covari, boltzmann, classifi.
- **Bayes-Learn:** infer, markov, graphic, segment, bandit, boundari, favor, carlo, prioriti, prop.
- **Bayes-Neuro:** spike, markov, cortex, dropout, recurr, iii, gibb, basin, circuit, subsystem.
- **Learn-Deep:** deep, forward, delay, subgroup, bandit, recept, invari, overlap, inequ, pia.
- **Learn-Neuro:** polici, interconnect, hardwar, decay, histolog, edg, period, basin, inject, human.

C Runtimes

In this section, we provide runtimes of all the experiments. The runtimes of the “Test power vs. sample n ” experiment are shown in Fig. 5. The runtimes of the “Test power vs. dimension d ” experiment are shown in Fig. 6. Table 4, 5 give the runtimes of the two real-data experiments.




Figure 5: Plots of runtimes in the “Test power vs. sample n ” experiment.




Figure 6: Plots of runtimes in the “Test power vs. dimension d ” experiment. The test sample size is 10000.

Table 4: Runtimes (in seconds) in the problem of distinguishing NIPS papers from two categories.

Problem	n^{te}	ME-full	ME-grid	SCF-full	SCF-grid	MMD-quad	MMD-lin
Bayes-Bayes	215	126.7	116	34.67	1.855	13.66	.6112
Bayes-Deep	216	118.3	111.7	36.41	1.933	13.59	.5105
Bayes-Learn	138	94.59	89.16	23.69	1.036	2.152	.36
Bayes-Neuro	394	142.5	130.3	69.19	3.533	32.71	.8643
Learn-Deep	149	105	99.59	24.99	1.253	2.417	.4744
Learn-Neuro	146	101.2	93.53	25.29	1.178	2.351	.3658

Table 5: Runtimes (in seconds) in the problem of distinguishing positive (+) and negative (-) facial expressions.

Problem	n^{te}	ME-full	ME-grid	SCF-full	SCF-grid	MMD-quad	MMD-lin
\pm vs. \pm	201	87.7	83.4	10.5	1.45	9.93	0.464
$+$ vs. $-$	201	85.0	80.6	11.7	1.42	10.4	0.482

In the cases where n is large (Fig. 5), MMD-quad has the largest runtime due to its quadratic dependency on the sample size. In the extreme case where the test sample size is 10000 (Fig. 6), it is computationally infeasible to run MMD-quad. We observe that the proposed ME-full and SCF-full have a slight overhead from the parameter optimization. However, since the optimization procedure is also linear in n , we are able to conduct an accurate test in less than 10 minutes even when the test sample size is 10000 and $d = 1500$ (see Fig. 6a, 6b). We note that the actual tests (after optimization) for all ME and SCF variants take less than one second in all cases. In the ME-full, we initialize the test locations with realizations from two multivariate normal distributions fitted to samples from P and Q . When d is large, this heuristic can be expensive. An alternative initialization scheme for \mathcal{V} is to randomly select J points from the two samples.

D Proof of theorem 2

Recall Theorem 2:

Theorem 2 (Consistency of $\hat{\lambda}_n$ in the ME test). *Let $\mathcal{X} \subseteq \mathbb{R}^d$ be a measurable set, and \mathbb{V} be a collection in which each element is a set of J test locations. All suprema over \mathcal{V} and k are to be understood as $\sup_{\mathcal{V} \in \mathbb{V}}$ and $\sup_{k \in \mathcal{K}}$ respectively. For a class of kernels \mathcal{K} on $\mathcal{X} \subseteq \mathbb{R}^d$, define*

$$\mathcal{F}_1 := \{\mathbf{x} \mapsto k(\mathbf{x}, \mathbf{v}) \mid k \in \mathcal{K}, \mathbf{v} \in \mathcal{X}\}, \quad \mathcal{F}_2 := \{\mathbf{x} \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{x}, \mathbf{v}') \mid k \in \mathcal{K}, \mathbf{v}, \mathbf{v}' \in \mathcal{X}\}, \quad (1)$$

$$\mathcal{F}_3 := \{(\mathbf{x}, \mathbf{y}) \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{y}, \mathbf{v}') \mid k \in \mathcal{K}, \mathbf{v}, \mathbf{v}' \in \mathcal{X}\}. \quad (2)$$

Assume that (1) \mathcal{K} is a uniformly bounded (by B) family of $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ measurable kernels, (2) $\tilde{c} := \sup_{\mathcal{V}, k} \|\Sigma^{-1}\|_F < \infty$, and (3) $\mathcal{F}_i = \{f_{\theta_i} \mid \theta_i \in \Theta_i\}$ is VC-subgraph with VC-index $VC(\mathcal{F}_i)$, and $\theta \mapsto f_{\theta_i}(m)$ is continuous ($\forall m, i = 1, 2, 3$). Let $\bar{c}_1 := 4B^2J\sqrt{J}\tilde{c}$, $\bar{c}_2 := 4B\sqrt{J}\tilde{c}$, and $\bar{c}_3 := 4B^2J\tilde{c}^2$. Let C_i -s ($i = 1, 2, 3$) be the universal constants associated to \mathcal{F}_i -s according to Theorem 2.6.7 in van der Vaart and Wellner (2000). Then for any $\delta \in (0, 1)$ with probability at least $1 - \delta$,

$$\begin{aligned} & \sup_{\mathcal{V}, k} \left| \bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \Sigma^{-1} \boldsymbol{\mu} \right| \\ & \leq 2T_{\mathcal{F}_1} \left(\frac{2}{\gamma_n} \bar{c}_1 B J \frac{2n-1}{n-1} + \bar{c}_2 \sqrt{J} \right) + \frac{2}{\gamma_n} \bar{c}_1 J (T_{\mathcal{F}_2} + T_{\mathcal{F}_3}) + \frac{8}{\gamma_n} \frac{\bar{c}_1 B^2 J}{n-1} + \bar{c}_3 \gamma_n, \text{ where} \\ T_{\mathcal{F}_j} &= \frac{16\sqrt{2}B^{\zeta_j}}{\sqrt{n}} \left(2\sqrt{\log[C_j \times VC(\mathcal{F}_j)(16e)^{VC(\mathcal{F}_j)}]} + \frac{\sqrt{2\pi[VC(\mathcal{F}_j) - 1]}}{2} \right) + B^{\zeta_j} \sqrt{\frac{2\log(5/\delta)}{n}}, \end{aligned}$$

for $j = 1, 2, 3$ and $\zeta_1 = 1, \zeta_2 = \zeta_3 = 2$.

A proof is given as follows.

D.1 Notations

Let $\langle \mathbf{A}, \mathbf{B} \rangle_F := \text{tr}(\mathbf{A}^\top \mathbf{B})$ be the Frobenius inner product, and $\|\mathbf{A}\|_F := \sqrt{\langle \mathbf{A}, \mathbf{A} \rangle_F}$. $\mathbf{A} \succeq \mathbf{0}$ means that $\mathbf{A} \in \mathbb{R}^{d \times d}$ is symmetric, positive semidefinite. For $\mathbf{a} \in \mathbb{R}^d$, $\|\mathbf{a}\|_2 = \langle \mathbf{a}, \mathbf{a} \rangle_2 = \mathbf{a}^\top \mathbf{a}$. $[\mathbf{a}_1; \dots; \mathbf{a}_N] \in \mathbb{R}^{d_1 + \dots + d_N}$ is the concatenation of the $\mathbf{a}_n \in \mathbb{R}^{d_n}$ vectors. \mathbb{R}^+ is the set of positive reals. $f \circ g$ is the composition of function f and g . Let \mathcal{M} denote a general metric space below. In measurability requirements metric spaces are meant to be endowed with their Borel σ -algebras.

Let \mathcal{C} be a collection of subsets of \mathcal{M} ($C \subseteq 2^{\mathcal{M}}$). \mathcal{C} is said to shatter an $\{p_1, p_2, \dots, p_i\} \subseteq \mathcal{M}$ set, if for any $S \subseteq \{p_1, p_2, \dots, p_i\}$ there exist $C \in \mathcal{C}$ such that $S = C \cap \{p_1, p_2, \dots, p_i\}$; in other words, arbitrary subset of $\{p_1, p_2, \dots, p_i\}$ can be cut out by an element of \mathcal{C} . The VC index of \mathcal{C} is the smallest i for which no set of size i is shattered:

$$VC(\mathcal{C}) = \inf \left\{ i : \max_{p_1, \dots, p_i} |\{C \cap \{p_1, \dots, p_i\} : C \in \mathcal{C}\}| < 2^i \right\}.$$

A collection \mathcal{C} of measurable sets is called VC-class if its index $VC(\mathcal{C})$ is finite. The subgraph of a real-valued function $f : \mathcal{M} \rightarrow \mathbb{R}$ is $\text{sub}(f) = \{(m, u) : u < f(m)\} \subseteq \mathcal{M} \times \mathbb{R}$. A collection of \mathcal{F} measurable functions is called VC-subgraph class, or shortly VC if the collection of all subgraphs of \mathcal{F} , $\{\text{sub}(f)\}_{f \in \mathcal{F}}$ is a VC-class of sets; its index is defined as $VC(\mathcal{F}) := VC(\{\text{sub}(f)\}_{f \in \mathcal{F}})$.

Let $L^0(\mathcal{M})$ be the set of $\mathcal{M} \rightarrow \mathbb{R}$ measurable functions. Given an i.i.d. (independent identically distributed) sample from \mathbb{P} ($w_i \stackrel{i.i.d.}{\sim} \mathbb{P}$), let $w_{1:n} = (w_1, \dots, w_n)$ and let $\mathbb{P}_n = \frac{1}{n} \sum_{i=1}^n \delta_{w_i}$ denote the empirical measure. $L^q(\mathcal{M}, \mathbb{P}_n) = \left\{ f \in L^0(\mathcal{M}) : \|f\|_{L^q(\mathcal{M}, \mathbb{P}_n)} := [\int_{\mathcal{M}} |f(w)|^q d\mathbb{P}_n(w)]^{\frac{1}{q}} = [\frac{1}{n} \sum_{i=1}^n |f(w_i)|^q]^{\frac{1}{q}} < \infty \right\}$ ($1 \leq q < \infty$), $\|f\|_{L^\infty(\mathcal{M})} := \sup_{m \in \mathcal{M}} |f(m)|$. Define $\mathbb{P}f := \int_{\mathcal{M}} f(w) d\mathbb{P}(w)$, where \mathbb{P} is a probability distribution on \mathcal{M} . Let $\|\mathbb{P}_n - \mathbb{P}\|_{\mathcal{F}} := \sup_{f \in \mathcal{F}} |\mathbb{P}_n f - \mathbb{P}f|$.

The diameter of a class $\mathcal{F} \subseteq L^2(\mathcal{M}, \mathbb{P}_n)$ is $\text{diam}(\mathcal{F}, L^2(\mathcal{M}, \mathbb{P}_n)) := \sup_{f, f' \in \mathcal{F}} \|f - f'\|_{L^2(\mathcal{M}, \mathbb{P}_n)}$, its r -covering number ($r > 0$) is the size of the smallest r -net

$$N(r, \mathcal{F}, L^2(\mathcal{M}, \mathbb{P}_n)) = \inf \{t \geq 1 : \exists f_1, \dots, f_t \in \mathcal{F} \text{ such that } \mathcal{F} \subseteq \cup_{i=1}^t B(r, f_i)\},$$

where $B(r, f) = \{g \in L^2(\mathcal{M}, \mathbb{P}_n) \mid \|f - g\|_{L^2(\mathcal{M}, \mathbb{P}_n)} \leq r\}$ is the ball with center f and radius r . $\times_{i=1}^N \mathbb{Q}_i$ is the N -fold product measure. For sets Q_i , $\times_{i=1}^n Q_i$ is their Cartesian product. For a function class $\mathcal{F} \subseteq L^0(\mathcal{M})$ and $w_{1:n} \in \mathcal{M}^n$, $R(\mathcal{F}, w_{1:n}) := \mathbb{E}_{\mathbf{r}} [\sup_{f \in \mathcal{F}} |\frac{1}{n} \sum_{i=1}^n r_i f(w_i)|]$ is the empirical Rademacher average, where $\mathbf{r} := r_{1:n}$ and r_i -s are i.i.d. samples from a Rademacher random variable $[\mathbb{P}(r_i = 1) = \mathbb{P}(r_i = -1) = \frac{1}{2}]$. Let (Θ, ρ) be a metric space; a collection of $\mathcal{F} = \{f_\theta \mid \theta \in \Theta\} \subseteq L^0(\mathcal{M})$ functions is called a separable Carathéodory family if Θ is separable and $\theta \mapsto f_\theta(m)$ is continuous for all $m \in \mathcal{M}$. $\text{span}(\cdot)$ denotes the linear hull of its arguments. $\Gamma(t) = \int_0^\infty u^{t-1} e^{-u} du$ denotes the Gamma function.

D.2 Bound in terms of \mathbf{S}_n and $\bar{\mathbf{z}}_n$

For brevity, we will interchangeably use \mathbf{S}_n for $\mathbf{S}_n(\mathcal{V})$ and $\bar{\mathbf{z}}_n$ for $\bar{\mathbf{z}}_n(\mathcal{V})$. $\mathbf{S}_n(\mathcal{V})$ and $\bar{\mathbf{z}}_n(\mathcal{V})$ will be used mainly when the dependency of \mathcal{V} needs to be emphasized. We start with $\sup_{\mathcal{V}, k} |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}|$ and upper bound the argument of $\sup_{\mathcal{V}, k}$ as

$$\begin{aligned} & |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}| \\ &= |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \boldsymbol{\mu} + \boldsymbol{\mu}^\top (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \boldsymbol{\mu} - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}| \\ &\leq |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \boldsymbol{\mu}| + |\boldsymbol{\mu}^\top (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \boldsymbol{\mu} - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}| \\ &:= (\square_1) + (\square_2). \end{aligned}$$

For (\square_1) , we have

$$\begin{aligned} & |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \boldsymbol{\mu}| \\ &= \left| \langle \bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top, (\mathbf{S}_n + \gamma_n I)^{-1} \rangle_F - \langle \boldsymbol{\mu} \boldsymbol{\mu}^\top, (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \rangle_F \right| \\ &= \left| \langle \bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top, (\mathbf{S}_n + \gamma_n I)^{-1} \rangle_F - \langle \bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top, (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \rangle_F + \langle \bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top, (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \rangle_F - \langle \boldsymbol{\mu} \boldsymbol{\mu}^\top, (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \rangle_F \right| \\ &\leq \left| \langle \bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top, (\mathbf{S}_n + \gamma_n I)^{-1} - (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \rangle_F \right| + \left| \langle \bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top - \boldsymbol{\mu} \boldsymbol{\mu}^\top, (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \rangle_F \right| \\ &= \|\bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top\|_F \|(\mathbf{S}_n + \gamma_n I)^{-1} - (\boldsymbol{\Sigma} + \gamma_n I)^{-1}\|_F + \|\bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top - \boldsymbol{\mu} \boldsymbol{\mu}^\top\|_F \|(\boldsymbol{\Sigma} + \gamma_n I)^{-1}\|_F \\ &\stackrel{(a)}{\leq} \|\bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top\|_F \|(\mathbf{S}_n + \gamma_n I)^{-1} [(\boldsymbol{\Sigma} + \gamma_n I) - (\mathbf{S}_n + \gamma_n I)] (\boldsymbol{\Sigma} + \gamma_n I)^{-1}\|_F + \|\bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top - \bar{\mathbf{z}}_n \boldsymbol{\mu}^\top + \bar{\mathbf{z}}_n \boldsymbol{\mu}^\top - \boldsymbol{\mu} \boldsymbol{\mu}^\top\|_F \|\boldsymbol{\Sigma}^{-1}\|_F \\ &\stackrel{(a)}{\leq} \|\bar{\mathbf{z}}_n \bar{\mathbf{z}}_n^\top\|_F \|(\mathbf{S}_n + \gamma_n I)^{-1}\|_F \|\boldsymbol{\Sigma} - \mathbf{S}_n\|_F \|\boldsymbol{\Sigma}^{-1}\|_F + \|\bar{\mathbf{z}}_n (\bar{\mathbf{z}}_n - \boldsymbol{\mu})^\top\|_F \|\boldsymbol{\Sigma}^{-1}\|_F + \|(\bar{\mathbf{z}}_n - \boldsymbol{\mu}) \boldsymbol{\mu}^\top\|_F \|\boldsymbol{\Sigma}^{-1}\|_F \\ &\stackrel{(b)}{\leq} \frac{\sqrt{J}}{\gamma_n} \|\bar{\mathbf{z}}_n\|_2^2 \|\boldsymbol{\Sigma} - \mathbf{S}_n\|_F \|\boldsymbol{\Sigma}^{-1}\|_F + \|\bar{\mathbf{z}}_n\|_2 \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 \|\boldsymbol{\Sigma}^{-1}\|_F + \|\boldsymbol{\mu}\|_2 \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 \|\boldsymbol{\Sigma}^{-1}\|_F, \end{aligned}$$

where at (a) we use $\|(\boldsymbol{\Sigma} + \gamma_n I)^{-1}\|_F \leq \|\boldsymbol{\Sigma}^{-1}\|_F$ and at (b) we use $\|(\mathbf{S}_n + \gamma_n I)^{-1}\|_F \leq \sqrt{J} \|(\mathbf{S}_n + \gamma_n I)^{-1}\|_2 \leq \sqrt{J}/\gamma_n$.

For (\square_2) , we have

$$\begin{aligned} & |\boldsymbol{\mu}^\top (\boldsymbol{\Sigma} + \gamma_n I)^{-1} \boldsymbol{\mu} - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}| = \left| \langle \boldsymbol{\mu} \boldsymbol{\mu}^\top, (\boldsymbol{\Sigma} + \gamma_n I)^{-1} - \boldsymbol{\Sigma}^{-1} \rangle_F \right| \\ &\leq \|\boldsymbol{\mu} \boldsymbol{\mu}^\top\|_F \|(\boldsymbol{\Sigma} + \gamma_n I)^{-1} - \boldsymbol{\Sigma}^{-1}\|_F \\ &= \|\boldsymbol{\mu}\|_2^2 \|(\boldsymbol{\Sigma} + \gamma_n I)^{-1} [\boldsymbol{\Sigma} - (\boldsymbol{\Sigma} + \gamma_n I)] \boldsymbol{\Sigma}^{-1}\|_F \\ &= \gamma_n \|\boldsymbol{\mu}\|_2^2 \|(\boldsymbol{\Sigma} + \gamma_n I)^{-1} \boldsymbol{\Sigma}^{-1}\|_F \\ &\leq \gamma_n \|\boldsymbol{\mu}\|_2^2 \|(\boldsymbol{\Sigma} + \gamma_n I)^{-1}\|_F \|\boldsymbol{\Sigma}^{-1}\|_F \\ &\stackrel{(a)}{\leq} \gamma_n \|\boldsymbol{\mu}\|_2^2 \|\boldsymbol{\Sigma}^{-1}\|_F^2. \end{aligned}$$

Combining the upper bounds for (\square_1) and (\square_2) , we arrive at

$$\begin{aligned} & |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}| \\ &\leq \frac{\sqrt{J}}{\gamma_n} \|\bar{\mathbf{z}}_n\|_2^2 \|\boldsymbol{\Sigma} - \mathbf{S}_n\|_F \|\boldsymbol{\Sigma}^{-1}\|_F + (\|\bar{\mathbf{z}}_n\|_2 + \|\boldsymbol{\mu}\|_2) \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 \|\boldsymbol{\Sigma}^{-1}\|_F + \gamma_n \|\boldsymbol{\mu}\|_2^2 \|\boldsymbol{\Sigma}^{-1}\|_F^2 \end{aligned}$$

$$\begin{aligned}
&\leq 4B^2 J \tilde{c} \frac{\sqrt{J}}{\gamma_n} \|\Sigma - \mathbf{S}_n\|_F + 4B\sqrt{J}\tilde{c} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 + 4B^2 J \tilde{c}^2 \gamma_n \\
&= \frac{\bar{c}_1}{\gamma_n} \|\Sigma - \mathbf{S}_n\|_F + \bar{c}_2 \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 + \bar{c}_3 \gamma_n
\end{aligned} \tag{3}$$

with $\bar{c}_1 := 4B^2 J \sqrt{J} \tilde{c}$, $\bar{c}_2 := 4B\sqrt{J}\tilde{c}$, $\bar{c}_3 := 4B^2 J \tilde{c}^2$, and $\tilde{c} := \sup_{\mathcal{V}, k} \|\Sigma^{-1}\|_F < \infty$, where we applied the triangle inequality, the CBS (Cauchy-Bunyakovskii-Schwarz) inequality, and $\|\mathbf{a}\mathbf{b}^\top\|_F = \|\mathbf{a}\|_2 \|\mathbf{b}\|_2$. The boundedness of kernel k with the Jensen inequality implies that

$$\|\bar{\mathbf{z}}_n\|_2^2 = \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i \right\|_2^2 \leq \frac{1}{n} \sum_{i=1}^n \|\mathbf{z}_i\|_2^2 = \frac{1}{n} \sum_{i=1}^n \|(k(\mathbf{x}_i, \mathbf{v}_j) - k(\mathbf{y}_i, \mathbf{v}_j))_{j=1}^J\|_2^2 \tag{4}$$

$$\begin{aligned}
&= \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^J [k(\mathbf{x}_i, \mathbf{v}_j) - k(\mathbf{y}_i, \mathbf{v}_j)]^2 \\
&\leq \frac{2}{n} \sum_{i=1}^n \sum_{j=1}^J k^2(\mathbf{x}_i, \mathbf{v}_j) + k^2(\mathbf{y}_i, \mathbf{v}_j) \leq 4B^2 J,
\end{aligned} \tag{5}$$

$$\|\boldsymbol{\mu}(\mathcal{V})\|_2^2 = \sum_{j=1}^J (\mathbb{E}_{\mathbf{x}\mathbf{y}} [k(\mathbf{x}, \mathbf{v}_j) - k(\mathbf{y}, \mathbf{v}_j)])^2 \leq \sum_{j=1}^J \mathbb{E}_{\mathbf{x}\mathbf{y}} [k(\mathbf{x}, \mathbf{v}_j) - k(\mathbf{y}, \mathbf{v}_j)]^2 \leq 4B^2 J. \tag{6}$$

Taking sup in (3), we get

$$\sup_{\mathcal{V}, k} |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \Sigma^{-1} \boldsymbol{\mu}| \leq \frac{\bar{c}_1}{\gamma_n} \sup_{\mathcal{V}, k} \|\Sigma - \mathbf{S}_n\|_F + \bar{c}_2 \sup_{\mathcal{V}, k} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 + \bar{c}_3 \gamma_n.$$

D.3 Empirical process bound on $\bar{\mathbf{z}}_n$

Recall that $\bar{\mathbf{z}}_n(\mathcal{V}) = \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i(\mathcal{V}) \in \mathbb{R}^J$, $\mathbf{z}_i(\mathcal{V}) = (k(\mathbf{x}_i, \mathbf{v}_j) - k(\mathbf{y}_i, \mathbf{v}_j))_{j=1}^J \in \mathbb{R}^J$, $\boldsymbol{\mu}(\mathcal{V}) = (\mathbb{E}_{\mathbf{x}\mathbf{y}} [k(\mathbf{x}, \mathbf{v}_j) - k(\mathbf{y}, \mathbf{v}_j)])_{j=1}^J$; thus

$$\sup_{\mathcal{V}} \sup_{k \in \mathcal{K}} \|\bar{\mathbf{z}}_n(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V})\|_2 = \sup_{\mathcal{V}} \sup_{k \in \mathcal{K}} \sup_{\mathbf{b} \in B(1, \mathbf{0})} \langle \mathbf{b}, \bar{\mathbf{z}}_n(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \rangle_2$$

using that $\|\mathbf{a}\|_2 = \sup_{\mathbf{b} \in B(1, \mathbf{0})} \langle \mathbf{a}, \mathbf{b} \rangle_2$. Let us bound the argument of the supremum:

$$\begin{aligned}
\langle \mathbf{b}, \bar{\mathbf{z}}_n(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \rangle_2 &\leq \sum_{j=1}^J |b_j| \left| \frac{1}{n} \sum_{i=1}^n [k(\mathbf{x}_i, \mathbf{v}_j) - k(\mathbf{y}_i, \mathbf{v}_j)] - \mathbb{E}_{\mathbf{x}\mathbf{y}} [k(\mathbf{x}, \mathbf{v}_j) - k(\mathbf{y}, \mathbf{v}_j)] \right| \\
&\leq \sum_{j=1}^J |b_j| \left(\left| \frac{1}{n} \sum_{i=1}^n k(\mathbf{x}_i, \mathbf{v}_j) - \mathbb{E}_{\mathbf{x}} k(\mathbf{x}, \mathbf{v}_j) \right| + \left| \frac{1}{n} \sum_{i=1}^n k(\mathbf{y}_i, \mathbf{v}_j) - \mathbb{E}_{\mathbf{y}} k(\mathbf{y}, \mathbf{v}_j) \right| \right) \\
&\leq \sqrt{J} \sup_{\mathbf{v} \in \mathcal{X}} \sup_{k \in \mathcal{K}} \left| \frac{1}{n} \sum_{i=1}^n k(\mathbf{x}_i, \mathbf{v}) - \mathbb{E}_{\mathbf{x}} k(\mathbf{x}, \mathbf{v}) \right| + \sqrt{J} \sup_{\mathbf{v} \in \mathcal{X}} \sup_{k \in \mathcal{K}} \left| \frac{1}{n} \sum_{i=1}^n k(\mathbf{y}_i, \mathbf{v}) - \mathbb{E}_{\mathbf{y}} k(\mathbf{y}, \mathbf{v}) \right| \\
&= \sqrt{J} \|P_n - P\|_{\mathcal{F}_1} + \sqrt{J} \|Q_n - Q\|_{\mathcal{F}_1}
\end{aligned} \tag{7}$$

by the triangle inequality and exploiting that $\|\mathbf{b}\|_1 \leq \sqrt{J} \|\mathbf{b}\|_2 \leq \sqrt{J}$ with $\mathbf{b} \in B(1, \mathbf{0})$. Thus, we have

$$\sup_{\mathcal{V}} \sup_{k \in \mathcal{K}} \|\bar{\mathbf{z}}_n(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V})\|_2 \leq \sqrt{J} \|P_n - P\|_{\mathcal{F}_1} + \sqrt{J} \|Q_n - Q\|_{\mathcal{F}_1}.$$

D.4 Empirical process bound on \mathbf{S}_n

Noting that

$$\Sigma(\mathcal{V}) = \mathbb{E}_{\mathbf{x}\mathbf{y}} [\mathbf{z}(\mathcal{V}) \mathbf{z}^\top(\mathcal{V})] - \boldsymbol{\mu}(\mathcal{V}) \boldsymbol{\mu}^\top(\mathcal{V}), \quad \mathbf{S}_n(\mathcal{V}) = \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \mathbf{z}_a^\top(\mathcal{V}) - \frac{1}{n(n-1)} \sum_{a=1}^n \sum_{b \neq a} \mathbf{z}_a \mathbf{z}_b^\top,$$

$$\mathbb{E}_{\mathbf{xy}} [\mathbf{z}(\mathcal{V}) \mathbf{z}^\top(\mathcal{V})] = \mathbb{E}_{\mathbf{xy}} \left[\frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \mathbf{z}_a^\top(\mathcal{V}) \right], \quad \boldsymbol{\mu}(\mathcal{V}) \boldsymbol{\mu}^\top(\mathcal{V}) = \mathbb{E}_{\mathbf{xy}} \left[\frac{1}{n(n-1)} \sum_{a=1}^n \sum_{b \neq a} \mathbf{z}_a(\mathcal{V}) \mathbf{z}_b^\top(\mathcal{V}) \right],$$

we bound our target quantity as

$$\|\mathbf{S}_n(\mathcal{V}) - \boldsymbol{\Sigma}(\mathcal{V})\|_F \leq \left\| \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \mathbf{z}_a^\top(\mathcal{V}) - \mathbb{E}_{\mathbf{xy}} [\mathbf{z}(\mathcal{V}) \mathbf{z}^\top(\mathcal{V})] \right\|_F + \left\| \frac{1}{n(n-1)} \sum_{a=1}^n \sum_{b \neq a} \mathbf{z}_a(\mathcal{V}) \mathbf{z}_b^\top(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \boldsymbol{\mu}^\top(\mathcal{V}) \right\|_F$$

$$=: (*_1) + (*_2). \quad (8)$$

$$\begin{aligned} (*_2) &= \left\| \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \left[\frac{1}{n-1} \sum_{b \neq a} \mathbf{z}_b^\top(\mathcal{V}) \right] - \boldsymbol{\mu}(\mathcal{V}) \boldsymbol{\mu}^\top(\mathcal{V}) \right\|_F \\ &\leq \left\| \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \left(\frac{1}{n-1} \sum_{b \neq a} \mathbf{z}_b^\top(\mathcal{V}) - \boldsymbol{\mu}^\top(\mathcal{V}) \right) \right\|_F + \left\| \left(\frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \right) \boldsymbol{\mu}^\top(\mathcal{V}) \right\|_F \\ &\leq \left\| \left(\frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \right) \left(\frac{1}{n-1} \sum_{b=1}^n \mathbf{z}_b(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \right)^\top \right\|_F + \left\| \left(\frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \right) \frac{\mathbf{z}_a^\top(\mathcal{V})}{n-1} \right\|_F \\ &\quad + \left\| \left(\frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \right) \boldsymbol{\mu}^\top(\mathcal{V}) \right\|_F \\ &= \|\bar{\mathbf{z}}_n(\mathcal{V})\|_2 \left\| \frac{1}{n-1} \sum_{b=1}^n \mathbf{z}_b(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \right\|_2 + \frac{1}{n-1} \|\bar{\mathbf{z}}_n(\mathcal{V})\|_2 \|\mathbf{z}_a(\mathcal{V})\|_2 + \|\bar{\mathbf{z}}_n(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V})\|_2 \|\boldsymbol{\mu}(\mathcal{V})\|_2 \\ &\leq 2B\sqrt{J} \left(\frac{n}{n-1} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}(\mathcal{V})\|_2 + \frac{2B\sqrt{J}}{n-1} \right) + \frac{1}{n-1} 4B^2J + 2B\sqrt{J} \|\bar{\mathbf{z}}_n(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V})\|_2 \\ &= \frac{8B^2J}{n-1} + 2B\sqrt{J} \frac{2n-1}{n-1} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}(\mathcal{V})\|_2 \quad (9) \end{aligned}$$

using the triangle inequality, the sub-additivity of sup, $\|\mathbf{ab}^T\|_F = \|\mathbf{a}\|_2 \|\mathbf{b}\|_2$, $\|\bar{\mathbf{z}}_n(\mathcal{V})\|_2 \leq 2B\sqrt{J}$, $\|\mathbf{z}_a(\mathcal{V})\|_2 \leq 2B\sqrt{J}$ [see Eq. (5)] and

$$\left\| \frac{1}{n-1} \sum_{b=1}^n \mathbf{z}_b(\mathcal{V}) - \boldsymbol{\mu}(\mathcal{V}) \right\|_2 = \left\| \frac{n}{n-1} \bar{\mathbf{z}}_n - \frac{n}{n-1} \boldsymbol{\mu}(\mathcal{V}) + \frac{1}{n-1} \boldsymbol{\mu}(\mathcal{V}) \right\|_2 \leq \frac{n}{n-1} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}(\mathcal{V})\|_2 + \frac{1}{n-1} \|\boldsymbol{\mu}(\mathcal{V})\|_2$$

with Eq. (6). Considering the first term in Eq. (8)

$$\begin{aligned} &\left\| \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \mathbf{z}_a^\top(\mathcal{V}) - \mathbb{E}_{\mathbf{xy}} [\mathbf{z}(\mathcal{V}) \mathbf{z}^\top(\mathcal{V})] \right\|_F = \sup_{\mathbf{B} \in B(1,0)} \left\langle \mathbf{B}, \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a(\mathcal{V}) \mathbf{z}_a^\top(\mathcal{V}) - \mathbb{E}_{\mathbf{xy}} [\mathbf{z}(\mathcal{V}) \mathbf{z}^\top(\mathcal{V})] \right\rangle_F \\ &\leq \sup_{\mathbf{B} \in B(1,0)} \sum_{i,j=1}^J |B_{ij}| \left| \frac{1}{n} \sum_{a=1}^n [k(\mathbf{x}_a, \mathbf{v}_i) - k(\mathbf{y}_a, \mathbf{v}_i)][k(\mathbf{x}_a, \mathbf{v}_j) - k(\mathbf{y}_a, \mathbf{v}_j)] - \mathbb{E}_{\mathbf{xy}} ([k(\mathbf{x}, \mathbf{v}_i) - k(\mathbf{y}, \mathbf{v}_i)][k(\mathbf{x}, \mathbf{v}_j) - k(\mathbf{y}, \mathbf{v}_j)]) \right| \\ &\leq \sup_{\mathbf{B} \in B(1,0)} \sum_{i,j=1}^J |B_{ij}| \left(\left| \frac{1}{n} \sum_{a=1}^n k(\mathbf{x}_a, \mathbf{v}_i) k(\mathbf{x}_a, \mathbf{v}_j) - \mathbb{E}_{\mathbf{x}} [k(\mathbf{x}, \mathbf{v}_i) k(\mathbf{x}, \mathbf{v}_j)] \right| \right. \\ &\quad + \left| \frac{1}{n} \sum_{a=1}^n k(\mathbf{x}_a, \mathbf{v}_i) k(\mathbf{y}_a, \mathbf{v}_j) - \mathbb{E}_{\mathbf{xy}} [k(\mathbf{x}, \mathbf{v}_i) k(\mathbf{y}, \mathbf{v}_j)] \right| \\ &\quad + \left| \frac{1}{n} \sum_{a=1}^n k(\mathbf{y}_a, \mathbf{v}_i) k(\mathbf{x}_a, \mathbf{v}_j) - \mathbb{E}_{\mathbf{xy}} [k(\mathbf{y}, \mathbf{v}_i) k(\mathbf{x}, \mathbf{v}_j)] \right| + \left. \left| \frac{1}{n} \sum_{a=1}^n k(\mathbf{y}_a, \mathbf{v}_i) k(\mathbf{y}_a, \mathbf{v}_j) - \mathbb{E}_{\mathbf{y}} [k(\mathbf{y}, \mathbf{v}_i) k(\mathbf{y}, \mathbf{v}_j)] \right| \right) \\ &\leq J \sup_{\mathbf{v}, \mathbf{v}' \in \mathcal{X}} \sup_{k \in \mathcal{K}} \left| \frac{1}{n} \sum_{a=1}^n k(\mathbf{x}_a, \mathbf{v}) k(\mathbf{x}_a, \mathbf{v}') - \mathbb{E}_{\mathbf{x}} [k(\mathbf{x}, \mathbf{v}) k(\mathbf{x}, \mathbf{v}')] \right| \end{aligned}$$

$$\begin{aligned}
& + 2J \sup_{\mathbf{v}, \mathbf{v}' \in \mathcal{X}} \sup_{k \in \mathcal{K}} \left| \frac{1}{n} \sum_{a=1}^n k(\mathbf{x}_a, \mathbf{v}) k(\mathbf{y}_a, \mathbf{v}') - \mathbb{E}_{\mathbf{x}\mathbf{y}} [k(\mathbf{x}, \mathbf{v}) k(\mathbf{y}, \mathbf{v}')] \right| \\
& + J \sup_{\mathbf{v}, \mathbf{v}' \in \mathcal{X}} \sup_{k \in \mathcal{K}} \left| \frac{1}{n} \sum_{a=1}^n k(\mathbf{y}_a, \mathbf{v}) k(\mathbf{y}_a, \mathbf{v}') - \mathbb{E}_{\mathbf{y}} [k(\mathbf{y}, \mathbf{v}) k(\mathbf{y}, \mathbf{v}')] \right|
\end{aligned}$$

by exploiting that $\|\mathbf{A}\|_F = \sup_{\mathbf{B} \in B(1, \mathbf{0})} \langle \mathbf{B}, \mathbf{A} \rangle_F$, and $\sum_{i,j=1}^J |B_{ij}| \leq J \|\mathbf{B}\|_F \leq J$ with $\mathbf{B} \in B(1, \mathbf{0})$. Using the bounds obtained for the two terms of Eq. (8), we get

$$\begin{aligned}
& \sup_{\mathcal{V}} \sup_{k \in \mathcal{K}} \|\mathbf{S}_n(\mathcal{V}) - \boldsymbol{\Sigma}(\mathcal{V})\|_F \leq \\
& \leq \frac{8B^2J}{n-1} + 2B\sqrt{J} \frac{2n-1}{n-1} \sup_{\mathcal{V}} \sup_{k \in \mathcal{K}} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}(\mathcal{V})\|_2 + J (\|P_n - P\|_{\mathcal{F}_2} + 2\|(P \times Q)_n - (P \times Q)\|_{\mathcal{F}_3} + \|Q_n - Q\|_{\mathcal{F}_2}).
\end{aligned} \tag{10}$$

D.5 Bounding by concentration and the VC property

By combining Eqs. (3), (7) and (10)

$$\begin{aligned}
& \sup_{\mathcal{V}} \sup_k |\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n - \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}| \leq \\
& \leq \frac{\bar{c}_1}{\gamma_n} \left[\frac{8B^2J}{n-1} + 2B\sqrt{J} \frac{2n-1}{n-1} \sqrt{J} (\|P_n - P\|_{\mathcal{F}_1} + \|Q_n - Q\|_{\mathcal{F}_1}) \right. \\
& \quad \left. + J (\|P_n - P\|_{\mathcal{F}_2} + 2\|(P \times Q)_n - (P \times Q)\|_{\mathcal{F}_3} + \|Q_n - Q\|_{\mathcal{F}_2}) \right] \\
& \quad + \bar{c}_2 \sqrt{J} (\|P_n - P\|_{\mathcal{F}_1} + \|Q_n - Q\|_{\mathcal{F}_1}) + \bar{c}_3 \gamma_n \\
& = (\|P_n - P\|_{\mathcal{F}_1} + \|Q_n - Q\|_{\mathcal{F}_1}) \left(\frac{2}{\gamma_n} \bar{c}_1 B J \frac{2n-1}{n-1} + \bar{c}_2 \sqrt{J} \right) + \bar{c}_3 \gamma_n \\
& \quad + \frac{\bar{c}_1}{\gamma_n} J [\|P_n - P\|_{\mathcal{F}_2} + \|Q_n - Q\|_{\mathcal{F}_2} + 2\|(P \times Q)_n - (P \times Q)\|_{\mathcal{F}_3}] + \frac{8}{\gamma_n} \frac{\bar{c}_1 B^2 J}{n-1}.
\end{aligned} \tag{11}$$

Applying Lemma 3 with $\frac{\delta}{5}$, we get the statement with a union bound. \square

Lemma 3 (Concentration of the empirical process for uniformly bounded separable Carathéodory VC classes).
Let \mathcal{F} be

1. VC-subgraph class of $\mathcal{M} \rightarrow \mathbb{R}$ functions with VC index $VC(\mathcal{F})$,
2. a uniformly bounded ($\|f\|_{L^\infty(\mathcal{M})} \leq K < \infty, \forall f \in \mathcal{F}$) separable Carathéodory family.

Let \mathbb{Q} be a probability measure, and let $\mathbb{Q}_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ be the corresponding empirical measure. Then for any $\delta \in (0, 1)$ with probability at least $1 - \delta$

$$\|\mathbb{Q}_n - \mathbb{Q}\|_{\mathcal{F}} \leq \frac{16\sqrt{2}K}{\sqrt{n}} \left[2\sqrt{\log [C \times VC(\mathcal{F})(16e)^{VC(\mathcal{F})}]} + \frac{\sqrt{2\pi[VC(\mathcal{F}) - 1]}}{2} \right] + K\sqrt{\frac{2\log(\frac{1}{\delta})}{n}}$$

where the universal constant C is associated according to Lemma 7(iv).

Proof. Notice that $g(x_1, \dots, x_n) = \|\mathbb{Q}_n - \mathbb{Q}\|_{\mathcal{F}}$ satisfies the bounded difference property with $b = \frac{2K}{n}$ [see Eq. (18)]:

$$\begin{aligned}
& |g(\mathbf{x}_1, \dots, \mathbf{x}_n) - g(\mathbf{x}_1, \dots, \mathbf{x}_j, \mathbf{x}'_j, \mathbf{x}_{j+1}, \dots, \mathbf{x}_n)| \\
& \leq \left| \sup_{f \in \mathcal{F}} \left| \mathbb{Q}f - \frac{1}{n} \sum_{i=1}^n f(\mathbf{x}_i) \right| - \sup_{f \in \mathcal{F}} \left| \mathbb{Q}f - \frac{1}{n} \sum_{i=1}^n f(\mathbf{x}_i) + \frac{1}{n} [f(\mathbf{x}_j) - f(\mathbf{x}'_j)] \right| \right| \\
& \leq \frac{1}{n} \sup_{f \in \mathcal{F}} |f(\mathbf{x}_j) - f(\mathbf{x}'_j)| \leq \frac{1}{n} \left(\sup_{f \in \mathcal{F}} |f(\mathbf{x}_j)| + \sup_{f \in \mathcal{F}} |f(\mathbf{x}'_j)| \right) \leq \frac{2K}{n}.
\end{aligned}$$

Hence, applying Lemma 8, and using symmetrization Steinwart and Christmann (2008) (Prop. 7.10) for the uniformly bounded separable Carathéodory \mathcal{F} class, for arbitrary $\delta \in (0, 1)$ with probability at least $1 - \delta$

$$\begin{aligned}\|\mathbb{Q}_n - \mathbb{Q}\|_{\mathcal{F}} &\leq \mathbb{E}_{x_{1:n}} \|\mathbb{Q}_n - \mathbb{Q}\|_{\mathcal{F}} + K \sqrt{\frac{2 \log(\frac{1}{\delta})}{n}} \\ &\leq 2\mathbb{E}_{x_{1:n}} R(\mathcal{F}, x_{1:n}) + K \sqrt{\frac{2 \log(\frac{1}{\delta})}{n}}.\end{aligned}$$

By the Dudley entropy bound Bousquet (2003) [see Eq. (4.4); $\text{diam}(\mathcal{F}, L^2(\mathcal{M}, \mathbb{Q}_n)) \leq 2 \sup_{f \in \mathcal{F}} \|f\|_{L^2(\mathcal{M}, \mathbb{Q}_n)} \leq 2 \sup_{f \in \mathcal{F}} \|f\|_{L^\infty(\mathcal{M})} \leq 2K < \infty$], Lemma 7(iv) [with $F \equiv K$, $q = 2$, $\mathbb{M} = \mathbb{Q}_n$] and the monotone decreasing property of the covering number, one arrives at

$$\begin{aligned}R(\mathcal{F}, x_{1:n}) &\leq \frac{8\sqrt{2}}{\sqrt{n}} \int_0^{2K} \sqrt{\log N(r, \mathcal{F}, L^2(\mathcal{M}, \mathbb{Q}_n))} dr \\ &\leq \frac{8\sqrt{2}}{\sqrt{n}} \left[\int_0^K \sqrt{\log N(r, \mathcal{F}, L^2(\mathcal{M}, \mathbb{Q}_n))} dr + K \sqrt{\log N(K, \mathcal{F}, L^2(\mathcal{M}, \mathbb{Q}_n))} \right] \\ &\leq \frac{8\sqrt{2}K}{\sqrt{n}} \left[\int_0^1 \sqrt{\log N(rK, \mathcal{F}, L^2(\mathcal{M}, \mathbb{Q}_n))} dr + \sqrt{\log N(K, \mathcal{F}, L^2(\mathcal{M}, \mathbb{Q}_n))} \right] \\ &\leq \frac{8\sqrt{2}K}{\sqrt{n}} \left[\int_0^1 \sqrt{\log \left[a_1 \left(\frac{1}{r} \right)^{a_2} \right]} dr + \sqrt{\log(a_1)} \right] = \frac{8\sqrt{2}K}{\sqrt{n}} \left[2\sqrt{\log(a_1)} + \int_0^1 \sqrt{a_2 \log \left(\frac{1}{r} \right)} dr \right] \\ &= \frac{8\sqrt{2}K}{\sqrt{n}} \left[2\sqrt{\log(a_1)} + \frac{\sqrt{\pi a_2}}{2} \right],\end{aligned}$$

where $a_1 := C \times VC(\mathcal{F})(16e)^{VC(\mathcal{F})}$, $a_2 := 2[VC(\mathcal{F}) - 1]$ and $\int_0^1 \sqrt{\log(\frac{1}{r})} dr = \int_0^\infty t^{\frac{1}{2}} e^{-t} dt = \Gamma(\frac{3}{2}) = \frac{\sqrt{\pi}}{2}$. \square

Lemma 4 (Properties of \mathcal{F}_i from \mathcal{K}).

1. **Uniform boundedness of \mathcal{F}_i -s** [see Eqs. (1)-(2)]: If \mathcal{K} is uniformly bounded, i.e., $\exists B < \infty$ such that $\sup_{k \in \mathcal{K}} \sup_{(\mathbf{x}, \mathbf{y}) \in \mathcal{X}^2} |k(\mathbf{x}, \mathbf{y})| \leq B$; then \mathcal{F}_1 , \mathcal{F}_2 and \mathcal{F}_3 [Eqs. (1)-(2)] are also uniformly bounded with B , B^2 , B^2 constants, respectively. That is, $\sup_{k \in \mathcal{K}, \mathbf{v} \in \mathcal{X}} |k(\mathbf{x}, \mathbf{v})| \leq B$, $\sup_{k \in \mathcal{K}, (\mathbf{v}, \mathbf{v}') \in \mathcal{X}^2} |k(\mathbf{x}, \mathbf{v})k(\mathbf{x}, \mathbf{v}')| \leq B^2$, $\sup_{k \in \mathcal{K}, (\mathbf{v}, \mathbf{v}') \in \mathcal{X}^2} |k(\mathbf{x}, \mathbf{v})k(\mathbf{y}, \mathbf{v}')| \leq B^2$.
2. **Separability of \mathcal{F}_i** : since \mathcal{F}_1 , \mathcal{F}_2 and \mathcal{F}_3 is parameterized by $\Theta = \mathcal{K} \times \mathcal{X}$, $\mathcal{K} \times \mathcal{X}^2$, $\mathcal{K} \times \mathcal{X}^2$, separability of \mathcal{K} implies that of Θ .
3. **Measurability of \mathcal{F}_i** : $\forall k \in \mathcal{K}$ is measurable, then the elements of \mathcal{F}_i ($i = 1, 2, 3$) are also measurable. \square

E Example kernel families

Below we give examples for \mathcal{K} kernel classes for which the associated \mathcal{F}_i -s are VC-subgraph and uniformly bounded separable Carathéodory families. The VC property will be a direct consequence of the VC indices of finite-dimensional function classes and preservation theorems (see Lemma 7); for a nice example application see Srebro and Ben-David (2006) (Section 5) who study the pseudo-dimension of $(\mathbf{x}, \mathbf{y}) \mapsto k(\mathbf{x}, \mathbf{y})$ kernel classes, for different Gaussian families. We take these Gaussian classes (isotropic, full) and use the preservation trick to bound the VC indices of the associated \mathcal{F}_i -s.

Lemma 5 (\mathcal{F}_i -s are VC-subgraph and uniformly bounded separable Carathéodory families for isotropic Gaussian kernel). Let $\mathcal{K} = \left\{ k_\sigma : (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{X} \subseteq \mathbb{R}^d \times \mathbb{R}^d \mapsto e^{-\frac{\|\mathbf{x} - \mathbf{y}\|_2^2}{2\sigma^2}} : \sigma > 0 \right\}$. Then the \mathcal{F}_1 , \mathcal{F}_2 , \mathcal{F}_3 classes [see Eqs. (1)-(2)] associated to \mathcal{K} are

- VC-subgraphs with indices $VC(\mathcal{F}_1) \leq d + 4$, $VC(\mathcal{F}_2) \leq d + 4$, $VC(\mathcal{F}_3) \leq 2d + 4$, and

- uniformly bounded separable Carathéodory families, with $\|f\|_{L^\infty(\mathcal{M})} \leq 1$ for all $f \in \{\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3\}$.⁴

Proof. **VC subgraph property:**

- \mathcal{F}_1 : Consider the function class $\mathcal{G} = \left\{ \mathbf{x} \mapsto \frac{\|\mathbf{x}-\mathbf{v}\|_2^2}{2\sigma^2} = \frac{1}{2\sigma^2} \left(\|\mathbf{x}\|_2^2 - 2\langle \mathbf{x}, \mathbf{v} \rangle_2 + \|\mathbf{v}\|_2^2 \right) : \sigma > 0, \mathbf{v} \in \mathcal{X} \right\} \subseteq L^0(\mathbb{R}^d)$. $\mathcal{G} \subseteq \tilde{\mathcal{G}} := \text{span} \left(\mathbf{x} \mapsto \|\mathbf{x}\|_2^2, \{\mathbf{x} \mapsto x_i\}_{i=1}^d, \mathbf{x} \mapsto 1 \right)$ vector space, $\dim(\mathcal{G}) \leq d+2$. Thus by Lemma 7(i)-(ii), \mathcal{G} is VC with $VC(\mathcal{G}) \leq d+4$; applying Lemma 7(iii) with $\phi(z) = e^{-z}$, $\mathcal{F}_1 = \phi \circ \mathcal{G}$ is also VC with index $VC(\mathcal{F}_1) \leq d+4$.

- \mathcal{F}_2 : Since $\mathcal{F}_2 = \left\{ \mathbf{x} \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{x}, \mathbf{v}') = e^{-\frac{\|\mathbf{x}-\mathbf{v}\|_2^2 + \|\mathbf{x}-\mathbf{v}'\|_2^2}{2\sigma^2}} : \sigma > 0, \mathbf{v} \in \mathcal{X}, \mathbf{v}' \in \mathcal{X} \right\}$, and $\left\{ \mathbf{x} \mapsto \frac{\|\mathbf{x}-\mathbf{v}\|_2^2 + \|\mathbf{x}-\mathbf{v}'\|_2^2}{2\sigma^2} : \sigma > 0, \mathbf{v} \in \mathcal{X}, \mathbf{v}' \in \mathcal{X} \right\} \subseteq S = \text{span} \left(\mathbf{x} \mapsto \|\mathbf{x}\|_2^2, \{\mathbf{x} \mapsto x_i\}_{i=1}^d, \mathbf{x} \mapsto 1 \right)$, $VC(\mathcal{F}_2) \leq d+4$.

- \mathcal{F}_3 : Since

$$\mathcal{F}_3 = \left\{ (\mathbf{x}, \mathbf{y}) \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{y}, \mathbf{v}') = e^{-\frac{\|\mathbf{x}-\mathbf{v}\|_2^2 + \|\mathbf{y}-\mathbf{v}'\|_2^2}{2\sigma^2}} = e^{-\frac{\|[\mathbf{x}; \mathbf{y}] - [\mathbf{v}; \mathbf{v}']\|_2^2}{2\sigma^2}} : \sigma > 0, \mathbf{v} \in \mathbb{R}^d, \mathbf{v}' \in \mathbb{R}^d \right\},$$

from the result on \mathcal{F}_1 we get that $VC(\mathcal{F}_3) \leq 2d+4$.

Uniformly bounded, separable Carathéodory family:

The result follows from Lemma 4 by noting that $|k(\mathbf{x}, \mathbf{y})| \leq 1 =: B$, $(\mathbf{x}, \mathbf{y}) \mapsto e^{-\frac{\|\mathbf{x}-\mathbf{y}\|_2^2}{2\sigma^2}}$ is continuous ($\forall \sigma > 0$), \mathbb{R}^+ is separable, and the $(\sigma, \mathbf{v}) \mapsto e^{-\frac{\|\mathbf{x}-\mathbf{v}\|_2^2}{2\sigma^2}}$, $(\sigma, \mathbf{v}, \mathbf{v}') \mapsto e^{-\frac{\|\mathbf{x}-\mathbf{v}\|_2^2}{2\sigma^2}} e^{-\frac{\|\mathbf{x}-\mathbf{v}'\|_2^2}{2\sigma^2}}$, $(\sigma, \mathbf{v}, \mathbf{v}') \mapsto e^{-\frac{\|\mathbf{x}-\mathbf{v}\|_2^2}{2\sigma^2}} e^{-\frac{\|\mathbf{y}-\mathbf{v}'\|_2^2}{2\sigma^2}}$ mappings are continuous ($\forall \mathbf{x}, \mathbf{y} \in \mathcal{X}$). \square

Lemma 6 (\mathcal{F}_i -s are VC-subgraph and uniformly bounded separable Carathéodory families for full Gaussian kernel). *Let $\mathcal{K} = \{k_{\mathbf{A}} : (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{X} \subseteq \mathbb{R}^d \times \mathbb{R}^d \mapsto e^{-(\mathbf{x}-\mathbf{y})^\top \mathbf{A}(\mathbf{x}-\mathbf{y})} : \mathbf{A} \succeq \mathbf{0}\}$. Then the $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ classes [see Eqs. (1)-(2)] associated to \mathcal{K} are*

- VC-subgraphs with indices $VC(\mathcal{F}_1) \leq \frac{d(d+1)}{2} + d + 2$, $VC(\mathcal{F}_2) \leq \frac{d(d+1)+2}{2} + d + 2$, $VC(\mathcal{F}_3) \leq d(d+1) + 2d + 3$,
- uniformly bounded separable Carathéodory families, with $\|f\|_{L^\infty(\mathcal{M})} \leq 1$ for all $f \in \{\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3\}$.⁴

Proof. We prove the VC index values; the rest is essentially identical to the proof of Lemma 5.

- \mathcal{F}_1 : Using that $\mathcal{G} = \left\{ \mathbf{x} \mapsto (\mathbf{x} - \mathbf{v})^\top \mathbf{A}(\mathbf{x} - \mathbf{v}) : \mathbf{A} \succeq \mathbf{0}, \mathbf{v} \in \mathcal{X} \right\} \subseteq S := \text{span} \left(\{\mathbf{x} \mapsto x_i x_j\}_{1 \leq i \leq j \leq d}, \{\mathbf{x} \mapsto x_i\}_{1 \leq i \leq d}, \mathbf{x} \mapsto 1 \right)$, we have $VC(\mathcal{F}_1) \leq VC(\mathcal{G}) \leq \dim(S) + 2 \leq \frac{d(d+1)}{2} + d + 3$.
- \mathcal{F}_2 : Since $\mathcal{F}_2 = \left\{ \mathbf{x} \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{x}, \mathbf{v}') = e^{-[(\mathbf{x}-\mathbf{v})^\top \mathbf{A}(\mathbf{x}-\mathbf{v}) + (\mathbf{x}-\mathbf{v}')^\top \mathbf{A}(\mathbf{x}-\mathbf{v}')] } : \mathbf{A} \succeq \mathbf{0}, \mathbf{v} \in \mathcal{X}, \mathbf{v}' \in \mathcal{X} \right\}$, and

$$\begin{aligned} & \left\{ (\mathbf{x}, \mathbf{y}) \mapsto (\mathbf{x} - \mathbf{v})^\top \mathbf{A}(\mathbf{x} - \mathbf{v}) + (\mathbf{x} - \mathbf{v}')^\top \mathbf{A}(\mathbf{x} - \mathbf{v}') \right\} \subseteq S \\ & := \text{span} \left(\{\mathbf{x} \mapsto x_i x_j\}_{1 \leq i \leq j \leq d}, \{\mathbf{x} \mapsto x_i\}_{1 \leq i \leq d}, \mathbf{x} \mapsto 1 \right), \end{aligned}$$

we have $VC(\mathcal{F}_2) \leq VC(S) = \dim(S) + 2 \leq \frac{d(d+1)}{2} + d + 3$.

⁴ $\mathcal{M} = \mathcal{X}$ for \mathcal{F}_1 and \mathcal{F}_2 , and $\mathcal{M} = \mathcal{X}^2$ in case of \mathcal{F}_3 .

- \mathcal{F}_3 : Exploiting that

$$\mathcal{F}_3 = \left\{ (\mathbf{x}, \mathbf{y}) \mapsto k(\mathbf{x}, \mathbf{v})k(\mathbf{y}, \mathbf{v}') = e^{-[(\mathbf{x}-\mathbf{v})^\top \mathbf{A}(\mathbf{x}-\mathbf{v}) + (\mathbf{y}-\mathbf{v}')^\top \mathbf{B}(\mathbf{y}-\mathbf{v}')] } : \mathbf{A} \succeq \mathbf{0}, \mathbf{B} \succeq \mathbf{0}, \mathbf{v} \in \mathcal{X}, \mathbf{v}' \in \mathcal{X} \right\},$$

and $\{(\mathbf{x}, \mathbf{y}) \mapsto (\mathbf{x} - \mathbf{v})^\top \mathbf{A}(\mathbf{x} - \mathbf{v}) + (\mathbf{y} - \mathbf{v}')^\top \mathbf{B}(\mathbf{y} - \mathbf{v}')\} \subseteq S := \text{span}(\{(\mathbf{x}, \mathbf{y}) \mapsto x_i x_j\}_{1 \leq i \leq j \leq d}, \{(\mathbf{x}, \mathbf{y}) \mapsto x_i\}_{1 \leq i \leq d}, \{(\mathbf{x}, \mathbf{y}) \mapsto 1\}, \{(\mathbf{x}, \mathbf{y}) \mapsto y_i y_j\}_{1 \leq i \leq j \leq d}, \{(\mathbf{x}, \mathbf{y}) \mapsto y_i\}_{1 \leq i \leq d})$,
we have $VC(\mathcal{F}_3) \leq VC(S) = \dim(S) + 2 \leq d(d+1) + 2d + 3$.

□

F Proof of proposition 1

Recall Proposition 1:

Proposition 1 (Lower bound on ME test power). *Let \mathcal{K} be a uniformly bounded (i.e., $\exists B < \infty$ such that $\sup_{k \in \mathcal{K}} \sup_{(\mathbf{x}, \mathbf{y}) \in \mathcal{X}^2} |k(\mathbf{x}, \mathbf{y})| \leq B$) family of $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ measurable kernels. Let \mathbb{V} be a collection in which each element is a set of J test locations. Assume that $\tilde{c} := \sup_{\mathcal{V} \in \mathbb{V}, k \in \mathcal{K}} \|\Sigma^{-1}\|_F < \infty$. Then, the test power $\mathbb{P}(\hat{\lambda}_n \geq T_\alpha)$ of the ME test satisfies $\mathbb{P}(\hat{\lambda}_n \geq T_\alpha) \geq L(\lambda_n)$ where*

$$L(\lambda_n) := 1 - 2e^{-\xi_1(\lambda_n - T_\alpha)^2/n} - 2e^{-\frac{[\gamma_n(\lambda_n - T_\alpha)(n-1) - \xi_2 n]^2}{\xi_3 n(2n-1)^2}} - 2e^{-[(\lambda_n - T_\alpha)/3 - \tilde{c}_3 n \gamma_n]^2 \gamma_n^2 / \xi_4},$$

and $\tilde{c}_3, \xi_1, \dots, \xi_4$ are positive constants depending on only B, J and \tilde{c} . The parameter $\lambda_n := n\boldsymbol{\mu}^\top \Sigma^{-1} \boldsymbol{\mu}$ is the population counterpart of $\hat{\lambda}_n := n\bar{\mathbf{z}}_n^\top (\mathbf{S}_n + \gamma_n I)^{-1} \bar{\mathbf{z}}_n$ where $\boldsymbol{\mu} = \mathbb{E}_{\mathbf{xy}}[\mathbf{z}_1]$ and $\Sigma = \mathbb{E}_{\mathbf{xy}}[(\mathbf{z}_1 - \boldsymbol{\mu})(\mathbf{z}_1 - \boldsymbol{\mu})^\top]$. For large n , $L(\lambda_n)$ is increasing in λ_n .

F.1 Proof

By (3), we have

$$|\hat{\lambda}_n - \lambda_n| \leq \frac{\bar{c}_1 n}{\gamma_n} \|\Sigma - \mathbf{S}_n\|_F + \bar{c}_2 n \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 + \bar{c}_3 n \gamma_n. \quad (12)$$

We will bound each of the three terms in (12).

Bounding $\|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2$ (second term in (12))

Let $g(\mathbf{x}, \mathbf{y}, \mathbf{v}) := k(\mathbf{x}, \mathbf{v}) - k(\mathbf{y}, \mathbf{v})$. Define $\mathbf{v}^* := \arg \max_{\mathbf{v} \in \{\mathbf{v}_1, \dots, \mathbf{v}_J\}} \left| \frac{1}{n} \sum_{i=1}^n g(\mathbf{x}_i, \mathbf{y}_i, \mathbf{v}) - \mathbb{E}_{\mathbf{xy}}[g(\mathbf{x}, \mathbf{y}, \mathbf{v})] \right|$. Define $G_i := g(\mathbf{x}_i, \mathbf{y}_i, \mathbf{v}^*)$.

$$\begin{aligned} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 &= \sup_{\mathbf{b} \in B(1, \mathbf{0})} \langle \mathbf{b}, \bar{\mathbf{z}}_n - \boldsymbol{\mu} \rangle_2 \\ &\leq \sup_{\mathbf{b} \in B(1, \mathbf{0})} \sum_{j=1}^J |b_j| \left| \frac{1}{n} \sum_{i=1}^n [k(\mathbf{x}_i, \mathbf{v}_j) - k(\mathbf{y}_i, \mathbf{v}_j)] - \mathbb{E}_{\mathbf{xy}}[k(\mathbf{x}, \mathbf{v}_j) - k(\mathbf{y}, \mathbf{v}_j)] \right| \\ &= \sup_{\mathbf{b} \in B(1, \mathbf{0})} \sum_{j=1}^J |b_j| \left| \frac{1}{n} \sum_{i=1}^n g(\mathbf{x}_i, \mathbf{y}_i, \mathbf{v}_j) - \mathbb{E}_{\mathbf{xy}}[g(\mathbf{x}, \mathbf{y}, \mathbf{v}_j)] \right| \\ &\leq \left| \frac{1}{n} \sum_{i=1}^n G_i - \mathbb{E}_{\mathbf{xy}}[G_1] \right| \sup_{\mathbf{b} \in B(1, \mathbf{0})} \sum_{j=1}^J |b_j| \\ &\leq \sqrt{J} \left| \frac{1}{n} \sum_{i=1}^n G_i - \mathbb{E}_{\mathbf{xy}}[G_1] \right| \sup_{\mathbf{b} \in B(1, \mathbf{0})} \|\mathbf{b}\|_2 \\ &= \sqrt{J} \left| \frac{1}{n} \sum_{i=1}^n G_i - \mathbb{E}_{\mathbf{xy}}[G_1] \right|, \end{aligned}$$

where we used the fact that $\|\mathbf{b}\|_1 \leq \sqrt{J}\|\mathbf{b}\|_2$. It can be seen that $-2B \leq G_i \leq 2B$ because

$$G_i = k(\mathbf{x}_i, \mathbf{v}^*) - k(\mathbf{y}_i, \mathbf{v}^*) \leq |k(\mathbf{x}_i, \mathbf{v}^*)| + |k(\mathbf{y}_i, \mathbf{v}^*)| \leq 2B.$$

Using Hoeffding's inequality (Lemma 9) to bound $|\frac{1}{n} \sum_{i=1}^n G_i - \mathbb{E}_{\mathbf{xy}}[G_1]|$, we have

$$\mathbb{P}(n\bar{c}_2\|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 \leq \alpha) \geq 1 - 2 \exp\left(-\frac{\alpha^2}{8B^2\bar{c}_2^2 Jn}\right). \quad (13)$$

Bounding first ($\|\boldsymbol{\Sigma} - \mathbf{S}_n\|_F$) and third terms in (12)

Let $\eta(\mathbf{v}_i, \mathbf{v}_j) := |\frac{1}{n} \sum_{a=1}^n g(\mathbf{x}_a, \mathbf{y}_a, \mathbf{v}_i)g(\mathbf{x}_a, \mathbf{y}_a, \mathbf{v}_j) - \mathbb{E}_{\mathbf{xy}}[g(\mathbf{x}, \mathbf{y}, \mathbf{v}_i)g(\mathbf{x}, \mathbf{y}, \mathbf{v}_j)]|$. Define $(\mathbf{v}_1^*, \mathbf{v}_2^*) = \arg \max_{(\mathbf{v}^{(1)}, \mathbf{v}^{(2)}) \in \{(\mathbf{v}_i, \mathbf{v}_j)\}_{i,j=1}^J} \eta(\mathbf{v}^{(1)}, \mathbf{v}^{(2)})$. Define $H_i := g(\mathbf{x}_i, \mathbf{y}_i, \mathbf{v}_1^*)g(\mathbf{x}_i, \mathbf{y}_i, \mathbf{v}_2^*)$. By (8), we have

$$\begin{aligned} \|\mathbf{S}_n - \boldsymbol{\Sigma}\|_F &\leq (*_1) + (*_2), \\ (*_1) &= \left\| \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a \mathbf{z}_a^\top - \mathbb{E}_{\mathbf{xy}}[\mathbf{z}_1 \mathbf{z}_1^\top] \right\|_F, \\ (*_2) &= \frac{8B^2 J}{n-1} + 2B_k \sqrt{J} \frac{2n-1}{n-1} \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2. \end{aligned}$$

We can upper bound $(*_2)$ by applying Hoeffding's inequality to bound $\|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2$ giving

$$\mathbb{P}\left(\frac{\bar{c}_1 n}{\gamma_n} (*_2) \leq \alpha\right) \geq 1 - 2 \exp\left(-\frac{(\alpha \gamma_n - \alpha \gamma_n n + 8B^2 \bar{c}_1 Jn)^2}{32B^4 \bar{c}_1^2 J^2 n(2n-1)^2}\right). \quad (14)$$

We can upper bound $(*_1)$ with

$$\begin{aligned} (*_1) &= \sup_{\mathbf{B} \in B(1, \mathbf{0})} \left\langle \mathbf{B}, \frac{1}{n} \sum_{a=1}^n \mathbf{z}_a \mathbf{z}_a^\top - \mathbb{E}_{\mathbf{xy}}[\mathbf{z}_1 \mathbf{z}_1^\top] \right\rangle_F \\ &\leq \sup_{\mathbf{B} \in B(1, \mathbf{0})} \sum_{i=1}^J \sum_{j=1}^J |B_{ij}| \left| \frac{1}{n} \sum_{a=1}^n g(\mathbf{x}_a, \mathbf{y}_a, \mathbf{v}_i)g(\mathbf{x}_a, \mathbf{y}_a, \mathbf{v}_j) - \mathbb{E}_{\mathbf{xy}}[g(\mathbf{x}, \mathbf{y}, \mathbf{v}_i)g(\mathbf{x}, \mathbf{y}, \mathbf{v}_j)] \right| \\ &\leq \left| \frac{1}{n} \sum_{a=1}^n H_a - \mathbb{E}_{\mathbf{xy}}[H_1] \right| \sup_{\mathbf{B} \in B(1, \mathbf{0})} \sum_{i=1}^J \sum_{j=1}^J |B_{ij}| \\ &\leq J \left| \frac{1}{n} \sum_{a=1}^n H_a - \mathbb{E}_{\mathbf{xy}}[H_1] \right| \sup_{\mathbf{B} \in B(1, \mathbf{0})} \|\mathbf{B}\|_F = J \left| \frac{1}{n} \sum_{a=1}^n H_a - \mathbb{E}_{\mathbf{xy}}[H_1] \right|, \end{aligned}$$

where we used the fact that $\sum_{i=1}^J \sum_{j=1}^J |B_{ij}| \leq J\|\mathbf{B}\|_F$. It can be seen that $-4B^2 \leq H_a \leq 4B^2$. Using Hoeffding's inequality (Lemma 9) to bound $|\frac{1}{n} \sum_{a=1}^n H_a - \mathbb{E}_{\mathbf{xy}}[H_1]|$, we have

$$\mathbb{P}\left(\frac{\bar{c}_1 n}{\gamma_n} (*_1) \leq \alpha\right) \geq 1 - 2 \exp\left(-\frac{\alpha^2 \gamma_n^2}{32B^4 J^2 \bar{c}_1^2 n}\right), \quad (15)$$

implying that

$$\mathbb{P}\left(\frac{\bar{c}_1 n}{\gamma_n} (*_1) + \bar{c}_3 n \gamma_n \leq \alpha\right) \geq 1 - 2 \exp\left(-\frac{(\alpha - \bar{c}_3 n \gamma_n)^2 \gamma_n^2}{32B^4 J^2 \bar{c}_1^2 n}\right). \quad (16)$$

Applying a union bound on (13), (14), and (16) with $t = \alpha/3$, we can conclude that

$$\begin{aligned} \mathbb{P}\left(\left|\hat{\lambda}_n - \lambda_n\right| \leq t\right) &\geq \mathbb{P}\left(\frac{\bar{c}_1 n}{\gamma_n} \|\boldsymbol{\Sigma} - \mathbf{S}_n\|_F + \bar{c}_2 n \|\bar{\mathbf{z}}_n - \boldsymbol{\mu}\|_2 + \bar{c}_3 n \gamma_n \leq t\right) \\ &\geq 1 - 2 \exp\left(-\frac{t^2}{32 \cdot 8B^2 \bar{c}_2^2 Jn}\right) - 2 \exp\left(-\frac{(t \gamma_n n - t \gamma_n - 24B^2 \bar{c}_1 Jn)^2}{32 \cdot 32B^4 \bar{c}_1^2 J^2 n(2n-1)^2}\right) - 2 \exp\left(-\frac{(t/3 - \bar{c}_3 n \gamma_n)^2 \gamma_n^2}{32B^4 J^2 \bar{c}_1^2 n}\right). \end{aligned}$$

A rearrangement yields

$$\mathbb{P}\left(\hat{\lambda}_n \geq T_\alpha\right)$$

$$\geq 1 - 2 \exp \left(-\frac{(\lambda_n - T_\alpha)^2}{3^2 \cdot 8B^2 \bar{c}_2^2 J n} \right) - 2 \exp \left(-\frac{(\gamma_n(\lambda_n - T_\alpha)(n-1) - 24B^2 \bar{c}_1 J n)^2}{3^2 \cdot 32B^4 \bar{c}_1^2 J^2 n(2n-1)^2} \right) - 2 \exp \left(-\frac{((\lambda_n - T_\alpha)/3 - \bar{c}_3 n \gamma_n)^2 \gamma_n^2}{32B^4 J^2 \bar{c}_1^2 n} \right).$$

Define $\xi_1 := \frac{1}{3^2 \cdot 8B^2 \bar{c}_2^2 J}$, $\xi_2 := 24B^2 \bar{c}_1 J$, $\xi_3 := 3^2 \cdot 32B^4 \bar{c}_1^2 J^2$, $\xi_4 := 32B^4 J^2 \bar{c}_1^2$. We have

$$\begin{aligned} & \mathbb{P}(\hat{\lambda}_n \geq T_\alpha) \\ & \geq 1 - 2 \exp \left(-\frac{\xi_1(\lambda_n - T_\alpha)^2}{n} \right) - 2 \exp \left(-\frac{(\gamma_n(\lambda_n - T_\alpha)(n-1) - \xi_2 n)^2}{\xi_3 n(2n-1)^2} \right) - 2 \exp \left(-\frac{((\lambda_n - T_\alpha)/3 - \bar{c}_3 n \gamma_n)^2 \gamma_n^2}{\xi_4} \right). \end{aligned}$$

□

G External lemmas

In this section we detail some external lemmas used in our proof.

Lemma 7 (properties of VC classes, see page 141, 146-147 in van der Vaart and Wellner (2000) and page 160-161 in Kosorok (2008)).

- (i) *Monotonicity:* $\mathcal{G} \subseteq \tilde{\mathcal{G}} \subseteq L^0(\mathcal{M}) \Rightarrow VC(\mathcal{G}) \leq VC(\tilde{\mathcal{G}})$.
- (ii) *Finite-dimensional vector space:* if \mathcal{G} is a finite-dimensional vector space of measurable functions, then $VC(\mathcal{G}) \leq \dim(\mathcal{G}) + 2$.
- (iii) *Composition with monotone function:* If \mathcal{G} is VC and $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is monotone, then for $\phi \circ \mathcal{G} := \{\phi \circ g : g \in \mathcal{G}\}$, $VC(\phi \circ \mathcal{G}) \leq VC(\mathcal{G})$.
- (iv) *The r -covering number of a VC class grows only polynomially in $\frac{1}{r}$:* Let \mathcal{F} be VC on the domain \mathcal{M} with measurable envelope F ($|f(m)| \leq F(m)$, $\forall m \in \mathcal{M}, f \in \mathcal{F}$). Then for any $q \geq 1$ and \mathbb{M} probability measure for which $\|F\|_{L^q(\mathcal{M}, \mathbb{M})} > 0$

$$N \left(r \|F\|_{L^q(\mathcal{M}, \mathbb{M})}, \mathcal{F}, L^q(\mathcal{M}, \mathbb{M}) \right) \leq C \times VC(\mathcal{F}) (16e)^{VC(\mathcal{F})} \left(\frac{1}{r} \right)^{q[VC(\mathcal{F})-1]} \quad (17)$$

for any $r \in (0, 1)$ with a universal constant C .

Lemma 8 (McDiarmid's inequality). Let $X_1, \dots, X_n \in \mathcal{M}$ be independent random variables and let $g : \mathcal{M}^n \rightarrow \mathbb{R}$ be a function such that the

$$\sup_{\mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{x}'_j \in \mathcal{M}} |g(\mathbf{x}_1, \dots, \mathbf{x}_n) - g(\mathbf{x}_1, \dots, \mathbf{x}_j, \mathbf{x}'_j, \mathbf{x}_{j+1}, \dots, \mathbf{x}_n)| \leq b \quad (18)$$

bounded difference property holds. Then for arbitrary $\delta \in (0, 1)$

$$\mathbb{P} \left(g(X_1, \dots, X_n) \leq \mathbb{E}[g(X_1, \dots, X_n)] + b \sqrt{\frac{\log \left(\frac{1}{\delta} \right) n}{2}} \right) \geq 1 - \delta.$$

Lemma 9 (Hoeffding's inequality). Let X_1, \dots, X_n be i.i.d. random variables with $\mathbb{P}(a \leq X_i \leq b) = 1$. Let $\bar{X} := \frac{1}{n} \sum_{i=1}^n X_i$. Then,

$$\mathbb{P}(|\bar{X} - \mathbb{E}[\bar{X}]| \leq t) \geq 1 - 2 \exp \left(-\frac{2nt^2}{(b-a)^2} \right).$$

Equivalently, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds that

$$|\bar{X} - \mathbb{E}[\bar{X}]| \leq \frac{b-a}{\sqrt{2n}} \sqrt{\log(2/\delta)}.$$