



HAL
open science

An Adaptive Test of Independence with Analytic Kernel Embeddings

Wittawat Jitkrittum, Zoltán Szabó, Arthur Gretton

► **To cite this version:**

Wittawat Jitkrittum, Zoltán Szabó, Arthur Gretton. An Adaptive Test of Independence with Analytic Kernel Embeddings. [Technical Report] University College London; Ecole Polytechnique (Palaiseau, France). 2016. hal-01385111v1

HAL Id: hal-01385111

<https://hal.science/hal-01385111v1>

Submitted on 20 Oct 2016 (v1), last revised 28 Jul 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An Adaptive Test of Independence with Analytic Kernel Embeddings

Wittawat Jitkrittum,¹ Zoltán Szabó,^{2*} Arthur Gretton¹

¹Gatsby Computational Neuroscience Unit, University College London

²Department of Applied Mathematics, CMAP, École Polytechnique

wittawat@gatsby.ucl.ac.uk
zoltan.szabo@polytechnique.edu
arthur.gretton@gmail.com

October 20, 2016

Abstract

A new computationally efficient dependence measure, and an adaptive statistical test of independence, are proposed. The dependence measure is the difference between analytic embeddings of the joint distribution and the product of the marginals, evaluated at a finite set of locations (features). These features are chosen so as to maximize a lower bound on the test power, resulting in a test that is data-efficient, and that runs in linear time (with respect to the sample size n). The optimized features can be interpreted as evidence to reject the null hypothesis, indicating regions in the joint domain where the joint distribution and the product of the marginals differ most. Consistency of the independence test is established, for an appropriate choice of features. In real-world benchmarks, independence tests using the optimized features perform comparably to the state-of-the-art quadratic-time HSIC test, and outperform competing $\mathcal{O}(n)$ and $\mathcal{O}(n \log n)$ tests.

1 Introduction

We consider the design of adaptive, nonparametric statistical tests of dependence: that is, tests of whether a joint distribution P_{xy} factorizes into the product of marginals $P_x P_y$. While classical tests of dependence, such as Pearson’s correlation and Kendall’s τ , are able to detect monotonic relations between univariate variables, more modern tests can address complex interactions, for instance changes in variance of X with the value of Y . Key to many recent tests is to examine covariance or correlation between data features. These interactions become significantly harder to detect, and the features

are more difficult to design when the data reside in high dimensions.

A basic nonlinear dependence measure is the Hilbert-Schmidt Independence Criterion (HSIC), which is the Hilbert-Schmidt norm of the covariance operator between feature mappings of the random variables [Gretton et al., 2005, 2008]. Each random variable X and Y is mapped to a respective reproducing kernel Hilbert space \mathcal{H}_k and \mathcal{H}_l . For sufficiently rich mappings, the covariance operator norm is zero if and only if the variables are independent. A second basic nonlinear dependence measure is the smoothed difference between the characteristic function of the joint distribution, and that of the product of marginals. When a particular smoothing function is used, the statistic corresponds to the covariance between distances of X and Y variable pairs [Feuerverger, 1993, Székely et al., 2007, Székely and Rizzo, 2009], yielding a simple test statistic. It has been shown by Sejdinovic et al. [2013] that the distance covariance (and its generalization to semi-metrics) is an instance of HSIC for an appropriate choice of kernels. A disadvantage of these feature covariance statistics, however, is that they require quadratic time to compute (besides in the special case of the distance covariance with univariate real-valued variables, where Huo and Székely [2014] achieve an $\mathcal{O}(n \log n)$ cost). Moreover, the feature covariance statistics have intractable null distributions, and either a permutation approach or the solution of an expensive eigenvalue problem [e.g. Zhang et al., 2011] is required for consistent estimation of the quantiles. Several approaches were proposed by Zhang et al. [2016] to obtain faster tests along the lines of HSIC. These include computing HSIC on finite-dimensional feature mappings chosen as random Fourier features (RFFs) [Rahimi and Recht, 2008], a block-averaged statistic, and a Nyström approximation to the statistic. Key to each of these approaches is a more efficient computation of the statistic and its threshold un-

*Zoltán Szabó’s ORCID ID: <http://orcid.org/0000-0001-6183-7603>.

der the null distribution: for RFFs, the null distribution is a finite weighted sum of χ^2 variables; for the block-averaged statistic, the null distribution is asymptotically normal; for Nyström, either a permutation approach is employed, or the spectrum of the Nyström approximation to the kernel matrix is used in approximating the null distribution. Each of these methods costs significantly less than the $\mathcal{O}(n^2)$ cost of the full HSIC (the cost is linear in n , but also depends quadratically on the number of features retained). A potential disadvantage of the Nyström and Fourier approaches is that the features are not optimized to maximize test power, but are chosen randomly. The block statistic performs worse than both, due to the large variance of the statistic under the null (which can be mitigated by observing more data).

In addition to feature covariances, correlation measures have also been developed in infinite dimensional feature spaces: in particular, Bach and Jordan [2002], Fukumizu et al. [2008] proposed statistics on the correlation operator in a reproducing kernel Hilbert space. While convergence has been established for certain of these statistics, their computational cost is high at $\mathcal{O}(n^3)$, and test thresholds have relied on permutation. A number of much faster approaches to testing based on feature correlations have been proposed, however. For instance, Dauxois and Nkiet [1998] compute statistics of the correlation between finite sets of basis functions, chosen for instance to be step functions or low order B-splines. The cost of this approach is $\mathcal{O}(n)$. This idea was extended by Lopez-Paz et al. [2013], who computed the canonical correlation between finite sets of basis functions chosen as random Fourier features; in addition, they performed a copula transform on the inputs, with a total cost of $\mathcal{O}(n \log n)$. Finally, space partitioning approaches have also been proposed, based on statistics such as the KL divergence, however these apply only to univariate variables [Heller et al., 2016], or to multivariate variables of low dimension [Gretton and Györfi, 2010] (that said, these tests have other advantages of theoretical interest, notably distribution-independent test thresholds).

The approach we take is most closely related to HSIC on a finite set of features. Our simplest test statistic, the Finite Set Independence Criterion (FSIC), is an average of covariances of analytic functions (i.e., features) defined on each of X and Y . A normalized version of the statistic (NFSIC) yields a distribution-independent asymptotic test threshold. We show that our test is consistent, despite a finite number of analytic features being used, via a generalization of arguments in Chwialkowski et al. [2015]. As in recent work on two-sample testing by Jitkrittum et al. [2016], our test is *adaptive* in the sense that we choose our features on a held-out validation set to optimize a lower bound on the test power. The design of features for independence testing turns out

to be quite different to the case of two-sample testing, however: the task is to find correlated feature *pairs* on the respective marginal domains, rather than attempting to find a single, high-dimensional feature representation for the *entire* (x, y) (as we would need to do if we were comparing distributions P_{xy} and Q_{xy} , rather than testing a specific property of P_{xy}). We demonstrate the performance of our tests on several challenging artificial and real-world datasets, including detection of dependence between music and its year of appearance, and between videos and captions. In these experiments, we outperform competing linear and $\mathcal{O}(n \log n)$ time tests.

2 Independence Criteria and Statistical Tests

We introduce two test statistics: first, the Finite Set Independence Criterion (FSIC), which builds on the principle that dependence can be measured in terms of the covariance between data features. Next, we propose a normalized version of this statistic (NFSIC), with a simpler asymptotic distribution when $P_{xy} = P_x P_y$. We show how to select features for the latter statistic to maximize a lower bound on the power of its corresponding statistical test.

2.1 The Finite Set Independence Criterion

We begin by introducing the Hilbert-Schmidt Independence Criterion (HSIC) as proposed in Gretton et al. [2005], since our unnormalized statistic is built along similar lines. Consider two random variables $X \in \mathcal{X} \subset \mathbb{R}^{d_x}$ and $Y \in \mathcal{Y} \subset \mathbb{R}^{d_y}$. Denote by P_{xy} the joint distribution between X and Y ; P_x and P_y are the marginal distributions of X and Y . Let \otimes denote the tensor product, such that $(a \otimes b)c = a \langle b, c \rangle$. Assume that $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $l : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ are positive definite kernels associated with reproducing kernel Hilbert spaces (RKHS) \mathcal{H}_k and \mathcal{H}_l , respectively. Let $\|\cdot\|_{HS}$ be the norm on the space of $\mathcal{H}_l \rightarrow \mathcal{H}_k$ Hilbert-Schmidt operators. Then, HSIC between X and Y is defined as

$$\begin{aligned} \text{HSIC}(X, Y) &= \|\mu_{xy} - \mu_x \otimes \mu_y\|_{\text{HS}}^2 \\ &= \mathbb{E}_{(\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}')} [k(\mathbf{x}, \mathbf{x}')l(\mathbf{y}, \mathbf{y}')] \\ &\quad + \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{x}'} [k(\mathbf{x}, \mathbf{x}')] \mathbb{E}_{\mathbf{y}} \mathbb{E}_{\mathbf{y}'} [l(\mathbf{y}, \mathbf{y}')] \\ &\quad - 2 \mathbb{E}_{(\mathbf{x}, \mathbf{y})} [\mathbb{E}_{\mathbf{x}'} [k(\mathbf{x}, \mathbf{x}')] \mathbb{E}_{\mathbf{y}'} [l(\mathbf{y}, \mathbf{y}')]], \end{aligned} \quad (1)$$

where $\mathbb{E}_{\mathbf{x}} := \mathbb{E}_{\mathbf{x} \sim P_x}$, $\mathbb{E}_{\mathbf{y}} := \mathbb{E}_{\mathbf{y} \sim P_y}$, $\mathbb{E}_{(\mathbf{x}, \mathbf{y})} := \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P_{xy}}$, and \mathbf{x}' is an independent copy of \mathbf{x} . The mean embedding of P_{xy} belongs to the space of Hilbert-Schmidt operators from \mathcal{H}_l to \mathcal{H}_k , $\mu_{xy} := \int_{\mathcal{X} \times \mathcal{Y}} k(\mathbf{x}, \cdot) \otimes l(\mathbf{y}, \cdot) dP_{xy}(\mathbf{x}, \mathbf{y}) \in \text{HS}(\mathcal{H}_l, \mathcal{H}_k)$, and the marginal mean

embeddings are $\mu_x := \int_{\mathcal{X}} k(\mathbf{x}, \cdot) dP_x(\mathbf{x}) \in \mathcal{H}_k$ and $\mu_y := \int_{\mathcal{Y}} l(\mathbf{y}, \cdot) dP_y(\mathbf{y}) \in \mathcal{H}_l$ [Smola et al., 2007]. Gretton et al. [2005, Theorem 4] show that if the kernels k and l are universal [Steinwart and Christmann, 2008] on compact domains \mathcal{X} and \mathcal{Y} , then $\text{HSIC}(X, Y) = 0$ if and only if X and Y are independent. Alternatively, Gretton [2015] shows that it is sufficient for each of k and l to be characteristic to their respective domains (meaning that distribution embeddings are injective in each marginal domain: see Sriperumbudur et al. [2010]). Given a joint sample $Z_n = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n \sim P_{xy}$, an empirical estimator of HSIC can be computed in $\mathcal{O}(n^2)$ time by replacing the population expectations in (1) with their corresponding empirical expectations based on Z_n .

We now propose our new linear-time dependence measure, the Finite Set Independence Criterion (FSIC). Let $\mathcal{X} \subset \mathbb{R}^{d_x}$ and $\mathcal{Y} \subset \mathbb{R}^{d_y}$ be open sets. Define the empirical measure $\nu := \frac{1}{J} \sum_{i=1}^J \delta_{(\mathbf{v}_i, \mathbf{w}_i)}$ over J test locations $V_J := \{(\mathbf{v}_i, \mathbf{w}_i)\}_{i=1}^J \subset \mathcal{X} \times \mathcal{Y}$ where $\delta_{\mathbf{t}}$ denotes the Dirac measure centered on \mathbf{t} , and $(\mathbf{v}_i, \mathbf{w}_i)$ are realizations from an absolutely continuous distribution (wrt the Lebesgue measure). Write $\mathbb{E}_{\mathbf{xy}}$ for $\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P_{xy}}$. The idea is to see $\mu_{xy}(\mathbf{v}, \mathbf{w}) = \mathbb{E}_{\mathbf{xy}}[k(\mathbf{x}, \mathbf{v})l(\mathbf{y}, \mathbf{w})]$, $\mu_x(\mathbf{v}) = \mathbb{E}_{\mathbf{x}}[k(\mathbf{x}, \mathbf{v})]$ and $\mu_y(\mathbf{w}) = \mathbb{E}_{\mathbf{y}}[l(\mathbf{y}, \mathbf{w})]$ as smooth functions, and consider an $L^2(\mathcal{X} \times \mathcal{Y}, \nu)$ distance between μ_{xy} and $\mu_x \mu_y$ instead of a Hilbert-Schmidt distance as in HSIC [Gretton et al., 2005]. Let $\mu_x \mu_y(\mathbf{x}, \mathbf{y}) := \mu_x(\mathbf{x})\mu_y(\mathbf{y})$. FSIC is defined as

$$\begin{aligned} \text{FSIC}^2(X, Y) &:= \|\mu_{xy} - \mu_x \mu_y\|_{L^2(\mathcal{X} \times \mathcal{Y}, \nu)}^2 \\ &= \int_{\mathcal{X}} \int_{\mathcal{Y}} (\mu_{xy}(\mathbf{x}, \mathbf{y}) - \mu_x(\mathbf{x})\mu_y(\mathbf{y}))^2 d\nu(\mathbf{x}, \mathbf{y}) \\ &:= \frac{1}{J} \sum_{i=1}^J u(\mathbf{v}_i, \mathbf{w}_i)^2 = \frac{1}{J} \|\mathbf{u}\|_2^2, \text{ where} \end{aligned}$$

$$\begin{aligned} u(\mathbf{v}, \mathbf{w}) &:= \mu_{xy}(\mathbf{v}, \mathbf{w}) - \mu_x(\mathbf{v})\mu_y(\mathbf{w}) \\ &= \mathbb{E}_{\mathbf{xy}}[k(\mathbf{x}, \mathbf{v})l(\mathbf{y}, \mathbf{w})] - \mathbb{E}_{\mathbf{x}}[k(\mathbf{x}, \mathbf{v})]\mathbb{E}_{\mathbf{y}}[l(\mathbf{y}, \mathbf{w})], \quad (2) \\ &= \text{cov}_{\mathbf{xy}}[k(\mathbf{x}, \mathbf{v}), l(\mathbf{y}, \mathbf{w})], \end{aligned}$$

and $\mathbf{u} := (u(\mathbf{v}_1, \mathbf{w}_1), \dots, u(\mathbf{v}_J, \mathbf{w}_J))^\top$.

Our first result in Proposition 2 states that $\text{FSIC}(X, Y)$ almost surely defines a dependence measure for the random variables X and Y , provided that the product kernel on the joint space $\mathcal{X} \times \mathcal{Y}$ is characteristic and analytic (see Definition 1).

Definition 1 (Analytic kernels [Chwialkowski et al., 2015]). Let \mathcal{X} be an open set in \mathbb{R}^d . A positive definite kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is said to be analytic on its domain $\mathcal{X} \times \mathcal{X}$ if for all $\mathbf{v} \in \mathcal{X}$, $f(\mathbf{x}) := k(\mathbf{x}, \mathbf{v})$ is an analytic function on \mathcal{X} .

Assumption A. The kernels $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $l : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ are bounded by B_k and B_l respectively, and

the product kernel $g((\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}')) := k(\mathbf{x}, \mathbf{x}')l(\mathbf{y}, \mathbf{y}')$ is characteristic [Sriperumbudur et al., 2010, Definition 6], and analytic (Definition 1) on $(\mathcal{X} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{Y})$.

Proposition 2 (FSIC is a dependence measure). Assume that

1. Assumption A holds.
2. The test locations $V_J = \{(\mathbf{v}_i, \mathbf{w}_i)\}_{i=1}^J$ are drawn from an absolutely continuous distribution.

Then, almost surely, $\text{FSIC}(X, Y) = \frac{1}{\sqrt{J}} \|\mathbf{u}\|_2 = 0$ if and only if X and Y are independent.

Proof. Since g is characteristic, the mean embedding map $\Pi_g : P \mapsto \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P}[g((\mathbf{x}, \mathbf{y}), \cdot)]$ is injective [Sriperumbudur et al., 2010, Section 3], where P is a probability distribution on $\mathcal{X} \times \mathcal{Y}$. Since g is analytic, by Lemma 10 (Appendix), μ_{xy} and $\mu_x \mu_y$ are analytic functions. Thus, Lemma 11 (Appendix, setting $\Lambda = \Pi_g$) guarantees that $\text{FSIC}(X, Y) = 0 \iff P_{xy} = P_x P_y \iff X$ and Y are independent almost surely. \square

FSIC uses μ_{xy} as a proxy for P_{xy} , and $\mu_x \mu_y$ as a proxy for $P_x P_y$. Proposition 2 suggests that, to detect the dependence between X and Y , it is sufficient to evaluate at a finite number of locations (defined by V_J) the difference of the population joint embedding μ_{xy} and the embedding of the product of the marginal distributions $\mu_x \mu_y$. A brief explanation to justify this property is as follows. If $P_{xy} = P_x P_y$, then $\rho(\mathbf{v}, \mathbf{w}) := \mu_{xy}(\mathbf{v}, \mathbf{w}) - \mu_x \mu_y(\mathbf{v}, \mathbf{w})$ is zero, and $\text{FSIC}(X, Y) = 0$ for any V_J . If $P_{xy} \neq P_x P_y$, then ρ will not be a zero function, since the mean embedding map is injective (require the product kernel to be characteristic). Using the same argument as in Chwialkowski et al. [2015], since k and l are analytic, ρ is also analytic, and the set of roots $R := \{(\mathbf{v}, \mathbf{w}) \mid \rho(\mathbf{v}, \mathbf{w}) = 0\}$ has Lebesgue measure zero. Thus, it is sufficient to draw (\mathbf{v}, \mathbf{w}) from an absolutely continuous distribution, as we are guaranteed that $(\mathbf{v}, \mathbf{w}) \notin R$ giving $\text{FSIC}(X, Y) > 0$.

For FSIC to be a dependence measure, the product kernel is required to be characteristic and analytic. We next show in Proposition 3 that Gaussian kernels k and l yield such a product kernel.

Proposition 3 (A product of Gaussian kernels is characteristic and analytic). Let $k(\mathbf{x}, \mathbf{x}') = \exp(-(\mathbf{x} - \mathbf{x}')^\top \mathbf{A}(\mathbf{x} - \mathbf{x}'))$ and $l(\mathbf{y}, \mathbf{y}') = \exp(-(\mathbf{y} - \mathbf{y}')^\top \mathbf{B}(\mathbf{y} - \mathbf{y}'))$ be Gaussian kernels on $\mathbb{R}^{d_x} \times \mathbb{R}^{d_x}$ and $\mathbb{R}^{d_y} \times \mathbb{R}^{d_y}$ respectively, for positive definite matrices \mathbf{A} and \mathbf{B} . Then, $g((\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}')) = k(\mathbf{x}, \mathbf{x}')l(\mathbf{y}, \mathbf{y}')$ is characteristic and analytic on $(\mathbb{R}^{d_x} \times \mathbb{R}^{d_y}) \times (\mathbb{R}^{d_x} \times \mathbb{R}^{d_y})$.

Proof (sketch). The main idea is to use the fact a Gaussian kernel is analytic, and a product of Gaussian kernels is a Gaussian kernel on the pair of variables. See the full proof in Appendix D. \square

Plug-in Estimator We now give an empirical estimator of FSIC. Assume that we observe a joint sample $Z_n := \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n \stackrel{i.i.d.}{\sim} P_{xy}$. Unbiased estimators of $\mu_{xy}(\mathbf{v}, \mathbf{w})$ and $\mu_x \mu_y(\mathbf{v}, \mathbf{w})$ are $\hat{\mu}_{xy}(\mathbf{v}, \mathbf{w}) := \frac{1}{n} \sum_{i=1}^n k(\mathbf{x}_i, \mathbf{v}) l(\mathbf{y}_i, \mathbf{w})$ and $\hat{\mu}_x \hat{\mu}_y(\mathbf{v}, \mathbf{w}) := \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i} k(\mathbf{x}_i, \mathbf{v}) l(\mathbf{y}_j, \mathbf{w})$, respectively. A straightforward empirical estimator of $\widehat{\text{FSIC}}^2$ is then given by

$$\begin{aligned} \widehat{\text{FSIC}}^2(Z_n) &= \frac{1}{J} \sum_{i=1}^J \hat{u}(\mathbf{v}_i, \mathbf{w}_i)^2, \\ \hat{u}(\mathbf{v}, \mathbf{w}) &:= \hat{\mu}_{xy}(\mathbf{v}, \mathbf{w}) - \hat{\mu}_x \hat{\mu}_y(\mathbf{v}, \mathbf{w}) \\ &= \frac{2}{n(n-1)} \sum_{i < j} h_{(\mathbf{v}, \mathbf{w})}((\mathbf{x}_i, \mathbf{y}_i), (\mathbf{x}_j, \mathbf{y}_j)), \end{aligned} \quad (3) \quad (4)$$

where $h_{(\mathbf{v}, \mathbf{w})}((\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}')) := \frac{1}{2}(k(\mathbf{x}, \mathbf{v}) - k(\mathbf{x}', \mathbf{v}))(l(\mathbf{y}, \mathbf{w}) - l(\mathbf{y}', \mathbf{w}))$. For conciseness, we define $\hat{\mathbf{u}} := (\hat{u}_1, \dots, \hat{u}_J)^\top \in \mathbb{R}^J$ where $\hat{u}_i := \hat{u}(\mathbf{v}_i, \mathbf{w}_i)$ so that $\widehat{\text{FSIC}}^2(Z_n) = \frac{1}{J} \hat{\mathbf{u}}^\top \hat{\mathbf{u}}$.

$\widehat{\text{FSIC}}^2$ can be efficiently computed in $\mathcal{O}((d_x + d_y)Jn)$ time [see (3)], assuming that the runtime complexity of evaluating $k(\mathbf{x}, \mathbf{v})$ is $\mathcal{O}(d_x)$ and that of $l(\mathbf{y}, \mathbf{w})$ is $\mathcal{O}(d_y)$. The unbiasedness of $\hat{\mu}_x \hat{\mu}_y$ is necessary for (4) to be a U-statistic. This fact and the rewriting of $\widehat{\text{FSIC}}^2$ in terms of $h_{(\mathbf{v}, \mathbf{w})}((\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}'))$ will be exploited when the asymptotic distribution of $\hat{\mathbf{u}}$ is derived (Proposition 4).

Since FSIC satisfies $\text{FSIC}(X, Y) = 0 \iff X \perp Y$, in principle its empirical estimator can be used as a test statistic for an independence test proposing a null hypothesis H_0 : “ X and Y are independent” against an alternative H_1 : “ X and Y are dependent”. The null distribution (i.e., distribution of the test statistic assuming that H_0 is true) is challenging to obtain, however and depends on the unknown P_{xy} . This prompts us to consider a normalized version of FSIC whose asymptotic null distribution is a convenient form. We first derive the asymptotic distribution of $\hat{\mathbf{u}}$ in Proposition 4, which we use to derive the normalized test statistic in Theorem 5. As a shorthand, we write $\mathbf{z} := (\mathbf{x}, \mathbf{y})$, and $\mathbf{t} := (\mathbf{v}, \mathbf{w})$.

Proposition 4 (Asymptotic distribution of $\hat{\mathbf{u}}$). *Define $\tilde{k}(\mathbf{x}, \mathbf{v}) := k(\mathbf{x}, \mathbf{v}) - \mathbb{E}_{\mathbf{x}'} k(\mathbf{x}', \mathbf{v})$, and $\tilde{l}(\mathbf{y}, \mathbf{w}) := l(\mathbf{y}, \mathbf{w}) - \mathbb{E}_{\mathbf{y}'} l(\mathbf{y}', \mathbf{w})$. Then, under both H_0 and H_1 , for any fixed locations \mathbf{t} and \mathbf{t}' ,*

$$\begin{aligned} \text{cov}_{\mathbf{z}}[\hat{u}(\mathbf{t}), \hat{u}(\mathbf{t}')] &\xrightarrow{n \rightarrow \infty} \text{cov}_{\mathbf{z}}[\tilde{k}(\mathbf{x}, \mathbf{v}) \tilde{l}(\mathbf{y}, \mathbf{w}), \tilde{k}(\mathbf{x}, \mathbf{v}') \tilde{l}(\mathbf{y}, \mathbf{w}')] \\ &= \mathbb{E}_{\mathbf{xy}}[(\tilde{k}(\mathbf{x}, \mathbf{v}) \tilde{l}(\mathbf{y}, \mathbf{w}) - u(\mathbf{t}))(\tilde{k}(\mathbf{x}, \mathbf{v}') \tilde{l}(\mathbf{y}, \mathbf{w}') - u(\mathbf{t}'))], \end{aligned}$$

where $u(\mathbf{v}, \mathbf{w})$ is given in (2), and $\hat{u}(\mathbf{v}, \mathbf{w})$ is defined in (4). Second, if $0 < \text{cov}_{\mathbf{z}}[\hat{u}(\mathbf{t}_i), \hat{u}(\mathbf{t}_i)] < \infty$ for $i = 1, \dots, J$, then $\sqrt{n}(\hat{\mathbf{u}} - \mathbf{u}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Sigma)$ as $n \rightarrow \infty$, where $\Sigma_{ij} = \text{cov}[\hat{u}(\mathbf{t}_i), \hat{u}(\mathbf{t}_j)]$ and $\mathbf{u} := (u(\mathbf{t}_1), \dots, u(\mathbf{t}_J))^\top$.

Proof. We first note that for a fixed $\mathbf{t} = (\mathbf{v}, \mathbf{w})$, $\hat{u}(\mathbf{v}, \mathbf{w})$ is a one-sample second-order U-statistic [Serfling, 2009, Section 5.1.3] with a U-statistic kernel $h_{\mathbf{t}}$ where $h_{\mathbf{t}}(\mathbf{a}, \mathbf{b}) =$

$h_{\mathbf{t}}(\mathbf{b}, \mathbf{a})$. Thus, by Kowalski and Tu [2008, Section 5.1, Theorem 1], it follows directly that $\text{cov}[\hat{u}(\mathbf{t}), \hat{u}(\mathbf{t}')] = 4\text{cov}_{\mathbf{z}}[\mathbb{E}_{\mathbf{a}} h_{\mathbf{t}}(\mathbf{z}, \mathbf{a}), \mathbb{E}_{\mathbf{b}} h_{\mathbf{t}'}(\mathbf{z}, \mathbf{b})]$. Substituting $h_{\mathbf{t}}$ with its definition yields the first claim, where we note that $\mathbb{E}_{\mathbf{xy}}[\tilde{k}(\mathbf{x}, \mathbf{v}) \tilde{l}(\mathbf{y}, \mathbf{w})] = u(\mathbf{v}, \mathbf{w})$.

For the second claim, since $\hat{\mathbf{u}}$ is a multivariate one-sample U-statistic, by Lehmann [1999, Theorem 6.1.6] and Kowalski and Tu [2008, Section 5.1, Theorem 1], it follows that $\sqrt{n}(\hat{\mathbf{u}} - \mathbf{u}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Sigma)$ as $n \rightarrow \infty$, where $\Sigma_{ij} = \text{cov}[\hat{u}(\mathbf{t}_i), \hat{u}(\mathbf{t}_j)]$. \square

Recall from Proposition 2 that $\mathbf{u} = \mathbf{0}$ holds almost surely under H_0 . The asymptotic normality in the second claim of Proposition 4 implies that $n\widehat{\text{FSIC}}^2 = \frac{n}{J} \hat{\mathbf{u}}^\top \hat{\mathbf{u}}$ converges in distribution to a sum of J dependent weighted χ^2 random variables. The dependence comes from the fact that the coordinates $\hat{u}_1, \dots, \hat{u}_J$ of $\hat{\mathbf{u}}$ all depend on the sample Z_n . This null distribution is not analytically tractable, and requires a large number of simulations to compute the rejection threshold T_α for a given significance value α .

2.2 Normalized FSIC and Adaptive Test

For the purpose of an independence test, we will consider a normalized variant of $\widehat{\text{FSIC}}^2$, which we call $\widehat{\text{NFSIC}}^2$, whose tractable asymptotic null distribution is $\chi^2(J)$, the chi-squared distribution with J degrees of freedom. We then show that the independence test defined by $\widehat{\text{NFSIC}}^2$ is consistent. These results are given in Theorem 5.

Theorem 5 (Independence test using $\widehat{\text{NFSIC}}^2$ is consistent). *Let $\hat{\Sigma}$ be a consistent estimate of Σ based on the joint sample Z_n . The $\widehat{\text{NFSIC}}^2$ statistic is defined as $\hat{\lambda}_n := n\hat{\mathbf{u}}^\top (\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}}$ where $\gamma_n \geq 0$ is a regularization parameter. Assume that*

1. Assumption A holds.
2. Σ is invertible almost surely with respect to $V_J = \{(\mathbf{v}_i, \mathbf{w}_i)\}_{i=1}^J$ drawn from an absolutely continuous distribution.
3. $\lim_{n \rightarrow \infty} \gamma_n = 0$.

Then, for any k, l and V_J satisfying the assumptions,

1. Under H_0 , $\hat{\lambda}_n \xrightarrow{d} \chi^2(J)$ as $n \rightarrow \infty$.
2. Under H_1 , for any $r \in \mathbb{R}$, $\lim_{n \rightarrow \infty} \mathbb{P}(\hat{\lambda}_n \geq r) = 1$ almost surely. That is, the independence test based on $\widehat{\text{NFSIC}}^2$ is consistent.

Proof (sketch). Under H_0 , $n\hat{\mathbf{u}}^\top (\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}}$ asymptotically follows $\chi^2(J)$ because $\sqrt{n}\hat{\mathbf{u}}$ is asymptotically normally distributed (see Proposition 4). Claim 2 builds on the result in Proposition 2 stating that $\mathbf{u} \neq \mathbf{0}$ under H_1 ;

it follows using the convergence of $\hat{\mathbf{u}}$ to \mathbf{u} . The full proof can be found in Appendix E. \square

Theorem 5 states that if H_1 holds, the statistic can be arbitrarily large as n increases, allowing H_0 to be rejected for any fixed threshold. Asymptotically the test threshold T_α is given by the $(1-\alpha)$ -quantile of $\chi^2(J)$ and is independent of n . The assumption on the consistency of $\hat{\Sigma}$ is required to obtain the asymptotic chi-squared distribution. The regularization parameter γ_n is to ensure that $(\hat{\Sigma} + \gamma_n \mathbf{I})^{-1}$ can be stably computed. In practice, γ_n requires no tuning, and can be set to be a very small constant.

The next proposition states that the computational complexity of the $\widehat{\text{NFSIC}}^2$ estimator is linear in both the input dimension and sample size, and that it can be expressed in terms of the $\mathbf{K} = [K_{ij}] = [k(\mathbf{v}_i, \mathbf{x}_j)] \in \mathbb{R}^{J \times n}$, $\mathbf{L} = [L_{ij}] = [l(\mathbf{w}_i, \mathbf{y}_j)] \in \mathbb{R}^{J \times n}$ matrices.

Proposition 6 (An empirical estimator of $\widehat{\text{NFSIC}}^2$). *Let $\mathbf{1}_n := (1, \dots, 1)^\top \in \mathbb{R}^n$. Denote by \circ the element-wise matrix product. Then,*

1. $\hat{\mathbf{u}} = \frac{(\mathbf{K} \circ \mathbf{L}) \mathbf{1}_n}{n-1} - \frac{(\mathbf{K} \mathbf{1}_n) \circ (\mathbf{L} \mathbf{1}_n)}{n(n-1)}$.

2. A consistent estimator for Σ is $\hat{\Sigma} = \frac{\Gamma \Gamma^\top}{n}$ where

$$\begin{aligned} \Gamma &:= (\mathbf{K} - n^{-1} \mathbf{K} \mathbf{1}_n \mathbf{1}_n^\top) \circ (\mathbf{L} - n^{-1} \mathbf{L} \mathbf{1}_n \mathbf{1}_n^\top) - \hat{\mathbf{u}}^b \mathbf{1}_n^\top, \\ \hat{\mathbf{u}}^b &:= n^{-1} (\mathbf{K} \circ \mathbf{L}) \mathbf{1}_n - n^{-2} (\mathbf{K} \mathbf{1}_n) \circ (\mathbf{L} \mathbf{1}_n). \end{aligned}$$

Assume that the complexity of the kernel evaluation is linear in the input dimension. Then the test statistic $\hat{\lambda}_n = n \hat{\mathbf{u}}^\top (\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}}$ can be computed in $\mathcal{O}(J^3 + J^2 n + (d_x + d_y) J n)$ time.

Proof (sketch). Claim 1 for $\hat{\mathbf{u}}$ is straightforward. The expression for $\hat{\Sigma}$ in claim 2 follows directly from the asymptotic covariance expression in Proposition 4. The consistency of $\hat{\Sigma}$ can be obtained by noting that the finite sample bound for $\mathbb{P}(\|\hat{\Sigma} - \Sigma\|_F > t)$ decreases as n increases. This is implicitly shown in Appendix F.2.2 and its following sections. \square

Although the dependency of the estimator on J is cubic, we empirically observe that only a small value of J is required (see Section 3). The number of test locations J relates to the number of regions in $\mathcal{X} \times \mathcal{Y}$ of p_{xy} and $p_x p_y$ that differ (see Figure 1). In particular, J need not increase with n for test consistency.

Our final theoretical result gives a lower bound on the test power of $\widehat{\text{NFSIC}}^2$ i.e., the probability of correctly rejecting H_0 . We will use this lower bound as the objective function to determine V_J and the kernel parameters. Let $\|\cdot\|_F$ be the Frobenius norm.

Theorem 7 (A lower bound on the test power). *Let $\widehat{\text{NFSIC}}^2(X, Y) := \lambda_n := n \mathbf{u}^\top \Sigma^{-1} \mathbf{u}$. Let \mathcal{K} be a kernel class for k , \mathcal{L} be a kernel class for l , and \mathcal{V} be a collection with each element being a set of J locations. Assume that*

1. *There exist finite B_k and B_l such that*

$$\sup_{k \in \mathcal{K}} \sup_{\mathbf{x}, \mathbf{x}' \in \mathcal{X}} |k(\mathbf{x}, \mathbf{x}')| \leq B_k \quad \text{and} \quad \sup_{l \in \mathcal{L}} \sup_{\mathbf{y}, \mathbf{y}' \in \mathcal{Y}} |l(\mathbf{y}, \mathbf{y}')| \leq B_l.$$

2. $\tilde{c} := \sup_{k \in \mathcal{K}} \sup_{l \in \mathcal{L}} \sup_{V_J \in \mathcal{V}} \|\Sigma^{-1}\|_F < \infty$.

Then, for any $k \in \mathcal{K}, l \in \mathcal{L}, V_J \in \mathcal{V}$, and $\lambda_n \geq r$, the test power satisfies $\mathbb{P}(\hat{\lambda}_n \geq r) \geq L(\lambda_n)$ where

$$\begin{aligned} L(\lambda_n) &= 1 - 62e^{-\xi_1 \gamma_n^2 (\lambda_n - r)^2 / n} - 2e^{-[0.5n] (\lambda_n - r)^2 / [\xi_2 n^2]} \\ &\quad - 2e^{-[(\lambda_n - r) \gamma_n (n-1)] / 3 - \xi_3 n - c_3 \gamma_n^2 n(n-1)]^2 / [\xi_4 n^2 (n-1)]}, \end{aligned}$$

$[\cdot]$ is the floor function, $\xi_1 := \frac{1}{32c_1^2 J^2 B^*}$, $\xi_2 := 72c_2^2 J B^2$, $B := B_k B_l$, $\xi_3 := 8c_1 B^2 J$, $c_3 := 4B^2 J \tilde{c}^2$, $\xi_4 := 2^8 B^4 J^2 c_1^2$, $c_1 := 4B^2 J \sqrt{J} \tilde{c}$, $c_2 := 4B \sqrt{J} \tilde{c}$, and B^* is a constant depending on only B_k and B_l . Moreover, for sufficiently large fixed n , $L(\lambda_n)$ is increasing in λ_n .

We provide the proof in Appendix F. To put Theorem 7 into perspective, let θ_x and θ_y be the parameters of the kernels k and l , respectively. We denote by $\theta = \{\theta_x, \theta_y, V_J\}$ the collection of all tuning parameters of the test. Assume that $\mathcal{K} = \left\{ (\mathbf{x}, \mathbf{v}) \mapsto \exp\left(-\frac{\|\mathbf{x}-\mathbf{v}\|^2}{2\sigma_x^2}\right) \mid \sigma_x^2 \in [\sigma_{x,l}^2, \sigma_{x,u}^2] \right\} =: \mathcal{K}_g$ for some $0 < \sigma_{x,l}^2 < \sigma_{x,u}^2 < \infty$ and $\mathcal{L} = \left\{ (\mathbf{y}, \mathbf{w}) \mapsto \exp\left(-\frac{\|\mathbf{y}-\mathbf{w}\|^2}{2\sigma_y^2}\right) \mid \sigma_y^2 \in [\sigma_{y,l}^2, \sigma_{y,u}^2] \right\} =: \mathcal{L}_g$ for some $0 < \sigma_{y,l}^2 < \sigma_{y,u}^2 < \infty$ are Gaussian kernel classes. Then, in Theorem 7, $B = B_k = B_l = 1$, and $B^* = 2$. The assumption $\tilde{c} < \infty$ is a technical condition to guarantee that the test power lower bound is finite for all θ defined by the feasible sets \mathcal{K}, \mathcal{L} , and \mathcal{V} . Let $\mathcal{V}_{\epsilon,r} := \{V_J \mid \|\mathbf{v}_i\|^2, \|\mathbf{w}_i\|^2 \leq r \text{ and } \|\mathbf{v}_i - \mathbf{v}_j\|_2^2 + \|\mathbf{w}_i - \mathbf{w}_j\|_2^2 \geq \epsilon, \text{ for all } i \neq j\}$. If we set $\mathcal{K} = \mathcal{K}_g, \mathcal{L} = \mathcal{L}_g$, and $\mathcal{V} = \mathcal{V}_{\epsilon,r}$ for some $\epsilon, r > 0$, then $\tilde{c} < \infty$ as $\mathcal{K}_g, \mathcal{L}_g$, and $\mathcal{V}_{\epsilon,r}$ are compact. In practice, these conditions do not necessarily create restrictions as they almost always hold implicitly. We show in Appendix C that the objective function used to choose V_J will discourage any two locations to be in the same neighborhood.

Parameter Tuning The test power lower bound $L(\lambda_n)$ in Theorem 7 is a function of $\lambda_n = n \mathbf{u}^\top \Sigma^{-1} \mathbf{u}$ which is the population counterpart of the test statistic $\hat{\lambda}_n$. As in FSIC, it can be shown that $\lambda_n = 0$ if and only if X and Y are independent (from Proposition 2). If X and Y are dependent, then $\lambda_n > 0$. According to Theorem 7, for a sufficiently large n , the test power lower bound is increasing in λ_n . One can therefore think of λ_n (a function of θ) as representing how easily the test rejects H_0 given a problem P_{xy} . The higher the λ_n , the

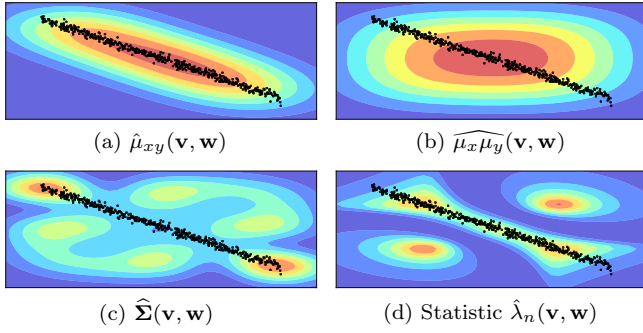


Figure 1: Illustration of $\widehat{\text{NFSIC}}^2$.

greater the lower bound on the test power, and thus the more likely it is that the test will reject H_0 when it is false.

In light of this reasoning, we propose setting θ to $\theta^* = \arg \max_{\theta} \lambda_n$. That this procedure is also valid under H_0 can be seen as follows. Under H_0 , $\theta^* = \arg \max_{\theta} 0$ will be arbitrary. Since Theorem 7 guarantees that $\hat{\lambda}_n \xrightarrow{d} \chi^2(J)$ as $n \rightarrow \infty$ for any θ , the asymptotic null distribution does not change by using θ^* . In practice, λ_n is a population quantity which is unknown. We propose dividing the sample Z_n into two disjoint sets: training and test sets. The training set is used to optimize for θ^* , and the test set is used for the actual independence test with the optimized θ^* . The splitting is to guarantee the independence of θ^* and the test sample, which is an assumption of Theorem 5.

To better understand $\widehat{\text{NFSIC}}^2$, we visualize $\hat{\mu}_{xy}(\mathbf{v}, \mathbf{w})$, $\widehat{\mu_x \mu_y}(\mathbf{v}, \mathbf{w})$ and $\widehat{\Sigma}(\mathbf{v}, \mathbf{w})$ as a function of one test location (\mathbf{v}, \mathbf{w}) on a simple toy problem. In this problem, $Y = -X + Z$ where $Z \sim \mathcal{N}(0, 0.3^2)$. As we consider only one location ($J = 1$), $\widehat{\Sigma}(\mathbf{v}, \mathbf{w})$ is a scalar. The statistic can be written as $\hat{\lambda}_n = n \frac{(\hat{\mu}_{xy}(\mathbf{v}, \mathbf{w}) - \widehat{\mu_x \mu_y}(\mathbf{v}, \mathbf{w}))^2}{\widehat{\Sigma}(\mathbf{v}, \mathbf{w})}$. These components are shown in Figure 1, where we use Gaussian kernels for both X and Y , and the horizontal and vertical axes correspond to $\mathbf{v} \in \mathbb{R}$ and $\mathbf{w} \in \mathbb{R}$, respectively.

Intuitively, $\hat{u}(\mathbf{v}, \mathbf{w}) = \hat{\mu}_{xy}(\mathbf{v}, \mathbf{w}) - \widehat{\mu_x \mu_y}(\mathbf{v}, \mathbf{w})$ captures the difference of the joint distribution and the product of the marginals as a function of (\mathbf{v}, \mathbf{w}) . Squaring $\hat{u}(\mathbf{v}, \mathbf{w})$ and dividing it by the variance shown in Figure 1c gives the statistic (also the parameter tuning objective) shown in Figure 1d. The latter figure suggests that the parameter tuning objective function can be non-convex. However, we note that the non-convexity arises since there are multiple ways to detect the difference between the joint distribution and the product of the marginals. In this case, the lower left and upper right regions equally indicate the largest difference.

3 Experiments

In this section, we empirically study the performance of the proposed method on both toy (Section 3.1) and real-life problems (Section 3.2). Our interest is in the performance of linear-time tests on challenging problems which require a large sample size to be able to accurately reveal the dependence. All the code is available at <https://github.com/wittawatj/fsic-test>.

We compare the proposed NFSIC with optimization (NFSIC-opt) to five multivariate nonparametric tests. The $\widehat{\text{NFSIC}}^2$ test without optimization (NFSIC-med) acts as a baseline, allowing the effect of parameter optimization to be clearly seen. For pedagogical reason, we consider the original HSIC test of Gretton et al. [2005] denoted by QHSIC, which is a quadratic-time test. Nyström HSIC (NyHSIC) uses a Nyström approximation to the kernel matrices of X and Y when computing the HSIC statistic. FHSIC is another variant of HSIC in which a random Fourier feature approximation [Rahimi and Recht, 2008] to the kernel is used. NyHSIC and FHSIC are studied in Zhang et al. [2016] and can be computed in $\mathcal{O}(n)$, with quadratic dependency on the number of inducing points in NyHSIC, and quadratic dependency in the number of random features in FHSIC. Finally, the Randomized Dependence Coefficient (RDC) proposed in Lopez-Paz et al. [2013] is also considered. The RDC can be seen as the primal form (with random Fourier features) of the kernel canonical correlation analysis of Bach and Jordan [2002] on copula-transformed data. We consider RDC as a linear-time test even though preprocessing by an empirical copula transform costs $\mathcal{O}((d_x + d_y)n \log n)$.

We use Gaussian kernel classes \mathcal{K}_g and \mathcal{L}_g for both X and Y in all the methods. Except NFSIC-opt, all other tests use full sample to conduct the independence test, where the Gaussian widths σ_x and σ_y are set according to the widely used median heuristic i.e., $\sigma_x = \text{median}(\{\|\mathbf{x}_i - \mathbf{x}_j\|_2 \mid 1 \leq i < j \leq n\})$, and σ_y is set in the same way using $\{\mathbf{y}_i\}_{i=1}^n$. The J locations for NFSIC-med are randomly drawn from the standard multivariate normal distribution in each trial. For a sample of size n , NFSIC-opt uses half the sample for parameter tuning, and the other disjoint half for the test. We permute the sample 300 times in RDC¹ and HSIC to simulate from the null distribution and compute the test threshold. The null distributions for FHSIC and NyHSIC are given by a finite sum of weighted $\chi^2(1)$ random variables given in Eq. 8 of Zhang et al. [2016]. Unless stated otherwise, we set the test threshold of the two NFSIC tests to be the $(1 - \alpha)$ -quantile of $\chi^2(J)$. To provide a fair comparison, we set $J = 10$, use 10 inducing points in NyHSIC, and 10

¹We use a permutation test for RDC, following the authors' implementation (<https://github.com/lopezpaz/randomized-dependence-coefficient>, referred commit: b0ac6c0).

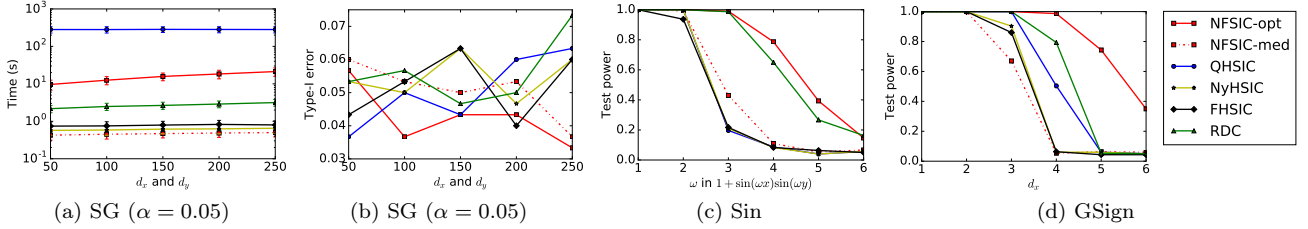


Figure 2: (a): Runtime. (b): Probability of rejecting H_0 as problem parameters vary. Fix $n = 4000$.

random Fourier features in FHSIC and RDC.

Optimization of NFSIC-opt The parameters of NFSIC-opt are σ_x, σ_y , and J locations of size $(d_x + d_y)J$. We treat all the parameters as a long vector in $\mathbb{R}^{2+(d_x+d_y)J}$ and use gradient ascent to optimize $\hat{\lambda}_{n/2}$. We observe that initializing V_J by randomly picking J points from the training sample yields good performance. The regularization parameter γ_n in NFSIC is fixed to a small value, and is not optimized. It is worth emphasizing that the complexity of the optimization procedure is still linear in n .

Since FSIC, NyHFSIC and RDC rely on a finite-dimensional kernel approximation, these tests are consistent only if both the number of features increases with n . By contrast, the proposed NFSIC requires only n to go to infinity to achieve consistency i.e., J can be fixed. We refer the reader to Appendix C for a brief investigation of the test power vs. increasing J . The test power does not necessarily monotonically increase with J .

3.1 Toy Problems

We consider three toy problems: Same Gaussian (SG), Sinusoid (Sin), and Gaussian Sign (GSign).

1. Same Gaussian (SG). The two variables are independently drawn from the standard multivariate normal distribution i.e., $X \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{d_x})$ and $Y \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{d_y})$ where \mathbf{I}_d is the $d \times d$ identity matrix. This problem represents a case in which H_0 holds.

2. Sinusoid (Sin). Let p_{xy} be the probability density of P_{xy} . In the Sinusoid problem, the dependency of X and Y is characterized by $(X, Y) \sim p_{xy}(x, y) \propto 1 + \sin(\omega x)\sin(\omega y)$, where the domains of $\mathcal{X}, \mathcal{Y} = (-\pi, \pi)$ and ω is the frequency of the sinusoid. As the frequency ω increases, the drawn sample becomes more similar to a sample drawn from $\text{Uniform}((-\pi, \pi)^2)$. That is, the higher ω , the harder to detect the dependency between X and Y . This problem was studied in [Sejdinovic et al. \[2013\]](#). Plots of the density for a few values of ω are shown in Figures 6 and 7 in the appendix. The main characteristic of interest in this problem is the local change in the density function.

3. Gaussian Sign (GSign). In this problem, $Y =$

$|Z| \prod_{i=1}^{d_x} \text{sgn}(X_i)$, where $X \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{d_x})$, $\text{sgn}(\cdot)$ is the sign function, and $Z \sim \mathcal{N}(0, 1)$ serves as a source of noise. The full interaction of $X = (X_1, \dots, X_{d_x})$ is what makes the problem challenging. That is, Y is dependent on X , yet it is independent of any proper subset of $\{X_1, \dots, X_d\}$. Thus, simultaneous consideration of all the coordinates of X is required to successfully detect the dependency.

We fix $n = 4000$ and vary the problem parameters. Each problem is repeated for 300 trials, and the sample is redrawn each time. The significance level α is set to 0.05. The results are shown in Figure 2. It can be seen that in the SG problem (Figure 2b) where H_0 holds, all the tests achieve roughly correct type-I errors at $\alpha = 0.05$. In the Sin problem, NFSIC-opt achieves the highest test power for all considered $\omega = 1, \dots, 6$, highlighting its strength in detecting local changes in the joint density. The performance of NFSIC-med is significantly lower than that of NFSIC-opt. This phenomenon clearly emphasizes the importance of the optimization to place the locations at the relevant regions in $\mathcal{X} \times \mathcal{Y}$. RDC has a remarkably high performance in both Sin and GSign (Figure 2c, 2d) despite no parameter tuning. Interestingly, both NFSIC-opt and RDC outperform the quadratic-time QHSIC in these two problems. The ability to simultaneously consider interacting features of NFSIC-opt is indicated by its superior test power in GSign, especially at the challenging settings of $d_x = 5, 6$. An average trial runtime for each test in the SG problem is shown in Figure 2a. We observe that the runtime does not increase with dimension, as the complexity of all the tests is linear in the dimension of the input. All the tests are implemented in Python using a common SciPy Stack.

To investigate the sample efficiency of all the tests, we fix $d_x = d_y = 250$ in SG, $\omega = 4$ in Sin, $d_x = 4$ in GSign, and increase n . Figure 3 shows the results. The quadratic dependency on n in QHSIC makes it infeasible both in terms of memory and runtime to consider n larger than 6000 (Figure 3a). In contrast, although not the most time-efficient, NFSIC-opt has the highest sample-efficiency for GSign, and for Sin in the low-sample regime, significantly outperforming QHSIC. Despite the small additional overhead from the optimization, we are yet able to conduct an accurate test with $n = 10^5, d_x = d_y =$

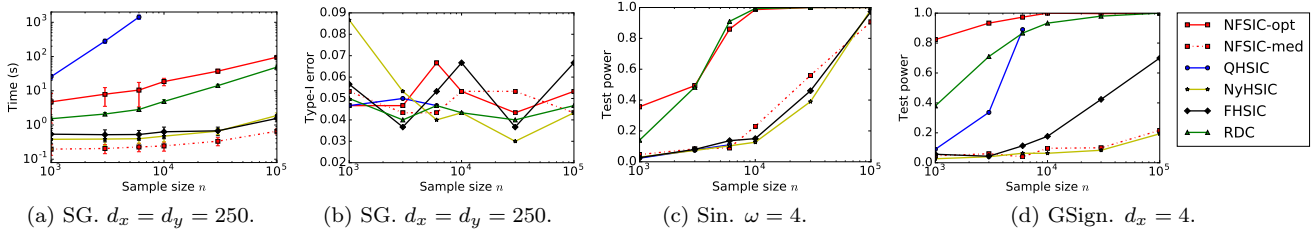


Figure 3: (a) Runtime. (b): Probability of rejecting H_0 as n increases in the toy problems.

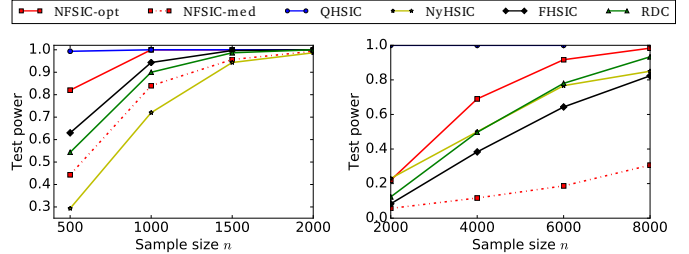
250 in less than 100 seconds. We observe in Figure 3b that the two NFSIC variants have correct type-I errors across all sample sizes, indicating that the asymptotic null distribution approximately holds by the time n reaches 1000. We recall from Theorem 5 that the NFSIC test with random test locations will asymptotically reject H_0 if it is false. A demonstration of this property is given in Figure 3c, where the test power of NFSIC-med eventually reaches 1 with n higher than 10^5 .

3.2 Real Problems

We now examine the performance of our proposed test on real problems.

Million Song Data (MSD) We consider a subset of the Million Song Data² [Bertin-Mahieux et al., 2011], in which each song (X) out of 515,345 is represented by 90 features, of which 12 features are timbre average (over all segments) of the song, and 78 features are timbre covariance. Most of the songs are western commercial tracks from 1922 to 2011. The goal is to detect the dependency between each song and its year of release (Y). We set $\alpha = 0.01$, and repeat for 300 trials where the full sample is randomly subsampled to n points in each trial. Other settings are the same as in the toy problems. To make sure that the type-I error is correct, we use the permutation approach in the NFSIC tests to compute the threshold. Figure 4b shows the test powers as n increases from 500 to 2000. To simulate the case where H_0 holds in the problem, we permute the sample to break the dependency of X and Y . The results are shown in Figure 5 in the appendix.

Evidently, NFSIC-opt has the highest test power among all the linear-time tests for all the sample sizes. Its test power is second to only QHSIC. We recall that NFSIC-opt uses half of the sample for parameter tuning. Thus, at $n = 500$, the actual sample for testing is 250, which is relatively small. The fact that there is a vast power gain from 0.4 (NFSIC-med) to 0.8 (NFSIC-opt) at $n = 500$ suggests that the optimization procedure can perform well even at a lower sample sizes.



(a) MSD problem. (b) Videos & Captions problem.

Figure 4: Probability of rejecting H_0 as n increases in the two real problems. $\alpha = 0.01$.

Videos and Captions Our last problem is based on the VideoStory46K³ dataset [Habibian et al., 2014]. The dataset contains 45,826 Youtube videos (X) of an average length of roughly one minute, and their corresponding text captions (Y) uploaded by the users. Each video is represented as a $d_x = 2000$ dimensional Fisher vector encoding of motion boundary histograms (MBH) descriptors of Wang and Schmid [2013]. Each caption is represented as a bag of words with each feature being the frequency of one word. After filtering only words which occur in at least six video captions, we obtain $d_y = 1878$ words. We examine the test powers as n increases from 2000 to 8000. The results are given in Figure 4. The problem is sufficiently challenging that all linear-time tests achieve a low power at $n = 2000$. QHSIC performs exceptionally well on this problem, achieving a maximum power throughout. NFSIC-opt has the highest sample efficiency among the linear-time tests, showing that the optimization procedure is also practical in a high dimensional setting.

Acknowledgement

We thank the Gatsby Charitable Foundation for the financial support. The major part of this work was carried out while Zoltán Szabó was a research associate at the Gatsby Computational Neuroscience Unit, University College London.

²Million Song Data subset: <https://archive.ics.uci.edu/ml/datasets/YearPredictionMSD>.

³VideoStory46K dataset: <https://ivi.fnwi.uva.nl/isis/mediamill/datasets/videostory.php>.

References

- T. W. Anderson. *An Introduction to Multivariate Statistical Analysis*. Wiley, 2003.
- F. R. Bach and M. I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.
- T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere. The million song dataset. In *International Conference on Music Information Retrieval (ISMIR)*, 2011.
- K. P. Chwialkowski, A. Ramdas, D. Sejdinovic, and A. Gretton. Fast Two-Sample Testing with Analytic Representations of Probability Measures. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1981–1989. 2015.
- J. Dauxois and G. M. Nkiet. Nonlinear canonical analysis and independence tests. *The Annals of Statistics*, 26(4):1254–1278, 1998.
- A. Feuerverger. A consistent test for bivariate dependence. *International Statistical Review*, 61(3):419–433, 1993.
- K. Fukumizu, A. Gretton, X. Sun, and B. Schölkopf. Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems (NIPS)*, pages 489–496, 2008.
- A. Gretton. A simpler condition for consistency of a kernel independence test. Technical report, 2015. URL <http://arxiv.org/abs/1501.06103>.
- A. Gretton and L. Györfi. Consistent nonparametric tests of independence. *Journal of Machine Learning Research*, 11:1391–1423, 2010.
- A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf. Measuring Statistical Dependence with Hilbert-Schmidt Norms. In *Algorithmic Learning Theory (ALT)*, pages 63–77. 2005.
- A. Gretton, K. Fukumizu, C. H. Teo, L. Song, B. Schölkopf, and A. J. Smola. A Kernel Statistical Test of Independence. In *Advances in Neural Information Processing Systems (NIPS)*, pages 585–592. 2008.
- A. Habibian, T. Mensink, and C. G. Snoek. Videostory: A new multimedia embedding for few-example recognition and translation of events. In *ACM International Conference on Multimedia*, pages 17–26, 2014.
- R. Heller, Y. Heller, S. Kaufman, B. Brill, and M. Gorfine. Consistent distribution-free k -sample and independence tests for univariate random variables. *Journal of Machine Learning Research*, 17(29):1–54, 2016.
- X. Huo and G. J. Székely. Fast computing for distance covariance. Technical report, 2014. URL <https://arxiv.org/abs/1410.1503>.
- W. Jitkrittum, Z. Szabó, K. Chwialkowski, and A. Gretton. Interpretable Distribution Features with Maximum Testing Power. 2016. URL <http://arxiv.org/abs/1605.06796>.
- J. Kowalski and X. M. Tu. *Modern Applied U-Statistics*. John Wiley & Sons, 2008.
- E. L. Lehmann. *Elements of Large-Sample Theory*. Springer Science & Business Media, 1999.
- D. Lopez-Paz, P. Hennig, and B. Schölkopf. The Randomized Dependence Coefficient. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1–9. 2013.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1177–1184. 2008.
- D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The Annals of Statistics*, 41(5):2263–2291, 2013.
- R. J. Serfling. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, 2009.
- A. Smola, A. Gretton, L. Song, and B. Schölkopf. A hilbert space embedding for distributions. In *International Conference on Algorithmic Learning Theory (ALT)*, pages 13–31, 2007.
- B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. G. Lanckriet. Hilbert Space Embeddings and Metrics on Probability Measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.
- I. Steinwart and A. Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- G. J. Székely and M. L. Rizzo. Brownian distance covariance. *The Annals of Applied Statistics*, 3(4):1236–1265, 2009.
- G. J. Székely, M. L. Rizzo, and N. K. Bakirov. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6):2769–2794, 2007.
- A. W. v. d. Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000.
- H. Wang and C. Schmid. Action recognition with improved trajectories. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3551–3558, 2013.

- K. Zhang, J. Peters, D. Janzing, B., and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 804–813, 2011.
- Q. Zhang, S. Filippi, A. Gretton, and D. Sejdinovic. Large-Scale Kernel Methods for Independence Testing. 2016. URL <http://arxiv.org/abs/1606.07892>.

An Adaptive Test of Independence with Analytic Kernel Embeddings

Supplementary Material

A Type-I Errors

In this section, we show that all the tests have correct type-I errors (i.e., the probability of reject H_0 when it is true) in real problems. We permute the joint sample so that the dependency is broken to simulate cases in which H_0 holds. The results are shown in Figure 5.

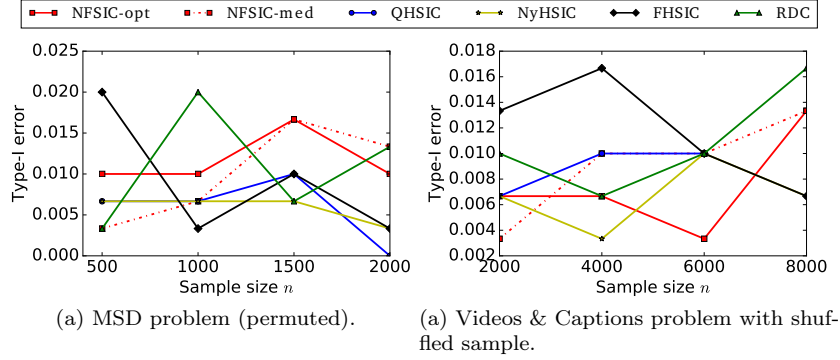


Figure 5: Probability of rejecting H_0 as n increases in the Million Song problem. $\alpha = 0.01$.

B Redundant Test Locations

Here, we provide a simple illustration to show that two locations $\mathbf{t}_1 = (\mathbf{v}_1, \mathbf{w}_1)$ and $\mathbf{t}_2 = (\mathbf{v}_2, \mathbf{w}_2)$ which are too close to each other will reduce the optimization objective. We consider the Sinusoid problem described in Section 3.1 with $\omega = 1$, and use $J = 2$ test locations. In Figure 6, \mathbf{t}_1 is fixed at the red star, while \mathbf{t}_2 is varied along the horizontal line. The objective value $\hat{\lambda}_n$ as a function of $(\mathbf{t}_1, \mathbf{t}_2)$ is shown in the bottom figure. It can be seen that $\hat{\lambda}_n$ decreases sharply when \mathbf{t}_2 is in the neighborhood of \mathbf{t}_1 . This property implies that two locations which are too close will not maximize the objective function (i.e., the second feature contains no additional information when it matches the first). For $J > 2$, the objective sharply decreases if any two locations are in the same neighborhood.

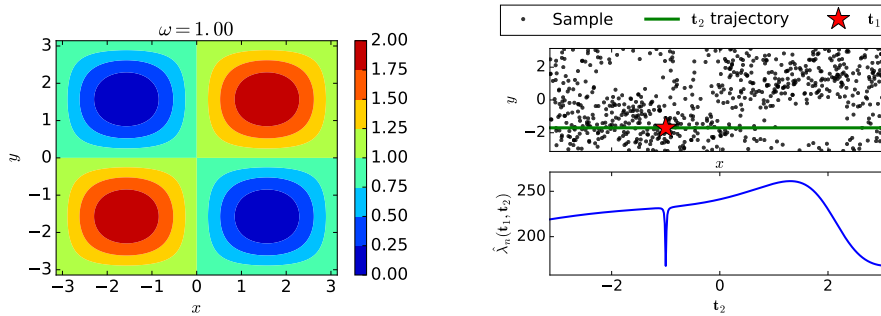


Figure 6: Plot of optimization objective values as location \mathbf{t}_2 moves along the green line. The objective sharply drops when the two locations are in the same neighborhood.

C Test Power vs. J

It might seem intuitive that as the number of locations J increases, the test power should also increase. Here, we empirically show that this statement is *not* always true. Consider the Sinusoid toy example described in Section 3.1 with $\omega = 2$ (also see the left figure of Figure 7). By construction, X and Y are dependent in this problem. We run

NFSIC test with a sample size of $n = 800$, varying J from 1 to 600. For each value of J , the test is repeated for 500 times. In each trial, the sample is redrawn and the J test locations are drawn from $\text{Uniform}((-\pi, \pi)^2)$. There is no optimization of the test locations. We use Gaussian kernels for both X and Y , and use the median heuristic to set the Gaussian widths to 1.8. Figure 7 shows the test power as J increases.

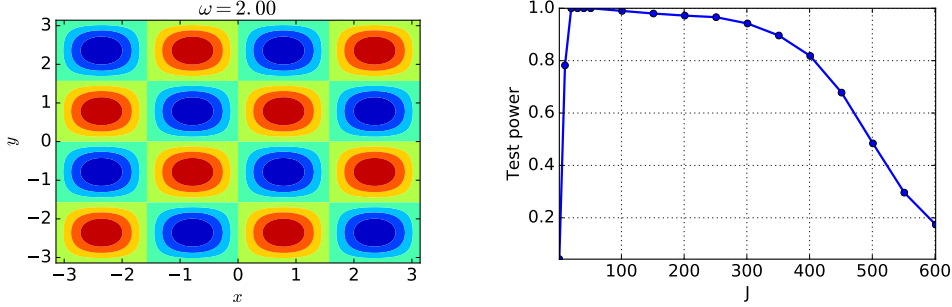


Figure 7: The Sinusoid problem and the plot of test power vs. the number of test locations.

We observe that the test power does not monotonically increase as J increases. When $J = 1$, the difference of p_{xy} and $p_x p_y$ cannot be adequately captured, resulting in a low power. The power increases rapidly to roughly 0.8 at $J = 10$, and stays at the maximum until about $J = 100$. Then, the power starts to drop sharply when J is higher than 400 in this problem.

Unlike random Fourier features, the number of test locations in NFSIC is not the number of Monte Carlo particles used to approximate an expectation. There is a tradeoff: if the test locations are in key regions (i.e., regions in which there is a big difference between p_{xy} and $p_x p_y$), then they increase power; yet the statistic gains in variance (thus reducing test power) as J increases. As can be seen in Figure 7, there are eight key regions (in blue) that can reveal the difference of p_{xy} and $p_x p_y$. Using an unnecessarily high J not only makes the covariance matrix $\hat{\Sigma}$ harder to estimate accurately, it also increases the computation as the complexity on J is $\mathcal{O}(J^3)$.

We note that NFSIC is not intended to be used with a large J . In practice, it should be set to be large enough so as to capture the key regions as stated. As a practical guide, with optimization of the test locations, a good starting point is $J = 5$ or 10.

D Proof of Proposition 3

Recall Proposition 3,

Proposition (A product of Gaussian kernels is characteristic and analytic). *Let $k(\mathbf{x}, \mathbf{x}') = \exp(-(\mathbf{x} - \mathbf{x}')^\top \mathbf{A}(\mathbf{x} - \mathbf{x}'))$ and $l(\mathbf{y}, \mathbf{y}') = \exp(-(\mathbf{y} - \mathbf{y}')^\top \mathbf{B}(\mathbf{y} - \mathbf{y}'))$ be Gaussian kernels on $\mathbb{R}^{d_x} \times \mathbb{R}^{d_x}$ and $\mathbb{R}^{d_y} \times \mathbb{R}^{d_y}$ respectively, for positive definite matrices \mathbf{A} and \mathbf{B} . Then, $g((\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}')) = k(\mathbf{x}, \mathbf{x}')l(\mathbf{y}, \mathbf{y}')$ is characteristic and analytic on $(\mathbb{R}^{d_x} \times \mathbb{R}^{d_y}) \times (\mathbb{R}^{d_x} \times \mathbb{R}^{d_y})$.*

Proof. Let $\mathbf{z} := (\mathbf{x}^\top, \mathbf{y}^\top)^\top$ and $\mathbf{z}' := (\mathbf{x}'^\top, \mathbf{y}'^\top)^\top$ be vectors in $\mathbb{R}^{d_x+d_y}$. We prove by reducing the product kernel to one Gaussian kernel with $g(\mathbf{z}, \mathbf{z}') = \exp(-(\mathbf{z} - \mathbf{z}')^\top \mathbf{C}(\mathbf{z} - \mathbf{z}'))$ where $\mathbf{C} := \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \end{pmatrix}$. Write $g(\mathbf{z}, \mathbf{z}') = \Psi(\mathbf{z} - \mathbf{z}')$ where $\Psi(\mathbf{t}) := \exp(-\mathbf{t}^\top \mathbf{C} \mathbf{t})$. Since \mathbf{C} is positive definite, we see that the finite measure ζ corresponding to Ψ as defined in Lemma 12 has support everywhere in $\mathbb{R}^{d_x+d_y}$. Thus, Sriperumbudur et al. [2010, Theorem 9] implies that g is characteristic.

To see that g is analytic, we observe that for each $\mathbf{z}' \in \mathbb{R}^{d_x+d_y}$, $\mathbf{z} \mapsto -(\mathbf{z} - \mathbf{z}')^\top \mathbf{C}(\mathbf{z} - \mathbf{z}')$ is a multivariate polynomial in \mathbf{z} , which is known to be analytic. Using the fact that $t \mapsto \exp(t)$ is analytic on \mathbb{R} , and that a composition of analytic functions is analytic, we see that $\mathbf{z} \mapsto \exp(-(\mathbf{z} - \mathbf{z}')^\top \mathbf{C}(\mathbf{z} - \mathbf{z}'))$ is analytic on $\mathbb{R}^{d_x+d_y}$ for each \mathbf{z}' . \square

E Proof of Theorem 5

Recall Theorem 5,

Theorem 5 (Independence test using $\widehat{\text{NFSIC}}^2$ is consistent). Let $\hat{\Sigma}$ be a consistent estimate of Σ based on the joint sample Z_n . The $\widehat{\text{NFSIC}}^2$ statistic is defined as $\hat{\lambda}_n := n\hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}}$ where $\gamma_n \geq 0$ is a regularization parameter. Assume that

1. Assumption A holds.
2. Σ is invertible almost surely with respect to $V_J = \{(\mathbf{v}_i, \mathbf{w}_i)\}_{i=1}^J$ drawn from an absolutely continuous distribution.
3. $\lim_{n \rightarrow \infty} \gamma_n = 0$.

Then, for any k, l and V_J satisfying the assumptions,

1. Under H_0 , $\hat{\lambda}_n \xrightarrow{d} \chi^2(J)$ as $n \rightarrow \infty$.
2. Under H_1 , for any $r \in \mathbb{R}$, $\lim_{n \rightarrow \infty} \mathbb{P}(\hat{\lambda}_n \geq r) = 1$ almost surely. That is, the independence test based on $\widehat{\text{NFSIC}}^2$ is consistent.

Proof. Assume that H_0 holds. The consistency of $\hat{\Sigma}$ and the continuous mapping theorem imply that $\left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \xrightarrow{p} \Sigma^{-1}$ which is a constant. Let \mathbf{a} be a random vector in \mathbb{R}^J following $\mathcal{N}(\mathbf{0}, \Sigma)$. By Vaart [2000, Theorem 2.7 (v)], it follows that $\left[\sqrt{n}\hat{\mathbf{u}}, \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \right] \xrightarrow{d} [\mathbf{a}, \Sigma^{-1}]$ where $\mathbf{u} = \mathbf{0}$ almost surely by Proposition 2, and $\sqrt{n}\hat{\mathbf{u}} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Sigma)$ by Proposition 4. Since $f(\mathbf{x}, \mathbf{S}) := \mathbf{x}^\top \mathbf{S} \mathbf{x}$ is continuous, $f\left(\sqrt{n}\hat{\mathbf{u}}, \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1}\right) \xrightarrow{d} f(\mathbf{a}, \Sigma^{-1})$. Equivalently, $n\hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} \xrightarrow{d} \mathbf{a}^\top \Sigma^{-1} \mathbf{a} \sim \chi^2(J)$ by Anderson [2003, Theorem 3.3.3]. This proves the first claim.

The proof of the second claim has a very similar structure to the proof of Proposition 2 of Chwialkowski et al. [2015]. Assume that H_1 holds. Then, $\mathbf{u} \neq \mathbf{0}$ almost surely by Proposition 2. Since k and l are bounded, it follows that $|h_t(\mathbf{z}, \mathbf{z}')| \leq 2B_k B_l$ for any \mathbf{z}, \mathbf{z}' (see (8)), and we have that $\hat{\mathbf{u}} \xrightarrow{a.s.} \mathbf{u}$ by Serfling [2009, Section 5.4, Theorem A]. Thus, $\hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} - \frac{r}{n} \xrightarrow{d} \mathbf{u}^\top \Sigma^{-1} \mathbf{u}$ by the continuous mapping theorem, and the consistency of $\hat{\Sigma}$. Consequently,

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{P}(\hat{\lambda}_n \geq r) \\ &= 1 - \lim_{n \rightarrow \infty} \mathbb{P}\left(\hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} - \frac{r}{n} < 0\right) \\ &\stackrel{(a)}{=} 1 - \mathbb{P}(\mathbf{u}^\top \Sigma^{-1} \mathbf{u} < 0) \stackrel{(b)}{=} 1, \end{aligned}$$

where at (a) we use the Portmanteau theorem [Vaart, 2000, Lemma 2.2 (i)] guaranteeing that $x_n \xrightarrow{d} x$ if and only if $\mathbb{P}(x_n < t) \rightarrow \mathbb{P}(x < t)$ for all continuity points of $t \mapsto \mathbb{P}(x < t)$. Step (b) is justified by noting that the covariance matrix Σ is positive definite so that $\mathbf{u}^\top \Sigma^{-1} \mathbf{u} > 0$, and $t \mapsto \mathbb{P}(\mathbf{u}^\top \Sigma^{-1} \mathbf{u} < t)$ (a step function) is continuous at 0. \square

F Proof of Theorem 7

Recall Theorem 7,

Theorem 7 (A lower bound on the test power). Let $\text{NFSIC}^2(X, Y) := \lambda_n := n\mathbf{u}^\top \Sigma^{-1} \mathbf{u}$. Let \mathcal{K} be a kernel class for k , \mathcal{L} be a kernel class for l , and \mathcal{V} be a collection with each element being a set of J locations. Assume that

1. There exist finite B_k and B_l such that $\sup_{k \in \mathcal{K}} \sup_{\mathbf{x}, \mathbf{x}' \in \mathcal{X}} |k(\mathbf{x}, \mathbf{x}')| \leq B_k$ and $\sup_{l \in \mathcal{L}} \sup_{\mathbf{y}, \mathbf{y}' \in \mathcal{Y}} |l(\mathbf{y}, \mathbf{y}')| \leq B_l$.
2. $\tilde{c} := \sup_{k \in \mathcal{K}} \sup_{l \in \mathcal{L}} \sup_{V_J \in \mathcal{V}} \|\Sigma^{-1}\|_F < \infty$.

Then, for any $k \in \mathcal{K}, l \in \mathcal{L}, V_J \in \mathcal{V}$, and $\lambda_n \geq r$, the test power satisfies $\mathbb{P}(\hat{\lambda}_n \geq r) \geq L(\lambda_n)$ where

$$\begin{aligned} L(\lambda_n) &= 1 - 62e^{-\xi_1 \gamma_n^2 (\lambda_n - r)^2 / n} - 2e^{-[0.5n](\lambda_n - r)^2 / [\xi_2 n^2]} \\ &\quad - 2e^{-[(\lambda_n - r)\gamma_n(n-1)/3 - \xi_3 n - c_3 \gamma_n^2 n(n-1)]^2 / [\xi_4 n^2(n-1)]}, \end{aligned}$$

$\lfloor \cdot \rfloor$ is the floor function, $\xi_1 := \frac{1}{3^2 c_1^2 J^2 B^*}$, $\xi_2 := 72 c_2^2 J B^2$, $B := B_k B_l$, $\xi_3 := 8 c_1 B^2 J$, $c_3 := 4 B^2 J \tilde{c}$, $\xi_4 := 2^8 B^4 J^2 c_1^2$, $c_1 := 4 B^2 J \sqrt{J} \tilde{c}$, $c_2 := 4 B \sqrt{J} \tilde{c}$, and B^* is a constant depending on only B_k and B_l . Moreover, for sufficiently large fixed n , $L(\lambda_n)$ is increasing in λ_n .

Overview of the proof We first derive a probabilistic bound for $|\hat{\lambda}_n - \lambda_n|/n$. The bound is in turn upper bounded by an expression involving $\|\hat{\mathbf{u}} - \mathbf{u}\|_2$ and $\|\hat{\Sigma} - \Sigma\|_F$. The difference $\|\hat{\mathbf{u}} - \mathbf{u}\|_2$ can be bounded by applying the bound for U-statistics given in Serfling [2009, Theorem A, p. 201]. For $\|\hat{\Sigma} - \Sigma\|_F$, we decompose it into a sum of smaller components, and bound each term with a product variant of the Hoeffding’s inequality (Lemma 9). $L(\lambda_n)$ is obtained by combining all the bounds with the union bound.

F.1 Notations

Let $\langle \mathbf{A}, \mathbf{B} \rangle_F := \text{tr}(\mathbf{A}^\top \mathbf{B})$ denote the Frobenius inner product, and $\|\mathbf{A}\|_F := \sqrt{\text{tr}(\mathbf{A}^\top \mathbf{A})}$ be the Frobenius norm. Write $\mathbf{z} := (\mathbf{x}, \mathbf{y})$ to denote a pair of points from $\mathcal{X} \times \mathcal{Y}$. We write $\mathbf{t} := (\mathbf{v}, \mathbf{w})$ to denote a pair of test locations from $\mathcal{X} \times \mathcal{Y}$. For brevity, an expectation over (\mathbf{x}, \mathbf{y}) (i.e., $\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P_{\mathbf{x}\mathbf{y}}}$) will be written as $\mathbb{E}_{\mathbf{x}}$ or $\mathbb{E}_{\mathbf{x}\mathbf{y}}$. Define $\tilde{k}(\mathbf{x}, \mathbf{v}) := k(\mathbf{x}, \mathbf{v}) - \mathbb{E}_{\mathbf{x}'} k(\mathbf{x}', \mathbf{v})$, and $\tilde{l}(\mathbf{y}, \mathbf{w}) := l(\mathbf{y}, \mathbf{w}) - \mathbb{E}_{\mathbf{y}'} l(\mathbf{y}', \mathbf{w})$. Let $B_2(r) := \{\mathbf{x} \mid \|\mathbf{x}\|_2 \leq r\}$ be a closed ball with radius r centered at the origin. Similarly, define $B_F(r) := \{\mathbf{A} \mid \|\mathbf{A}\|_F \leq r\}$ to be a closed ball with radius r of $J \times J$ matrices under the Frobenius norm. Denote the max operation by $(x_1, \dots, x_m)_+ = \max(x_1, \dots, x_m)$.

For a product of marginal mean embeddings $\mu_x(\mathbf{v})\mu_y(\mathbf{w})$, we write $\widehat{\mu_x \mu_y}(\mathbf{v}, \mathbf{w}) := \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n k(\mathbf{x}_i, \mathbf{v}) l(\mathbf{y}_j, \mathbf{w})$ to denote the unbiased plug-in estimator, and write $\hat{\mu}_x(\mathbf{v})\hat{\mu}_y(\mathbf{w}) := \frac{1}{n} \sum_{i=1}^n k(\mathbf{x}_i, \mathbf{v}) \frac{1}{n} \sum_{j=1}^n l(\mathbf{y}_j, \mathbf{w})$ which is a biased estimator. Define $\hat{u}^b(\mathbf{v}, \mathbf{w}) := \hat{\mu}_{xy}(\mathbf{v}, \mathbf{w}) - \hat{\mu}_x(\mathbf{v})\hat{\mu}_y(\mathbf{w})$ so that $\hat{\mathbf{u}}^b := (\hat{u}^b(\mathbf{t}_1), \dots, \hat{u}^b(\mathbf{t}_J))^\top$ where the superscript b stands for “biased”. To avoid confusing with a positive definite kernel, we will refer to a U-statistic kernel as a *core*.

F.2 Proof

We will first derive a bound for $\mathbb{P}(|\hat{\lambda}_n - \lambda_n| \geq t)$, which will then be reparametrized to get a bound for the target quantity $\mathbb{P}(\hat{\lambda}_n \geq r)$. We closely follow the proof in Jitkrittum et al. [2016, Section C.1] up to (12), then we diverge. We start by considering $|\hat{\lambda}_n - \lambda_n|/n$.

$$\begin{aligned} |\hat{\lambda}_n - \lambda_n|/n &= \left| \hat{\mathbf{u}}^\top (\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top \Sigma^{-1} \mathbf{u} \right| \\ &= \left| \hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} + \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} - \mathbf{u}^\top \Sigma^{-1} \mathbf{u} \right| \\ &\leq \left| \hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} \right| + \left| \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} - \mathbf{u}^\top \Sigma^{-1} \mathbf{u} \right| \\ &:= (\star)_1 + (\star)_2. \end{aligned}$$

We next bound $(\star)_1$ and $(\star)_2$ separately.

$$\begin{aligned} (\star)_1 &= \left| \hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} \right| \\ &= \left| \hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} - \hat{\mathbf{u}}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}} + \hat{\mathbf{u}}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} \right| \\ &\leq \left| \hat{\mathbf{u}}^\top \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} \hat{\mathbf{u}} - \hat{\mathbf{u}}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}} \right| + \left| \hat{\mathbf{u}}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} \right| \\ &= \left| \left\langle \hat{\mathbf{u}} \hat{\mathbf{u}}^\top, \left(\hat{\Sigma} + \gamma_n \mathbf{I} \right)^{-1} - (\Sigma + \gamma_n \mathbf{I})^{-1} \right\rangle_F \right| + \left| \left\langle \hat{\mathbf{u}} \hat{\mathbf{u}}^\top - \mathbf{u} \mathbf{u}^\top, (\Sigma + \gamma_n \mathbf{I})^{-1} \right\rangle_F \right| \\ &\leq \|\hat{\mathbf{u}} \hat{\mathbf{u}}^\top\|_F \|(\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} - (\Sigma + \gamma_n \mathbf{I})^{-1}\|_F + \|\hat{\mathbf{u}} \hat{\mathbf{u}}^\top - \mathbf{u} \mathbf{u}^\top\|_F \|(\Sigma + \gamma_n \mathbf{I})^{-1}\|_F \\ &= \|\hat{\mathbf{u}} \hat{\mathbf{u}}^\top\|_F \|(\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} [(\Sigma + \gamma_n \mathbf{I}) - (\hat{\Sigma} + \gamma_n \mathbf{I})] (\Sigma + \gamma_n \mathbf{I})^{-1}\|_F + \|\hat{\mathbf{u}} \hat{\mathbf{u}}^\top - \mathbf{u} \mathbf{u}^\top + \hat{\mathbf{u}} \hat{\mathbf{u}}^\top - \mathbf{u} \mathbf{u}^\top\|_F \|(\Sigma + \gamma_n \mathbf{I})^{-1}\|_F \end{aligned}$$

$$\begin{aligned}
&\stackrel{(a)}{\leq} \|\hat{\mathbf{u}}\hat{\mathbf{u}}^\top\|_F \|(\hat{\Sigma} + \gamma_n \mathbf{I})^{-1}\|_F \|\Sigma - \hat{\Sigma}\|_F \|\Sigma^{-1}\|_F + \|\hat{\mathbf{u}}\hat{\mathbf{u}}^\top - \hat{\mathbf{u}}\mathbf{u}^\top + \hat{\mathbf{u}}\mathbf{u}^\top - \mathbf{u}\mathbf{u}^\top\|_F \|\Sigma^{-1}\|_F \\
&\stackrel{(b)}{\leq} \frac{\sqrt{J}}{\gamma_n} \|\hat{\mathbf{u}}\|_2^2 \|\Sigma - \hat{\Sigma}\|_F \|\Sigma^{-1}\|_F + (\|\hat{\mathbf{u}}(\hat{\mathbf{u}} - \mathbf{u})^\top\|_F + \|(\hat{\mathbf{u}} - \mathbf{u})\mathbf{u}^\top\|_F) \|\Sigma^{-1}\|_F \\
&\leq \frac{\sqrt{J}}{\gamma_n} \|\hat{\mathbf{u}}\|_2^2 \|\Sigma - \hat{\Sigma}\|_F \|\Sigma^{-1}\|_F + (\|\hat{\mathbf{u}}\|_2 + \|\mathbf{u}\|_2) \|\hat{\mathbf{u}} - \mathbf{u}\|_2 \|\Sigma^{-1}\|_F,
\end{aligned} \tag{5}$$

where at (a) we used $\|(\Sigma + \gamma_n \mathbf{I})^{-1}\|_F \leq \|\Sigma^{-1}\|_F$, at (b) we used $\|(\hat{\Sigma} + \gamma_n \mathbf{I})^{-1}\|_F \leq \sqrt{J} \|(\hat{\Sigma} + \gamma_n \mathbf{I})^{-1}\|_2 \leq \sqrt{J}/\gamma_n$. For $(\star)_2$, we have

$$\begin{aligned}
(\star)_2 &= \left| \mathbf{u}^\top (\Sigma + \gamma_n \mathbf{I})^{-1} \mathbf{u} - \mathbf{u}^\top \Sigma^{-1} \mathbf{u} \right| \\
&= \left| \langle \mathbf{u}\mathbf{u}^\top, (\Sigma + \gamma_n \mathbf{I})^{-1} - \Sigma^{-1} \rangle_F \right| \\
&\leq \|\mathbf{u}\mathbf{u}^\top\|_F \|(\Sigma + \gamma_n \mathbf{I})^{-1} - \Sigma^{-1}\|_F \\
&= \|\mathbf{u}\|_2^2 \|(\Sigma + \gamma_n \mathbf{I})^{-1} [\Sigma - (\Sigma + \gamma_n \mathbf{I})] \Sigma^{-1}\|_F \\
&\leq \gamma_n \|\mathbf{u}\|_2^2 \|(\Sigma + \gamma_n \mathbf{I})^{-1}\|_F \|\Sigma^{-1}\|_F \\
&\stackrel{(a)}{\leq} \gamma_n \|\mathbf{u}\|_2^2 \|\Sigma^{-1}\|_F^2,
\end{aligned} \tag{6}$$

where at (a) we used $\|(\Sigma + \gamma_n \mathbf{I})^{-1}\|_F \leq \|\Sigma^{-1}\|_F$.

Combining (5) and (6), we have

$$\begin{aligned}
&\left| \hat{\mathbf{u}}^\top (\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top \Sigma^{-1} \mathbf{u} \right| \\
&\leq \frac{\sqrt{J}}{\gamma_n} \|\hat{\mathbf{u}}\|_2^2 \|\Sigma - \hat{\Sigma}\|_F \|\Sigma^{-1}\|_F + (\|\hat{\mathbf{u}}\|_2 + \|\mathbf{u}\|_2) \|\hat{\mathbf{u}} - \mathbf{u}\|_2 \|\Sigma^{-1}\|_F + \gamma_n \|\mathbf{u}\|_2^2 \|\Sigma^{-1}\|_F^2.
\end{aligned} \tag{7}$$

Bounding $\|\hat{\mathbf{u}}\|_2^2$ and $\|\mathbf{u}\|_2^2$ Here, we show that by the boundedness of the kernels k and l , it follows that $\|\hat{\mathbf{u}}\|_2^2$ is bounded. Recall that $\sup_{\mathbf{x}, \mathbf{x}' \in \mathcal{X}} |k(\mathbf{x}, \mathbf{x}')| \leq B_k$, $\sup_{\mathbf{y}, \mathbf{y}' \in \mathcal{Y}} |l(\mathbf{y}, \mathbf{y}')| \leq B_l$, our notation $\mathbf{t} = (\mathbf{v}, \mathbf{w})$ for the test locations, and $\mathbf{z}_i := (\mathbf{x}_i, \mathbf{y}_i)$. We first show that the U-statistic core h is bounded.

$$\begin{aligned}
|h_{\mathbf{t}}((\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}'))| &= \left| \frac{1}{2} (k(\mathbf{x}, \mathbf{v}) - k(\mathbf{x}', \mathbf{v})) (l(\mathbf{y}, \mathbf{w}) - l(\mathbf{y}', \mathbf{w})) \right| \\
&\leq \frac{1}{2} (|k(\mathbf{x}, \mathbf{v})| + |k(\mathbf{x}', \mathbf{v})|) (|l(\mathbf{y}, \mathbf{w})| + |l(\mathbf{y}', \mathbf{w})|) \\
&\leq 2B_k B_l := 2B,
\end{aligned} \tag{8}$$

where we define $B := B_k B_l$. It follows that

$$\|\hat{\mathbf{u}}\|_2^2 = \sum_{m=1}^J \left[\frac{2}{n(n-1)} \sum_{i < j} h_{\mathbf{t}_m}(\mathbf{z}_i, \mathbf{z}_j) \right]^2 \leq \sum_{m=1}^J [2B_k B_l]^2 = 4B^2 J, \tag{9}$$

$$\|\mathbf{u}\|_2^2 = \sum_{m=1}^J [\mathbb{E}_{\mathbf{z}} \mathbb{E}_{\mathbf{z}'} h_{\mathbf{t}_m}(\mathbf{z}, \mathbf{z}')]^2 \leq 4B^2 J. \tag{10}$$

Using the upper bounds on $\|\hat{\mathbf{u}}\|_2^2$, $\|\mathbf{u}\|_2^2$, (7) and the definition of \tilde{c} , we have

$$\begin{aligned}
&\left| \hat{\mathbf{u}}^\top (\hat{\Sigma} + \gamma_n \mathbf{I})^{-1} \hat{\mathbf{u}} - \mathbf{u}^\top \Sigma^{-1} \mathbf{u} \right| \\
&\leq \frac{\sqrt{J}}{\gamma_n} 4B^2 J \tilde{c} \|\Sigma - \hat{\Sigma}\|_F + 4B \sqrt{J} \tilde{c} \|\hat{\mathbf{u}} - \mathbf{u}\|_2 + 4B^2 J \tilde{c}^2 \gamma_n \\
&=: \frac{c_1}{\gamma_n} \|\Sigma - \hat{\Sigma}\|_F + c_2 \|\hat{\mathbf{u}} - \mathbf{u}\|_2 + c_3 \gamma_n,
\end{aligned} \tag{11}$$

where we define $c_1 := 4B^2 J\sqrt{J}\tilde{c}$, $c_2 := 4B\sqrt{J}\tilde{c}$, and $c_3 := 4B^2 J\tilde{c}^2$. This upper bound implies that

$$|\hat{\lambda}_n - \lambda_n| \leq \frac{c_1}{\gamma_n} n \|\Sigma - \hat{\Sigma}\|_F + c_2 n \|\hat{\mathbf{u}} - \mathbf{u}\|_2 + c_3 n \gamma_n. \quad (12)$$

We will separately upper bound $\|\Sigma - \hat{\Sigma}\|_F$ and $\|\hat{\mathbf{u}} - \mathbf{u}\|_2$, and combine them with a union bound.

F.2.1 Bounding $\|\hat{\mathbf{u}} - \mathbf{u}\|_2$

Let $\mathbf{t}^* = \arg \max_{\mathbf{t} \in \{\mathbf{t}_1, \dots, \mathbf{t}_J\}} |\hat{u}(\mathbf{t}) - u(\mathbf{t})|$. Recall that $\mathbf{u} = (u(\mathbf{t}_1), \dots, u(\mathbf{t}_J))^\top = (u_1, \dots, u_J)^\top$.

$$\begin{aligned} \|\hat{\mathbf{u}} - \mathbf{u}\|_2 &= \sup_{\mathbf{b} \in B_2(1)} \langle \mathbf{b}, \hat{\mathbf{u}} - \mathbf{u} \rangle_2 \leq \sup_{\mathbf{b} \in B_2(1)} \sum_{j=1}^J |b_j| |\hat{u}(\mathbf{t}_j) - u(\mathbf{t}_j)| \\ &\leq |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)| \sup_{\mathbf{b} \in B_2(1)} \sum_{j=1}^J |b_j| \\ &\stackrel{(a)}{\leq} \sqrt{J} |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)| \sup_{\mathbf{b} \in B_2(1)} \|\mathbf{b}\|_2 \\ &= \sqrt{J} |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)|, \end{aligned} \quad (13)$$

where at (a) we used $\|\mathbf{a}\|_1 \leq \sqrt{J} \|\mathbf{a}\|_2$ for any $\mathbf{a} \in \mathbb{R}^J$. From (13), it can be seen that bounding $\|\hat{\mathbf{u}} - \mathbf{u}\|_2$ amounts to bounding the difference of a U-statistic $\hat{u}(\mathbf{t}^*)$ (see (4)) to its expectation $u(\mathbf{t}^*)$. Combining (13) and (12), we have

$$|\hat{\lambda}_n - \lambda_n| \leq \frac{c_1}{\gamma_n} n \|\Sigma - \hat{\Sigma}\|_F + c_2 n \sqrt{J} |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)| + c_3 n \gamma_n. \quad (14)$$

F.2.2 Bounding $\|\hat{\Sigma} - \Sigma\|_F$

The plan is to write $\hat{\Sigma} = \hat{\mathbf{S}} - \hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top}$, $\Sigma = \mathbf{S} - \mathbf{u} \mathbf{u}^\top$, so that $\|\hat{\Sigma} - \Sigma\|_F \leq \|\hat{\mathbf{S}} - \mathbf{S}\|_F + \|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F$ and bound separately $\|\hat{\mathbf{S}} - \mathbf{S}\|_F$ and $\|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F$.

Recall that $\Sigma_{ij} = \eta(\mathbf{t}_i, \mathbf{t}_j)$, $\eta(\mathbf{t}, \mathbf{t}') = \mathbb{E}_{\mathbf{x}\mathbf{y}}[(\tilde{k}(\mathbf{x}, \mathbf{v})\tilde{l}(\mathbf{y}, \mathbf{w}) - u(\mathbf{v}, \mathbf{w}))(\tilde{k}(\mathbf{x}, \mathbf{v}')\tilde{l}(\mathbf{y}, \mathbf{w}') - u(\mathbf{v}', \mathbf{w}'))]$ where $\tilde{k}(\mathbf{x}, \mathbf{v}) = k(\mathbf{x}, \mathbf{v}) - \mathbb{E}_{\mathbf{x}'} k(\mathbf{x}', \mathbf{v})$, and $\tilde{l}(\mathbf{y}, \mathbf{w}) = l(\mathbf{y}, \mathbf{w}) - \mathbb{E}_{\mathbf{y}'} l(\mathbf{y}', \mathbf{w})$. Its empirical estimator (see Proposition 6) is $\hat{\Sigma}_{ij} = \hat{\eta}(\mathbf{t}_i, \mathbf{t}_j)$ where

$$\begin{aligned} \hat{\eta}(\mathbf{t}, \mathbf{t}') &= \frac{1}{n} \sum_{i=1}^n [(\bar{k}(\mathbf{x}_i, \mathbf{v})\bar{l}(\mathbf{y}_i, \mathbf{w}) - \hat{u}^b(\mathbf{v}, \mathbf{w}))(\bar{k}(\mathbf{x}_i, \mathbf{v}')\bar{l}(\mathbf{y}_i, \mathbf{w}') - \hat{u}^b(\mathbf{v}', \mathbf{w}'))] \\ &= \frac{1}{n} \sum_{i=1}^n \bar{k}(\mathbf{x}_i, \mathbf{v})\bar{l}(\mathbf{y}_i, \mathbf{w})\bar{k}(\mathbf{x}_i, \mathbf{v}')\bar{l}(\mathbf{y}_i, \mathbf{w}') - \hat{u}^b(\mathbf{v}, \mathbf{w})\hat{u}^b(\mathbf{v}', \mathbf{w}'), \end{aligned}$$

$\bar{k}(\mathbf{x}, \mathbf{v}) := k(\mathbf{x}, \mathbf{v}) - \frac{1}{n} \sum_{i=1}^n k(\mathbf{x}_i, \mathbf{v})$, and $\bar{l}(\mathbf{y}, \mathbf{w}) := l(\mathbf{y}, \mathbf{w}) - \frac{1}{n} \sum_{i=1}^n l(\mathbf{y}_i, \mathbf{w})$. We note that $\frac{1}{n} \sum_{i=1}^n \bar{k}(\mathbf{x}_i, \mathbf{v})\bar{l}(\mathbf{y}_i, \mathbf{w}) = \hat{u}^b(\mathbf{v}, \mathbf{w})$. We define $\hat{\mathbf{S}} \in \mathbb{R}^{J \times J}$ such that $\hat{S}_{ij} := \frac{1}{n} \sum_{m=1}^n \bar{k}(\mathbf{x}_m, \mathbf{v}_i)\bar{l}(\mathbf{y}_m, \mathbf{w}_i)\bar{k}(\mathbf{x}_m, \mathbf{v}_j)\bar{l}(\mathbf{y}_m, \mathbf{w}_j)$, and define similarly its population counterpart \mathbf{S} such that $S_{ij} := \mathbb{E}_{\mathbf{x}\mathbf{y}}[\tilde{k}(\mathbf{x}, \mathbf{v})\tilde{l}(\mathbf{y}, \mathbf{w})\tilde{k}(\mathbf{x}, \mathbf{v}')\tilde{l}(\mathbf{y}, \mathbf{w}')]$. We have

$$\hat{\Sigma} = \hat{\mathbf{S}} - \hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top},$$

$$\Sigma = \mathbf{S} - \mathbf{u} \mathbf{u}^\top,$$

$$\|\hat{\Sigma} - \Sigma\|_F = \|\hat{\mathbf{S}} - \mathbf{S} - (\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top)\|_F \quad (15)$$

$$\leq \|\hat{\mathbf{S}} - \mathbf{S}\|_F + \|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F. \quad (16)$$

With (16), (14) becomes

$$|\hat{\lambda}_n - \lambda_n| \leq \frac{c_1 n}{\gamma_n} \|\hat{\mathbf{S}} - \mathbf{S}\|_F + \frac{c_1 n}{\gamma_n} \|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F + c_2 n \sqrt{J} |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)| + c_3 n \gamma_n. \quad (17)$$

We will further separately bound $\|\hat{\mathbf{S}} - \mathbf{S}\|_F$ and $\|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F$.

F.2.3 Bounding $\|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F$

$$\begin{aligned}
\|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F &= \|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \hat{\mathbf{u}}^b \mathbf{u}^\top + \hat{\mathbf{u}}^b \mathbf{u}^\top - \mathbf{u} \mathbf{u}^\top\|_F \\
&\leq \|\hat{\mathbf{u}}^b (\hat{\mathbf{u}}^b - \mathbf{u})^\top\|_F + \|(\hat{\mathbf{u}}^b - \mathbf{u}) \mathbf{u}^\top\|_F \\
&= \|\hat{\mathbf{u}}^b\|_2 \|\hat{\mathbf{u}}^b - \mathbf{u}\|_2 + \|\hat{\mathbf{u}}^b - \mathbf{u}\|_2 \|\mathbf{u}\|_2 \\
&\leq 4B\sqrt{J} \|\hat{\mathbf{u}}^b - \mathbf{u}\|_2,
\end{aligned}$$

where we used (10) and the fact that $\|\hat{\mathbf{u}}^b\|_2 \leq 2B\sqrt{J}$ which can be shown similarly to (9) as

$$\|\hat{\mathbf{u}}^b\|_2^2 = \sum_{m=1}^J [\hat{\mu}_{xy}(\mathbf{v}_m, \mathbf{w}_m) - \hat{\mu}_x(\mathbf{v}_m) \hat{\mu}_y(\mathbf{w}_m)]^2 = \sum_{m=1}^J \left[\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n h_{\mathbf{t}_m}(\mathbf{z}_i, \mathbf{z}_j) \right]^2 \leq \sum_{m=1}^J [2B_k B_l]^2 = 4B^2 J.$$

Let $(\tilde{\mathbf{v}}, \tilde{\mathbf{w}}) := \tilde{\mathbf{t}} = \arg \max_{\mathbf{t} \in \{\mathbf{t}_1, \dots, \mathbf{t}_J\}} |\hat{u}^b(\mathbf{t}) - u(\mathbf{t})|$. We bound $\|\hat{\mathbf{u}}^b - \mathbf{u}\|_2$ by

$$\begin{aligned}
\|\hat{\mathbf{u}}^b - \mathbf{u}\|_2 &\stackrel{(a)}{\leq} \sqrt{J} |\hat{u}^b(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})| \\
&= \sqrt{J} |\hat{\mu}_{xy}(\tilde{\mathbf{t}}) - \hat{\mu}_x(\tilde{\mathbf{v}}) \hat{\mu}_y(\tilde{\mathbf{w}}) - u(\tilde{\mathbf{t}})| \\
&= \sqrt{J} |\hat{\mu}_{xy}(\tilde{\mathbf{t}}) - \widehat{\mu_x \mu_y}(\tilde{\mathbf{t}}) + \widehat{\mu_x \mu_y}(\tilde{\mathbf{t}}) - \hat{\mu}_x(\tilde{\mathbf{v}}) \hat{\mu}_y(\tilde{\mathbf{w}}) - u(\tilde{\mathbf{t}})| \\
&\leq \sqrt{J} |\hat{\mu}_{xy}(\tilde{\mathbf{t}}) - \widehat{\mu_x \mu_y}(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})| + \sqrt{J} |\widehat{\mu_x \mu_y}(\tilde{\mathbf{t}}) - \hat{\mu}_x(\tilde{\mathbf{v}}) \hat{\mu}_y(\tilde{\mathbf{w}})| \\
&= \sqrt{J} |\hat{u}(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})| + \sqrt{J} |\widehat{\mu_x \mu_y}(\tilde{\mathbf{t}}) - \hat{\mu}_x(\tilde{\mathbf{v}}) \hat{\mu}_y(\tilde{\mathbf{w}})|, \tag{18}
\end{aligned}$$

where at (a) we used the same reasoning as in (13). The bias $|\widehat{\mu_x \mu_y}(\tilde{\mathbf{t}}) - \hat{\mu}_x(\tilde{\mathbf{v}}) \hat{\mu}_y(\tilde{\mathbf{w}})|$ in the second term can be bounded as

$$\begin{aligned}
&|\widehat{\mu_x \mu_y}(\tilde{\mathbf{t}}) - \hat{\mu}_x(\tilde{\mathbf{v}}) \hat{\mu}_y(\tilde{\mathbf{w}})| \\
&= \left| \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_j, \tilde{\mathbf{w}}) - \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_j, \tilde{\mathbf{w}}) \right| \\
&= \left| \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_j, \tilde{\mathbf{w}}) - \frac{1}{n(n-1)} \sum_{i=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_i, \tilde{\mathbf{w}}) - \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_j, \tilde{\mathbf{w}}) \right| \\
&= \left| \left(1 - \frac{n}{n-1}\right) \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_j, \tilde{\mathbf{w}}) + \frac{1}{n(n-1)} \sum_{i=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_i, \tilde{\mathbf{w}}) \right| \\
&\leq \left| \left(1 - \frac{n}{n-1}\right) \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_j, \tilde{\mathbf{w}}) \right| + \left| \frac{1}{n(n-1)} \sum_{i=1}^n k(\mathbf{x}_i, \tilde{\mathbf{v}}) l(\mathbf{y}_i, \tilde{\mathbf{w}}) \right| \\
&\leq \frac{B}{n-1} + \frac{B}{n-1} = \frac{2B}{n-1}.
\end{aligned}$$

Combining this upper bound with (18), we have

$$\|\hat{\mathbf{u}}^b \hat{\mathbf{u}}^{b\top} - \mathbf{u} \mathbf{u}^\top\|_F \leq 4BJ |\hat{u}(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})| + \frac{8B^2 J}{n-1}. \tag{19}$$

With (19), (17) becomes

$$|\hat{\lambda}_n - \lambda_n| \leq \frac{c_1 n}{\gamma_n} \|\hat{\mathbf{S}} - \mathbf{S}\|_F + \frac{4BJc_1 n}{\gamma_n} |\hat{u}(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})| + \frac{c_1 n}{\gamma_n} \frac{8B^2 J}{n-1} + c_2 n \sqrt{J} |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)| + c_3 n \gamma_n. \tag{20}$$

$$S(\mathbf{t}, \mathbf{t}') = \mathbb{E}_{\mathbf{xy}} \left[(a - \tilde{a})(b - \tilde{b})(a' - \tilde{a}')(b' - \tilde{b}') \right].$$

By expanding $S(\mathbf{t}, \mathbf{t}')$, we have

$$\begin{aligned} S(\mathbf{t}, \mathbf{t}') &= \mathbb{E}_{\mathbf{xy}} \left[+aba'b' - aba'\tilde{b}' - ab\tilde{a}'b' + ab\tilde{a}'\tilde{b}' \right. \\ &\quad - \tilde{a}\tilde{b}a'b' + \tilde{a}\tilde{b}a'\tilde{b}' + \tilde{a}\tilde{b}\tilde{a}'b' - \tilde{a}\tilde{b}\tilde{a}'\tilde{b}' \\ &\quad - \tilde{a}ba'b' + \tilde{a}ba'\tilde{b}' + \tilde{a}b\tilde{a}'b' - \tilde{a}b\tilde{a}'\tilde{b}' \\ &\quad \left. + \tilde{a}\tilde{b}a'b' - \tilde{a}\tilde{b}a'\tilde{b}' - \tilde{a}\tilde{b}\tilde{a}'b' + \tilde{a}\tilde{b}\tilde{a}'\tilde{b}' \right] \\ &= +\widetilde{aba'b'} - \widetilde{aba'\tilde{b}'} - \widetilde{ab\tilde{a}'b'} + \widetilde{ab\tilde{a}'\tilde{b}'} \\ &\quad - \widetilde{\tilde{a}\tilde{b}a'b'} + \widetilde{\tilde{a}\tilde{b}a'\tilde{b}'} + \widetilde{\tilde{a}\tilde{b}\tilde{a}'b'} - \widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} \\ &\quad - \widetilde{a'\tilde{b}\tilde{b}'\tilde{a}} + \widetilde{a'\tilde{b}\tilde{a}\tilde{b}'} + \widetilde{\tilde{a}\tilde{a}'\tilde{b}\tilde{b}'} - \widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} \\ &\quad + \widetilde{a'b'\tilde{a}\tilde{b}} - \widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} - \widetilde{\tilde{a}\tilde{b}\tilde{a}'b'} + \widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} \\ &= +\widetilde{aba'b'} - \widetilde{aba'\tilde{b}'} - \widetilde{ab\tilde{a}'b'} + \widetilde{ab\tilde{a}'\tilde{b}'} \\ &\quad - \widetilde{\tilde{a}\tilde{b}a'b'} + \widetilde{\tilde{a}\tilde{b}a'\tilde{b}'} + \widetilde{\tilde{a}\tilde{b}\tilde{a}'b'} + \widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} \\ &\quad - \widetilde{a'\tilde{b}\tilde{b}'\tilde{a}} + \widetilde{a'\tilde{b}\tilde{a}\tilde{b}'} + \widetilde{\tilde{a}\tilde{a}'\tilde{b}\tilde{b}'} - 3\widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} \end{aligned}$$

The expansion of $\hat{S}(\mathbf{t}, \mathbf{t}')$ can be done in the same way. By the triangle inequality, we have

$$\begin{aligned} \left| \hat{S}(\mathbf{t}, \mathbf{t}') - S(\mathbf{t}, \mathbf{t}') \right| &\leq \left| \overline{aba'b'} - \widetilde{aba'b'} \right| + \left| \overline{aba'\tilde{b}'} - \widetilde{aba'\tilde{b}'} \right| + \left| \overline{ab\tilde{a}'b'} - \widetilde{ab\tilde{a}'b'} \right| + \left| \overline{ab\tilde{a}'\tilde{b}'} - \widetilde{ab\tilde{a}'\tilde{b}'} \right| \\ &\quad \left| \overline{\tilde{a}\tilde{b}a'b'} - \widetilde{\tilde{a}\tilde{b}a'b'} \right| + \left| \overline{\tilde{a}\tilde{b}a'\tilde{b}'} - \widetilde{\tilde{a}\tilde{b}a'\tilde{b}'} \right| + \left| \overline{\tilde{a}\tilde{b}\tilde{a}'b'} - \widetilde{\tilde{a}\tilde{b}\tilde{a}'b'} \right| + \left| \overline{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} - \widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} \right| \\ &\quad \left| \overline{a'\tilde{b}\tilde{b}'\tilde{a}} - \widetilde{a'\tilde{b}\tilde{b}'\tilde{a}} \right| + \left| \overline{a'\tilde{b}\tilde{a}\tilde{b}'} - \widetilde{a'\tilde{b}\tilde{a}\tilde{b}'} \right| + \left| \overline{\tilde{a}\tilde{a}'\tilde{b}\tilde{b}'} - \widetilde{\tilde{a}\tilde{a}'\tilde{b}\tilde{b}'} \right| + 3 \left| \overline{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} - \widetilde{\tilde{a}\tilde{b}\tilde{a}'\tilde{b}'} \right|. \end{aligned}$$

The first term $\left| \overline{aba'b'} - \widetilde{aba'b'} \right|$ can be bounded by applying the Hoeffding's inequality. Other terms can be bounded by applying Lemma 9. Recall that we write $(x_1, \dots, x_m)_+$ for $\max(x_1, \dots, x_m)$.

Bounding $\left| \overline{aba'b'} - \widetilde{aba'b'} \right|$ (1st term). Since $-B^2 \leq aba'b' \leq B^2$, by the Hoeffding's inequality (Lemma 14), we have

$$\mathbb{P} \left(\left| \overline{aba'b'} - \widetilde{aba'b'} \right| \leq t \right) \geq 1 - 2 \exp \left(-\frac{nt^2}{2B^4} \right).$$

Bounding $\left| \overline{aba'\tilde{b}'} - \widetilde{aba'\tilde{b}'} \right|$ (2nd term). Let $f_1(\mathbf{x}, \mathbf{y}) = aba' = k(\mathbf{x}, \mathbf{v})l(\mathbf{y}, \mathbf{w})k(\mathbf{x}, \mathbf{v}')$ and $f_2(\mathbf{y}) = b' = l(\mathbf{y}, \mathbf{w}')$. We note that $|f_1(\mathbf{x}, \mathbf{y})| \leq (BB_k, B_l)_+$ and $|f_2(\mathbf{y})| \leq (BB_k, B_l)_+$. Thus, by Lemma 9 with $E = 2$, we have

$$\mathbb{P} \left(\left| \overline{aba'\tilde{b}'} - \widetilde{aba'\tilde{b}'} \right| \leq t \right) \geq 1 - 4 \exp \left(-\frac{nt^2}{8(BB_k, B_l)_+^4} \right).$$

Bounding $\left| \overline{ab\tilde{a}'b'} - \widetilde{ab\tilde{a}'b'} \right|$ (4th term). Let $f_1(\mathbf{x}, \mathbf{y}) = ab = k(\mathbf{x}, \mathbf{v})l(\mathbf{y}, \mathbf{w})$, $f_2(\mathbf{x}) = a' = k(\mathbf{x}, \mathbf{v}')$ and $f_3(\mathbf{y}) = b' = l(\mathbf{y}, \mathbf{w}')$. We can see that $|f_1(\mathbf{x}, \mathbf{y})|, |f_2(\mathbf{x})|, |f_3(\mathbf{y})| \leq (B, B_k, B_l)_+$. Thus, by Lemma 9 with $E = 3$, we have

$$\mathbb{P} \left(\left| \overline{ab\tilde{a}'b'} - \widetilde{ab\tilde{a}'b'} \right| \leq t \right) \geq 1 - 6 \exp \left(-\frac{nt^2}{18(B, B_k, B_l)_+^6} \right).$$

Bounding $|\bar{a}\bar{b}\bar{a}'\bar{b}' - \tilde{a}\tilde{b}\tilde{a}'\tilde{b}'|$ (last term). Let $f_1(\mathbf{x}) = a = k(\mathbf{x}, \mathbf{v})$, $f_2(\mathbf{y}) = b = l(\mathbf{y}, \mathbf{w})$, $f_3(\mathbf{x}) = a' = k(\mathbf{x}, \mathbf{v}')$ and $f_4(\mathbf{y}) = b' = l(\mathbf{y}, \mathbf{w}')$. It can be seen that $|f_1(\mathbf{x})|, |f_2(\mathbf{y})|, |f_3(\mathbf{x})|, |f_4(\mathbf{y})| \leq (B_k, B_l)_+$. Thus, by Lemma 9 with $E = 4$, we have

$$\mathbb{P}\left(3\left|\bar{a}\bar{b}\bar{a}'\bar{b}' - \tilde{a}\tilde{b}\tilde{a}'\tilde{b}'\right| \leq t\right) \geq 1 - 8 \exp\left(-\frac{nt^2}{32 \cdot 3^2(B_k, B_l)_+^8}\right).$$

Bounds for other terms can be derived in a similar way to yield

$$\begin{aligned} (3^{rd} \text{ term}) \quad & \mathbb{P}\left(\left|\overline{abb'a'} - \widetilde{abb'a'}\right| \leq t\right) \geq 1 - 4 \exp\left(-\frac{nt^2}{8(BB_l, B_k)_+^4}\right), \\ (5^{th} \text{ term}) \quad & \mathbb{P}\left(\left|\overline{aa'b'\bar{b}} - \widetilde{aa'b'\bar{b}}\right| \leq t\right) \geq 1 - 4 \exp\left(-\frac{nt^2}{8(BB_k, B_l)_+^4}\right), \\ (6^{th} \text{ term}) \quad & \mathbb{P}\left(\left|\overline{aa'\bar{b}\bar{b}'} - \widetilde{aa'\bar{b}\bar{b}'}\right| \leq t\right) \geq 1 - 6 \exp\left(-\frac{nt^2}{18(B_k^2, B_l)_+^6}\right), \\ (7^{th} \text{ term}) \quad & \mathbb{P}\left(\left|\overline{ab'a'\bar{b}} - \widetilde{ab'a'\bar{b}}\right| \leq t\right) \geq 1 - 6 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right), \\ (8^{th} \text{ term}) \quad & \mathbb{P}\left(\left|\overline{a'b'\bar{a}\bar{b}} - \widetilde{a'b'\bar{a}\bar{b}}\right| \leq t\right) \geq 1 - 6 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right), \\ (9^{th} \text{ term}) \quad & \mathbb{P}\left(\left|\overline{a'bb'\bar{a}} - \widetilde{a'bb'\bar{a}}\right| \leq t\right) \geq 1 - 4 \exp\left(-\frac{nt^2}{8(BB_l, B_k)_+^4}\right), \\ (10^{th} \text{ term}) \quad & \mathbb{P}\left(\left|\overline{a'\bar{b}\bar{a}\bar{b}'} - \widetilde{a'\bar{b}\bar{a}\bar{b}'}\right| \leq t\right) \geq 1 - 6 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right), \\ (11^{th} \text{ term}) \quad & \mathbb{P}\left(\left|\bar{a}\bar{a}'\bar{b}\bar{b}' - \tilde{a}\tilde{a}'\tilde{b}\tilde{b}'\right| \leq t\right) \geq 1 - 6 \exp\left(-\frac{nt^2}{18(B_k, B_l^2)_+^6}\right). \end{aligned}$$

By the union bound, we have

$$\begin{aligned} & \mathbb{P}\left(\left|\hat{S}(\mathbf{t}, \mathbf{t}') - S(\mathbf{t}, \mathbf{t}')\right| \leq 12t\right) \\ & \geq 1 - \left[2 \exp\left(-\frac{nt^2}{2B^4}\right) + 4 \exp\left(-\frac{nt^2}{8(BB_k, B_l)_+^4}\right) + 4 \exp\left(-\frac{nt^2}{8(BB_l, B_k)_+^4}\right) + 6 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right) \right. \\ & \quad 4 \exp\left(-\frac{nt^2}{8(BB_k, B_l)_+^4}\right) + 6 \exp\left(-\frac{nt^2}{18(B_k^2, B_l)_+^6}\right) + 6 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right) + 6 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right) \\ & \quad \left. 4 \exp\left(-\frac{nt^2}{8(BB_l, B_k)_+^4}\right) + 6 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right) + 6 \exp\left(-\frac{nt^2}{18(B_k, B_l^2)_+^6}\right) + 8 \exp\left(-\frac{nt^2}{32 \cdot 3^2(B_k, B_l)_+^8}\right)\right] \\ & = 1 - \left[2 \exp\left(-\frac{nt^2}{2B^4}\right) + 8 \exp\left(-\frac{nt^2}{8(BB_k, B_l)_+^4}\right) + 8 \exp\left(-\frac{nt^2}{8(BB_l, B_k)_+^4}\right) + 24 \exp\left(-\frac{nt^2}{18(B, B_k, B_l)_+^6}\right) \right. \\ & \quad \left. + 6 \exp\left(-\frac{nt^2}{18(B_k^2, B_l)_+^6}\right) + 6 \exp\left(-\frac{nt^2}{18(B_k, B_l^2)_+^6}\right) + 8 \exp\left(-\frac{nt^2}{32 \cdot 3^2(B_k, B_l)_+^8}\right)\right] \\ & \geq 1 - \left[2 \exp\left(-\frac{12^2 nt^2}{B^*}\right) + 8 \exp\left(-\frac{12^2 nt^2}{B^*}\right) + 8 \exp\left(-\frac{12^2 nt^2}{B^*}\right) + 24 \exp\left(-\frac{12^2 nt^2}{B^*}\right) \right. \\ & \quad \left. + 6 \exp\left(-\frac{12^2 nt^2}{B^*}\right) + 6 \exp\left(-\frac{12^2 nt^2}{B^*}\right) + 8 \exp\left(-\frac{12^2 nt^2}{B^*}\right)\right] \\ & = 1 - 62 \exp\left(-\frac{12^2 nt^2}{B^*}\right), \end{aligned}$$

where

$$B^* := \frac{1}{12^2} \max(2B^4, 8(BB_k, B_l)_+^4, 8(BB_l, B_k)_+^4, 18(B, B_k, B_l)_+^6, 18(B_k^2, B_l)_+^6, 18(B_k, B_l^2)_+^6, 32 \cdot 3^2(B_k, B_l)_+^8).$$

By reparameterization, it follows that

$$\mathbb{P}\left(\frac{c_1 J n}{\gamma_n} \left| \hat{S}(\mathbf{t}, \mathbf{t}') - S(\mathbf{t}, \mathbf{t}') \right| \leq t\right) \geq 1 - 62 \exp\left(-\frac{\gamma_n^2 t^2}{c_1^2 J^2 n B^*}\right). \quad (23)$$

F.2.6 Union Bound for $|\hat{\lambda}_n - \lambda_n|$ and Final Lower Bound

Recall from (22) that

$$\begin{aligned} |\hat{\lambda}_n - \lambda_n| &\leq \frac{c_1 J n}{\gamma_n} \left| \hat{S}(\mathbf{t}^{(1)}, \mathbf{t}^{(2)}) - S(\mathbf{t}^{(1)}, \mathbf{t}^{(2)}) \right| + \frac{4B J c_1 n}{\gamma_n} |\hat{u}(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})| \\ &\quad + \frac{c_1 n}{\gamma_n} \frac{8B^2 J}{n-1} + c_2 n \sqrt{J} |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)| + c_3 n \gamma_n. \end{aligned}$$

We will bound terms in (22) separately and combine all the bounds with the union bound. As shown in (8), the U-statistic core h is bounded between $-2B$ and $2B$. Thus, by Lemma 13 (with $m = 2$), we have

$$\mathbb{P}\left(c_2 n \sqrt{J} |\hat{u}(\mathbf{t}^*) - u(\mathbf{t}^*)| \leq t\right) \geq 1 - 2 \exp\left(-\frac{\lfloor 0.5n \rfloor t^2}{8c_2^2 n^2 J B^2}\right). \quad (24)$$

Bounding $\frac{c_1 n}{\gamma_n} \frac{8B^2 J}{n-1} + c_3 n \gamma_n + \frac{4B J c_1 n}{\gamma_n} |\hat{u}(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})|$. By Lemma 13 (with $m = 2$), it follows that

$$\begin{aligned} &\mathbb{P}\left(\frac{c_1 n}{\gamma_n} \frac{8B^2 J}{n-1} + c_3 n \gamma_n + \frac{4B J c_1 n}{\gamma_n} |\hat{u}(\tilde{\mathbf{t}}) - u(\tilde{\mathbf{t}})| \leq t\right) \\ &\geq 1 - 2 \exp\left(-\frac{\lfloor 0.5n \rfloor \gamma_n^2 \left[t - \frac{c_1 n}{\gamma_n} \frac{8B^2 J}{n-1} - c_3 n \gamma_n\right]^2}{2^7 B^4 J^2 c_1^2 n^2}\right) \\ &= 1 - 2 \exp\left(-\frac{\lfloor 0.5n \rfloor [t \gamma_n (n-1) - 8c_1 B^2 n J - c_3 n (n-1) \gamma_n^2]^2}{2^7 B^4 J^2 c_1^2 n^2 (n-1)^2}\right) \\ &\stackrel{(a)}{\geq} 1 - 2 \exp\left(-\frac{\lfloor t \gamma_n (n-1) - 8c_1 B^2 n J - c_3 n (n-1) \gamma_n^2 \rfloor^2}{2^8 B^4 J^2 c_1^2 n^2 (n-1)}\right), \end{aligned} \quad (25)$$

where at (a) we used $\lfloor 0.5n \rfloor \geq (n-1)/2$. Combining (23), (24), and (25) with the union bound (set $T = 3t$), we can bound (22) with

$$\begin{aligned} \mathbb{P}\left(|\hat{\lambda}_n - \lambda_n| \leq T\right) &\geq 1 - 62 \exp\left(-\frac{\gamma_n^2 T^2}{3^2 c_1^2 J^2 n B^*}\right) - 2 \exp\left(-\frac{\lfloor 0.5n \rfloor T^2}{72 c_2^2 n^2 J B^2}\right) \\ &\quad - 2 \exp\left(-\frac{\lfloor T \gamma_n (n-1) / 3 - 8c_1 B^2 n J - c_3 \gamma_n^2 n (n-1) \rfloor^2}{2^8 B^4 J^2 c_1^2 n^2 (n-1)}\right). \end{aligned}$$

Since $|\hat{\lambda}_n - \lambda_n| \leq T$ implies $\hat{\lambda}_n \geq \lambda_n - T$, a reparametrization with $r = \lambda_n - T$ gives

$$\begin{aligned} \mathbb{P}\left(\hat{\lambda}_n \geq r\right) &\geq 1 - 62 \exp\left(-\frac{\gamma_n^2 (\lambda_n - r)^2}{3^2 c_1^2 J^2 n B^*}\right) - 2 \exp\left(-\frac{\lfloor 0.5n \rfloor (\lambda_n - r)^2}{72 c_2^2 n^2 J B^2}\right) \\ &\quad - 2 \exp\left(-\frac{\lfloor (\lambda_n - r) \gamma_n (n-1) / 3 - 8c_1 B^2 n J - c_3 \gamma_n^2 n (n-1) \rfloor^2}{2^8 B^4 J^2 c_1^2 n^2 (n-1)}\right) \\ &:= L(\lambda_n). \end{aligned}$$

Grouping constants into ξ_1, \dots, ξ_5 gives the result.

The lower bound $L(\lambda_n)$ takes the form

$$1 - 62 \exp(-C_1(\lambda_n - T_\alpha)^2) - 2 \exp(-C_2(\lambda_n - T_\alpha)^2) - 2 \exp\left(-\frac{[(\lambda_n - T_\alpha)C_3 - C_4]^2}{C_5}\right),$$

where C_1, \dots, C_5 are positive constants. For fixed large enough n such that $\lambda_n > T_\alpha$, and fixed significance level α , increasing λ_n will increase $L(\lambda_n)$. Specifically, since n is fixed, increasing $\mathbf{u}^\top \boldsymbol{\Sigma}^{-1} \mathbf{u}$ in $\lambda_n = n \mathbf{u}^\top \boldsymbol{\Sigma}^{-1} \mathbf{u}$ will increase $L(\lambda_n)$.

G Helper Lemmas

This section contains lemmas used to prove the main results in this work.

Lemma 8 (Product to sum). *Assume that $|a_i| \leq B$, $|b_i| \leq B$ for $i = 1, \dots, E$. Then $\left| \prod_{i=1}^E a_i - \prod_{i=1}^E b_i \right| \leq B^{E-1} \sum_{j=1}^E |a_j - b_j|$.*

Proof.

$$\begin{aligned} \left| \prod_{i=1}^E a_i - \prod_{j=1}^E b_j \right| &\leq \left| \prod_{i=1}^E a_i - \prod_{i=1}^{E-1} a_i b_E \right| + \left| \prod_{i=1}^{E-1} a_i b_E - \prod_{i=1}^{E-2} a_i b_{E-1} b_E \right| + \dots + \left| a_1 \prod_{j=2}^E b_j - \prod_{j=1}^E b_j \right| \\ &\leq |a_E - b_E| \left| \prod_{i=1}^{E-1} a_i \right| + |a_{E-1} - b_{E-1}| \left| \left(\prod_{i=1}^{E-2} a_i \right) b_E \right| + \dots + |a_1 - b_1| \left| \prod_{j=2}^E b_j \right| \\ &\leq |a_E - b_E| B^{E-1} + |a_{E-1} - b_{E-1}| B^{E-1} + \dots + |a_1 - b_1| B^{E-1} \\ &= B^{E-1} \sum_{j=1}^E |a_j - b_j| \end{aligned}$$

applying triangle inequality, and the boundedness of a_i and b_i -s. \square

Lemma 9 (Product variant of the Hoeffding's inequality). *For $i = 1, \dots, E$, let $\{\mathbf{x}_j^{(i)}\}_{j=1}^{n_i} \subset \mathcal{X}_i$ be an i.i.d. sample from a distribution P_i , and $f_i : \mathcal{X}_i \mapsto \mathbb{R}$ be a measurable function. Note that it is possible that $P_1 = P_2 = \dots = P_E$ and $\{\mathbf{x}_j^{(1)}\}_{j=1}^{n_1} = \dots = \{\mathbf{x}_j^{(E)}\}_{j=1}^{n_E}$. Assume that $|f_i(\mathbf{x})| \leq B < \infty$ for all $\mathbf{x} \in \mathcal{X}_i$ and $i = 1, \dots, E$. Write \hat{P}_i to denote an empirical distribution based on the sample $\{\mathbf{x}_j^{(i)}\}_{j=1}^{n_i}$. Then,*

$$\mathbb{P} \left(\left| \left[\prod_{i=1}^E \mathbb{E}_{\mathbf{x}^{(i)} \sim \hat{P}_i} f_i(\mathbf{x}^{(i)}) \right] - \left[\prod_{i=1}^E \mathbb{E}_{\mathbf{x}^{(i)} \sim P_i} f_i(\mathbf{x}^{(i)}) \right] \right| \leq T \right) \geq 1 - 2 \sum_{i=1}^E \exp\left(-\frac{n_i T^2}{2E^2 B^2 E}\right).$$

Proof. By Lemma 8, we have

$$\left| \left[\prod_{i=1}^E \mathbb{E}_{\mathbf{x}^{(i)} \sim \hat{P}_i} f_i(\mathbf{x}^{(i)}) \right] - \left[\prod_{i=1}^E \mathbb{E}_{\mathbf{x}^{(i)} \sim P_i} f_i(\mathbf{x}^{(i)}) \right] \right| \leq B^{E-1} \sum_{i=1}^E \left| \mathbb{E}_{\mathbf{x}^{(i)} \sim \hat{P}_i} f_i(\mathbf{x}^{(i)}) - \mathbb{E}_{\mathbf{x}^{(i)} \sim P_i} f_i(\mathbf{x}^{(i)}) \right|.$$

By applying the Hoeffding's inequality to each term in the sum, we have $\mathbb{P} \left(\left| \mathbb{E}_{\mathbf{x}^{(i)} \sim \hat{P}_i} f_i(\mathbf{x}^{(i)}) - \mathbb{E}_{\mathbf{x}^{(i)} \sim P_i} f_i(\mathbf{x}^{(i)}) \right| \leq t \right) \geq 1 - 2 \exp\left(-\frac{2n_i t^2}{4B^2}\right)$. The result is obtained with a union bound. \square

H External Lemmas

In this section, we provide known results referred to in this work.

Lemma 10 (Chwialkowski et al. [2015, Lemma 1]). *If k is a bounded, analytic kernel (in the sense given in Definition 1) on $\mathbb{R}^d \times \mathbb{R}^d$, then all functions in the RKHS defined by k are analytic.*

Lemma 11 (Chwialkowski et al. [2015, Lemma 3]). *Let Λ be an injective mapping from the space of probability measures into a space of analytic functions on \mathbb{R}^d . Define*

$$d_{V_J}^2(P, Q) = \sum_{j=1}^J |[\Lambda P](\mathbf{v}_j) - [\Lambda Q](\mathbf{v}_j)|^2,$$

where $V_J = \{\mathbf{v}_i\}_{i=1}^J$ are vector-valued i.i.d. random variables from a distribution which is absolutely continuous with respect to the Lebesgue measure. Then, $d_{V_J}(P, Q)$ is almost surely (w.r.t. V_J) a metric.

Lemma 12 (Bochner's theorem). *A continuous function $\Psi : \mathbb{R}^d \rightarrow \mathbb{R}$ is positive definite if and only if it is the Fourier transform of a finite nonnegative Borel measure ζ on \mathbb{R}^d , that is, $\Psi(\mathbf{x}) = \int_{\mathbb{R}^d} e^{-i\mathbf{x}^\top \boldsymbol{\omega}} d\zeta(\boldsymbol{\omega})$, $\mathbf{x} \in \mathbb{R}^d$.*

Lemma 13 (A bound for U-statistics [Serfling, 2009, Theorem A, p. 201]). *Let $h(\mathbf{x}_1, \dots, \mathbf{x}_m)$ be a U-statistic kernel for an m -order U-statistic such that $h(\mathbf{x}_1, \dots, \mathbf{x}_m) \in [a, b]$ where $a \leq b < \infty$. Let $U_n = \binom{n}{m}^{-1} \sum_{i_1 < \dots < i_m} h(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_m})$ be a U-statistic computed with a sample of size n , where the summation is over the $\binom{n}{m}$ combinations of m distinct elements $\{i_1, \dots, i_m\}$ from $\{1, \dots, n\}$. Then, for $t > 0$ and $n \geq m$,*

$$\begin{aligned} \mathbb{P}(U_n - \mathbb{E}h(\mathbf{x}_1, \dots, \mathbf{x}_m) \geq t) &\leq \exp(-2\lfloor n/m \rfloor t^2 / (b-a)^2), \\ \mathbb{P}(|U_n - \mathbb{E}h(\mathbf{x}_1, \dots, \mathbf{x}_m)| \geq t) &\leq 2 \exp(-2\lfloor n/m \rfloor t^2 / (b-a)^2), \end{aligned}$$

where $\lfloor x \rfloor$ denotes the greatest integer which is smaller than or equal to x . Hoeffding's inequality is a special case when $m = 1$.

Lemma 14 (Hoeffding's inequality). *Let X_1, \dots, X_n be i.i.d. random variables such that $a \leq X_i \leq b$ almost surely. Define $\bar{X} := \frac{1}{n} \sum_{i=1}^n X_i$. Then,*

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