



HAL
open science

Extending Kernel Testing To General Designs

Anthony Ozier-Lafontaine, Franck Picard, Bertrand Michel

► **To cite this version:**

Anthony Ozier-Lafontaine, Franck Picard, Bertrand Michel. Extending Kernel Testing To General Designs. 2024. hal-04747959

HAL Id: hal-04747959

<https://hal.science/hal-04747959v1>

Preprint submitted on 22 Oct 2024

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.

Extending Kernel Testing To General Designs

Anthony Ozier-Lafontaine *

Laboratoire de Mathématiques Jean Leray (UMR 6629 CNRS)
Nantes France
anthony.ozier-lafontaine@ec-nantes.fr

Bertrand Michel †

Laboratoire de Mathématiques Jean Leray (UMR 6629 CNRS)
Nantes France
bertrand.michel@ec-nantes.fr

Franck Picard ‡

LBMC, CNRS
Lyon France
franck.picard@ens-lyon.fr

Abstract

Kernel-based testing has revolutionized the field of non-parametric tests through the embedding of distributions in an RKHS. This strategy has proven to be powerful and flexible, yet its applicability has been limited to the standard two-sample case, while practical situations often involve more complex experimental designs. To extend kernel testing to any design, we propose a linear model in the RKHS that allows for the decomposition of mean embeddings into additive functional effects. We then introduce a truncated kernel Hotelling-Lawley statistic to test the effects of the model, demonstrating that its asymptotic distribution is chi-square, which remains valid with its Nystrom approximation. We discuss a homoscedasticity assumption that, although absent in the standard two-sample case, is necessary for general designs. Finally, we illustrate our framework using a single-cell RNA sequencing dataset and provide kernel-based generalizations of classical diagnostic and exploration tools to broaden the scope of kernel testing in any experimental design.

1 Introduction

Statistical hypothesis testing has witnessed significant advancements in recent years, largely due to the integration of kernel methods with non-parametric testing. Indeed, it has now become common practice to develop testing procedures that involve embedding probabilistic distributions into Reproducing Kernel Hilbert Spaces (RKHS) to assess the equality of distributions [8], dependencies [9], causality [25], and essentially any probabilistic feature that may distinguish two populations. Kernel testing has also facilitated the development of tests for complex data types, such as graphs and sequences, which surpass traditional non-parametric procedures like the Wilcoxon or Kolmogorov tests [17, 23]. Typically, kernel testing involves embedding data into a feature space using a kernel, with the statistical test based on the distribution of the embedded data, such as means and variances. A

*<https://anthoozier.github.io/>

†<http://bertrand.michel.perso.math.cnrs.fr/>

‡<https://franckpicard.github.io/>

significant milestone in the field was the Maximum Mean Discrepancy (MMD) test, which measures the distance between mean embeddings of two conditions [7]. The MMD statistic asymptotically follows a mixture of chi-square distributions, that can be approximated using permutations.

Parallel to the work of Gretton et. al [8], a studentized version of the MMD was proposed to account for the variance of the embeddings directly in the test statistic [11], which constitutes a kernelized version of the Hotelling-Lawley trace test [19]. This strategy results in a metric that can be interpreted as the distance between mean embeddings projected onto the kernel-Fisher discriminant axis. In other terms, the normalized version of the MMD reduces to the norm of a Kernel Fisher Discriminant Analysis classifier used for non-linear hypothesis testing. The original version is based on ridge-based regularization of the residual covariance operator [11]. This approach has recently been proven optimal with respect to the test’s ability to best separate two distributions within a class of alternatives [10]. These recent theoretical developments underscore the theoretical benefits of properly account for a noise model within the RKHS to gain power in the testing procedure through studentization. However, in practice, the ridge-based regularization test lacks interpretability and the asymptotic distribution is complex to compute. This motivated the development of the truncation-based test that is based on the eigen-decomposition of the residual covariance operator into T principal components [12]. This results in a statistic that asymptotically follows a $\chi^2(T)$ distribution, which greatly increases its applicability since it does not require permutation-based strategies, at the cost of a $\mathcal{O}(n^3)$ complexity. This complexity can be mitigated through Nystrom-type approximations. Despite methodological and theoretical developments in line with the MMD framework, the application of kernel testing in practice has been limited to the simple scenario of two-sample testing, whereas practical situations almost always involve more complex experimental designs. If kernel testing is to become a standard in various application areas, we need to offer extensions that are both statistically sound and efficiently implemented.

In this work, we propose a unified framework that enables kernel testing for any experimental design. We introduce a linear model in the RKHS that decomposes the mean embedding of observed distributions into linear functional effects. Testing is then performed using the kernelized Hotelling-Lawley statistic to test linear combinations of effects. Indeed, the truncated studentized-MMD can be easily interpreted as a truncated Hotelling-Lawley test based on data embeddings. Moreover, this connection between kernel testing and the standard multivariate framework allows us to reinterpret the standard MMD procedure as a linear model within the RKHS, linking the mean embeddings to any experimental design. Interestingly, according to our results, the comparison of linear combinations of any effect (contrast testing) cannot be conducted without specific conditions on the residual covariance operator—conditions that do not arise in the simpler two-sample test. These conditions relate to the standard homoscedasticity hypothesis required in linear models, which states that the variance of the noise must be independent of the expectation provided by the model. We reinterpret this hypothesis for covariance operators in the RKHS and propose a restricted heteroscedasticity hypothesis that states that the covariance operator of the noise is homogeneous up to its first T eigen-components. Under this hypothesis, we provide the asymptotic distribution of the test for linear combinations of effects (contrast tests), demonstrating that it follows a χ^2 distribution that is straightforward to compute. We also verify that this distribution remains valid when using a Nyström approximation, which enhances the practical applicability of kernel testing. To advance the full operability of kernel testing, we offer a comprehensive framework that benefits from powerful representations and diagnostic plots, enabling the complete application of kernel testing in any experimental situation. Moreover, all these quantities are implemented in our package *ktest*⁴.

To illustrate our framework we fully analyze a dataset from single-cell transcriptomic data. Thanks to the combination of single-cell isolation techniques and massive parallel sequencing, it is now possible to create high-dimensional molecular portraits of cell populations [16, 26]. These advances have resulted in the production of complex, high-dimensional data, revolutionizing our understanding of the complexity of living tissues. The field of single-cell data science presents new methodological challenges, key among them being the statistical comparison of single-cell RNA sequencing (scRNA-Seq) datasets between conditions or tissues. This step is crucial for distinguishing biological from technical variabilities and for asserting meaningful expression differences. Comparative analysis of single-cell datasets, regardless of their type, is an essential component of single-cell data science, providing biological insights and opening therapeutic perspectives through the identification of biomarkers and therapeutic targets. The effectiveness of kernel testing in this context has recently

⁴The *ktest* package is available at <https://github.com/LMJL-Alea/ktest>.

been demonstrated [20]; however, the practical application of kernel testing in the field hinges on the ability to conduct a comprehensive analysis tailored to any experimental design.

2 Testing linear combinations of functional effects in the feature space

2.1 A linear model in the feature space

We consider $\mathbf{Y} = (Y_1, \dots, Y_n)$, a set of n observations from a measurable space \mathcal{Y} , jointly recorded with some explanatory variables encoded in the design matrix $\mathbf{X} = (x_1, \dots, x_n)' \in \mathcal{M}_{n,p}(\mathbb{R})$, where for $i \in \{1, \dots, n\}$, $x_i = (1, x_i^1, \dots, x_i^{p-1})' \in \mathbb{R}^p$ is the experimental design that is supposed to be fixed. To generalize kernel testing to any experimental design we start by considering a featurization of the response variable. For this purpose we introduce $k(\cdot, \cdot) : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, a positive definite kernel associated to a separable Reproducing Kernel Hilbert Space (RKHS) $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$ called the feature space, and the feature map $\phi(\cdot) : Y \in \mathcal{Y} \mapsto k(Y, \cdot) \in \mathcal{H}$. Then we propose a linear model in the feature space, such that for $i \in \{1, \dots, n\}$, we have:

$$\phi(Y_i) = \theta_0 + x_i^1 \theta_1 + \dots + x_i^{p-1} \theta_{p-1} + e_i,$$

where $\mathbb{E}(e_i) = 0_{\mathcal{H}}$ and $\mathbb{V}(e_i) = \Sigma_i$, the residual covariance operator. This model corresponds to the decomposition of the mean embedding of the distribution of the response variable into linear combination of mean embeddings $\Theta = (\theta_0, \theta_1, \dots, \theta_{p-1})' \in \mathcal{H}^p$ that correspond to different factors. A matrix formulation of this model can also be considered, with $\Phi(\mathbf{Y}) = (\phi(Y_1), \dots, \phi(Y_n))' \in \mathcal{H}^n$, $\mathbf{E} = (e_1, \dots, e_n)' \in \mathcal{H}^n$, which provides the linear model in the feature space that resumes to:

$$\Phi(\mathbf{Y}) = \mathbf{X}\Theta + \mathbf{E}, \quad (1)$$

where we use a natural extension of the matrix product between $\mathbf{X} \in \mathcal{M}_{n,p}(\mathbb{R})$ and $\Theta \in \mathcal{H}^p$. Functional parameters can then be easily inferred using a least squares estimator such that

$$\widehat{\Theta} = (\mathbf{X}'\mathbf{X})^{-}\mathbf{X}'\Phi(\mathbf{Y}) \in \mathcal{H}^p, \quad (2)$$

where $(\mathbf{X}'\mathbf{X})^{-}$ is a generalized inverse of $\mathbf{X}'\mathbf{X}$, see Appendix B.1 for a proof. Then the predicted embeddings according to the model are defined as $\widehat{\Phi}(\mathbf{Y}) = \mathbf{P}_{\mathbf{X}}\Phi(\mathbf{Y})$, where the orthogonal projection $\mathbf{P}_{\mathbf{X}} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-}\mathbf{X}' \in \mathcal{M}_n(\mathbb{R})$ on $\text{Im}(\mathbf{X}) \subset \mathbb{R}^n$ is also seen as an operator on \mathcal{H}^n . The residuals of the model are defined by $\widehat{\mathbf{E}} = (\widehat{e}_1, \dots, \widehat{e}_n) = \mathbf{P}_{\mathbf{X}}^{\perp}\Phi(\mathbf{Y}) = \mathbf{P}_{\mathbf{X}}^{\perp}\mathbf{E} \in \mathcal{H}^n$ where $\mathbf{P}_{\mathbf{X}}^{\perp} = \mathbf{I}_n - \mathbf{P}_{\mathbf{X}}$ is the orthogonal projection on $\text{Im}(\mathbf{X})^{\perp}$.

One-way and two-way designs. We provide two examples of linear models in the feature space that are most commonly encountered in practice. The one-way design model with $I \geq 2$ levels of sizes n_1, \dots, n_I with $n = \sum_{i=1}^I n_i$. For $i \in \{1, \dots, I\}$ and $j \in \{1, \dots, n_i\}$, $Y_{i,j}$ stands for the j^{th} observation of level i . The model is such that: $\mathbb{E}(\phi(Y_{i,j})) = \mu + \alpha_i$, where $\Theta = (\mu, \alpha_1, \dots, \alpha_I)$ are the model functional parameters. The standard two-sample test corresponds to the model with $I = 2$ levels, which can easily be generalized to multiple comparisons. A natural extension would be to consider a two-way design to introduce blocking factors for instance. In this setting, the first effect takes I distinct values and the second one takes J distinct values, and for $i, j \in \{1, \dots, I\} \times \{1, \dots, J\}$, we observe $\ell = 1, \dots, n_{i,j}$ observations associated to effects (i, j) . We propose to model these effects with the following additive linear model in the feature space: $\mathbb{E}(\phi(Y_{i,j,\ell})) = \mu + \alpha_i + \beta_j$. We then can design contrast matrices to test if there is an influence of the first effect and the second effect respectively, as see an application in Section 5.

Conditional Mean Embedding. For simplicity, our model is proposed for a deterministic design, but it can also be adapted for a random design. In this latter setting, Model (1) can be easily reinterpreted in terms of conditional mean embedding [24]. Specifically, our model assumes that the conditional mean embedding $\mathcal{U}_{Y|X}$ is the linear operator mapping $x \in \mathbb{R}^p$ to $x'\Theta \in \mathcal{H}$, by considering the identity embedding for $x \in \mathbb{R}^p$.

2.2 Hypothesis testing with the truncated kernel Hotelling-Lawley statistic

We propose to perform hypothesis testing on the model functional parameters with a standard contrast testing approach. Let $\mathbf{L} \in \mathbb{R}^{d \times p}$ a surjective matrix. As for \mathbf{X} , \mathbf{L} can be seen both as a linear operator

from \mathbb{R}^p to \mathbb{R}^d , or as a linear operator from \mathcal{H}^p to \mathcal{H}^d . We use the contrast matrix \mathbf{L} to test if a combination of Θ is null or not and we formulate the null and alternative hypotheses as :

$$H_0 : \mathbf{L}\Theta = 0 \quad \text{vs.} \quad H_1 : \mathbf{L}\Theta \neq 0.$$

In the setting of the multivariate linear model, several popular statistical tests exist for this type of hypothesis. We propose to extend the Hotelling-Lawley test to the linear model in the feature space (1), which requires considering general operators defined on Hilbert spaces. In this section we introduce the statistic without providing all the justifications, a rigorous definition of the kernelized Hotelling-Lawley statistic is proposed in Appendices A and B, and more specifically in Section B.2.

We first introduce the so-called test operator (Hilbert-Schmidt) associated with \mathbf{L} :

$$\widehat{H}_{\mathbf{L}} = (\mathbf{L}\widehat{\Theta})^*(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')^{-1}(\mathbf{L}\widehat{\Theta}), \quad (3)$$

which generalizes the "between sum of squares" from ANOVA and the "cross products matrix" from MANOVA models to more complex designs in the RKHS setting of this paper. Next, we define the residual covariance operator as $\widehat{\Sigma} = n^{-1} \sum_{i=1}^n \widehat{e}_i \otimes \widehat{e}_i$. The spectral theorem applies to this operator: $\widehat{\Sigma} = \sum_{t=1}^r \widehat{\lambda}_t (\widehat{f}_t \otimes \widehat{f}_t)$, where $(\widehat{f}_t)_{t \geq 0}$ is an orthonormal basis of \mathcal{H} , where $\widehat{\lambda}_t > 0$ and r is the rank of $\widehat{\Sigma}$. We then introduce a generalized inverse for $\widehat{\Sigma}$ as $\widehat{\Sigma}^- = \sum_{t=1}^r \widehat{\lambda}_t^{-1} (\widehat{f}_t \otimes \widehat{f}_t)$. We finally introduce the truncated kernelized Hotelling-Lawley (TKHL) statistic as

$$\widehat{\mathcal{F}}_T = n^{-1} \text{tr} \left(\widehat{\Sigma}_T^{-1} \widehat{H}_{\mathbf{L}} \right),$$

where $\widehat{\Sigma}_T^{-1} := \sum_{t=1}^{T \wedge r} \widehat{\lambda}_t^{-1} (\widehat{f}_t \otimes \widehat{f}_t)$ is defined from $\widehat{\Sigma}^-$ by a spectral truncation of depth T . An alternative ridge-based regularization has been proposed for $\widehat{\Sigma}^-$, in the spirit of kernel FDA [11]. The spectral truncation we propose has the advantages of *i)* providing a tractable chi-square asymptotic distribution under the null, as we show in Theorem 1, and *ii)* providing a natural framework for data and model exploration (Section 5). Computing this statistic is based on a kernel trick detailed in the Appendix (Section F).

One-way and two-way designs. In the one-way design case, the following "one-versus-all" contrast matrix allows us to test $H_0 : \alpha_1 = \alpha_2 = \dots = \alpha_I$ versus $H_1 : \exists i, i', \alpha_i \neq \alpha_{i'}$:

$$\mathbf{L}_{\alpha} = \begin{pmatrix} 0 & 1 - \frac{n_1}{n} & -\frac{n_2}{n} & \dots & -\frac{n_I}{n} \\ 0 & -\frac{n_1}{n} & 1 - \frac{n_2}{n} & \dots & -\frac{n_I}{n} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & -\frac{n_1}{n} & \dots & 1 - \frac{n_{I-1}}{n} & -\frac{n_I}{n} \end{pmatrix} \quad (4)$$

When $I = 2$ we exactly retrieve the two-sample test, since the model writes $\mathbb{E}(\phi(Y_{i,j})) = \mu_i$ and testing the null $H_0 : \alpha_1 = \alpha_2$ resumes to testing $H_0 : \mu_1 = \mu_2$. In this setting, the truncated Hotelling-Lawley statistic is exactly the truncated KFDA statistic [11, 12]. Our framework provides a straightforward generalization to the comparison of I levels thanks to the truncated Hotelling-Lawley test. When two additive factors are present (two-way design), we can design contrast matrices to test if one factor influences the featurized response variable independently of the other. For instance, this allows for testing whether the blocking factor has a significant effect and thus should be considered when evaluating the treatment factor. Our framework enables the application of reasoning similar to that used in standard linear model procedures, thereby greatly extending the scope of kernel testing. We provide an example in Section 5.

3 Asymptotic distribution of the TKHL statistic

We derive the asymptotic distribution of the TKHL statistic, demonstrating that it follows a chi-square distribution under the null, thereby generalizing the known results for the HL trace statistic to the RKHS framework. Additionally, we introduce a homoscedasticity hypothesis for the error terms, as is typical in vectorial models, and provide a functional generalization within the RKHS by assuming restricted homoscedasticity to the first T eigen-directions of the residual covariance operator.

Assumption 1. *The covariances Σ_i of the errors are all identical on their T first directions: there exists a symmetric positive operator $\Sigma_T = \sum_{t=1}^T \lambda_t f_t^{\otimes 2}$ such that $\Sigma_i = \Sigma_T + \Sigma_{i,T^c}$ for all $i \in \{1, \dots, n\}$, where $\Sigma_{i,T^c} = \sum_{t>T} \lambda_{i,t} f_{i,t}^{\otimes 2}$ is a symmetric positive operator.*

Assumption 2. *The eigenvalues of λ_t of Σ_T are all simple, and there exists $\mu_{T+1} > 0$ such that $\lambda_{i,T+1} \leq \mu_{T+1} < \lambda_T$ for any $i \in \{1, \dots, n\}$.*

Under Assumptions 1 and 2 and for all $i \in \{1, \dots, n\}$, the eigenvalues of Σ_i satisfy $\lambda_1 > \dots > \lambda_T > \lambda_{T+1} \geq \lambda_{i,T+1}$. Next assumption require the kernel function to be bounded:

Assumption 3. *The kernel is bounded : $\|k(\cdot, \cdot)\|_\infty = M_k < +\infty$.*

Finally, we make the three assumptions on the design $\mathbf{X} = \mathbf{X}(n) \in \mathcal{M}_{n,p}(\mathbb{R})$. In the following we will omit the dependency on n since it is clear that it depends on n .

Assumption 4. *There exist $M_x \in \mathbb{R}$ such that for all $n \geq 1$, the design matrix \mathbf{X} is such that $\|\mathbf{X}\|_\infty := \sup_{i,j} |x_{i,j}| \leq M_x$.*

Assumption 5. *The maximal coefficient on the diagonal of $\mathbf{P}_\mathbf{X}$ satisfies $\max_{i=1\dots n} |(\mathbf{P}_\mathbf{X})_{i,i}| \rightarrow 0$ as n tends to infinity.*

Assumption 6. *The sampling gives a convergent design : $\frac{1}{n} \mathbf{X}' \mathbf{X} \rightarrow W^{-1} \in \mathbb{R}^{p \times p}$ as n tends to infinity.*

Assumptions are discussed after the following Theorem which demonstrates that the asymptotic distribution of the TKHL statistic is not only distribution-free but also tractable.

Theorem 1. *Under the linear model (1) in the feature space \mathcal{H} , assume that Assumptions 1 to 6 are satisfied. If H_0 is true then $n\widehat{\mathcal{F}}_T \xrightarrow[n \rightarrow \infty]{\mathcal{D}} \chi_{dT}^2$.*

Our result generalizes the asymptotic convergence of the HL trace statistic, commonly used in multivariate linear models (see, for instance, Theorem 12.8 from [19]), to cases where data lies in a RKHS. In this functional framework, establishing the asymptotic properties is more intricate since the T directions used to define the TKHL are stochastic.

The proof of the Theorem is given in Appendix C, see in particular Section C.1 for a summary of the proof. The proof relies on the convergence in probability of $\widehat{\Sigma}$ towards Σ , which we establish thanks to an approach inspired from [22] and [28]. This approach combines a perturbation bound with a bounded difference theorem. In our framework we must control the residual covariance operator, rather than a standard covariance operator as is typically done in kernel PCA.

Regarding the assumptions of the theorem, Assumption 1 relates to the classical homoscedasticity hypothesis required in linear models, generalized here to covariance operators with a restriction on the first T directions. This hypothesis is typically verified visually using diagnostic plots, which we generalize for the feature space (Section 5). In scenarios where testing involves the equality of I distributions with a characteristic kernel, the null hypothesis entails the equality of the I covariance operators. In such cases, the homoscedasticity assumption may be omitted as specified in Theorem 1. However, when extending this approach to general designs, maintaining the homoscedasticity assumption is crucial and cannot simply be substituted with a *full homoscedasticity* assumption, which demands equality of all covariances across all directions in \mathcal{H} . Given that the covariance operator happens to be itself a mean embedding that may characterize the distribution in certain contexts (see for instance Proposition 1 in [1] and also [3]), a full homoscedasticity assumption would impose a stronger condition than the null hypothesis itself.

Assumptions 2 and 3 are strong but standard hypotheses for deriving statistical results on kernel methods. Assumption 4 is used in addition to Assumption 3 to obtain a bound on the errors and to show that their fourth moment is finite (see Lemma 4). Note that the fourth moment of the errors are also assumed to be finite in the standard HL test (see Theorem 12.8 from [19]).

Finally, Assumptions 5 and 6 are also needed to obtain the asymptotic distribution of the standard HL statistic (see for instance Theorem 12.8 from [19]).

4 The Nyström TKHL statistic

The computational complexity of kernel testing based on the TKHL statistic depends on the diagonalization of the residual covariance operator $\widehat{\Sigma}_T$, or its kernel-trick equivalent, the gram matrix \mathbf{K}_e of general term $\langle \widehat{e}_i, \widehat{e}_{i'} \rangle_{\mathcal{H}}$, which requires $\mathcal{O}(n^3)$ operations. To reduce the computational burden,

the Nyström approximation consists in finding a low-rank approximation such that $\mathbf{K}_e \simeq \mathbf{A}\mathbf{A}^\top$ with $\mathbf{A} \in \mathcal{M}_{n,m}(\mathbb{R})$, $m \ll n$, which reduces the cost to $\mathcal{O}(m^3)$. To choose a matrix \mathbf{A} that ensures the same asymptotic distribution as the original TKHL statistic, we uniformly sample q landmarks from the n observations and define the $m \geq T$ first unit eigenfunctions $(\hat{a}_1, \dots, \hat{a}_m)$ of their associated residual covariance operator as Nyström anchors. Finally we build matrix \mathbf{A} , as the $n \times m$ matrix whose rows are the projections of residuals \hat{e}_i onto the eigenfunctions \hat{a}_j , with general term $\mathbf{A}_{i,j} = \langle \hat{e}_i, \hat{a}_j \rangle_{\mathcal{H}}$, $i = 1, \dots, n$, $j = 1, \dots, m$. We then define the Nyström residual covariance operator $\widehat{\Sigma}^a$ as :

$$\widehat{\Sigma}^a = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^a \otimes \hat{e}_i^a,$$

where for $i \in \{1, \dots, n\}$, \hat{e}_i^a is the orthogonal projection of \hat{e}_i onto $\text{Span}(\hat{a}_1, \dots, \hat{a}_m)$. We denote by $\widehat{f}_1^a, \dots, \widehat{f}_m^a$ the orthonormal eigenfunctions of $\widehat{\Sigma}^a$ associated to the non-increasing eigenvalues $\widehat{\lambda}_1^a, \dots, \widehat{\lambda}_m^a$, so that $\widehat{\Sigma}_T^a = \sum_{t=1}^T \widehat{\lambda}_t^a \widehat{f}_t^a \otimes \widehat{f}_t^a$ is the spectral truncation of $\widehat{\Sigma}^a$. We can then introduce as before a generalized inverse for $\widehat{\Sigma}_T^a$ as

$$\widehat{\Sigma}_T^{a^{-1}} = \sum_{t \leq T, \widehat{\lambda}_t^a > 0} \widehat{\lambda}_t^{a^{-1}} \widehat{f}_t^a \otimes \widehat{f}_t^a.$$

We finally define a Nyström version of the kernel Hotelling-Lawley Statistic as

$$\widehat{\mathcal{F}}_T^{\text{NY}} = \frac{1}{n} \text{tr} \left(\widehat{\Sigma}_T^{a^{-1}} \widehat{H}_{\mathbf{L}} \right). \quad (5)$$

Note that according to this version of Nyström we don't use the landmarks to approximate $\widehat{H}_{\mathbf{L}}$, as the computation cost of this last is not the worst part in the TKHL statistic.

We now study the asymptotic distribution of the Nyström TKHL statistic defined in the linear model (1). We consider the Nyström TKHL statistic defined with q landmarks, m anchors and a truncate parameter $T \leq m$. In the following the parameters T and m are fixed, contrary to n and q , and we assume that $q \leq n$. We need to rewrite Assumptions 1 and 2 according to the numbers of anchors:

Assumption 1'. *The covariances Σ_i of the errors are all identical on their m first directions: there exists a symmetric positive operator $\Sigma_m = \sum_{t=1}^m \lambda_t f_t^{\otimes 2}$ such that for $i \in \{1, \dots, n\}$*

$$\Sigma_i = \Sigma_m + \Sigma_{i,T^c},$$

where $\Sigma_{i,T^c} = \sum_{t>T} \lambda_{i,t} f_{i,t}^{\otimes 2}$ is a symmetric positive operator.

Assumption 2'. *The eigenvalues of Σ_m are all simple, and there exists a positive constant μ_{m+1} such that for any $i \in \{1, \dots, n\}$, $\lambda_{i,m+1} \leq \mu_{m+1} < \lambda_m$.*

Theorem 2. *Under Assumptions 1', 2', 3, 4, 5 and 6, if H_0 is true then:*

$$n \widehat{\mathcal{F}}_T^{\text{NY}} \xrightarrow[q, n \rightarrow \infty]{\mathcal{D}} \chi_{dT}^2,$$

where in the convergence both q and n tends to infinity and $q \leq n$.

The proof of the Theorem is given in Section D.

In the theorem, landmarks are selected uniformly from the observations. However, in practice, choosing landmarks uniformly across observations may not always be appropriate. For instance, when comparing unbalanced populations, under-represented groups could be over-sampled to ensure that their within-covariance matrices are accurately estimated. This approach also improves the accuracy of reconstructing the within covariance operator using the anchors.

5 Leveraging kernel testing in single-cell transcriptomics

Experimental design and model. We consider the analysis of a gene expression dataset⁵ that consists of the expression measurements of 83 genes on 685 cells [27]. In this context single-cell

⁵Data are available under licence CC0, under the SRA repository number SRP076011

transcriptomics has been performed to investigate the cell differentiation process of chicken primary erythroid progenitor cells (T2EC). The experimental design is the following: to investigate the impact of the medium on cell differentiation, undifferentiated cells were initially put in a self-renewal medium (0H level of the medium effect), then put in a differentiation-inducing medium for 24h (24H). The population was then split into a first population maintained in the same medium for an additional 24h to achieve differentiation (48HDIFF), the second population was put back in the self-renewal medium to investigate potential reversion (48HREV). Before testing the medium effect, we must account for the fact that data were acquired on different dates, which induces a 8-level batch effect. We consider the linear model in the feature space:

$$\phi(Y_{i,j,\ell}) = \mu + \alpha_i + \beta_j + e_{i,j,\ell}, \quad (6)$$

where $Y_{i,j,\ell}$ is the vector of log-normalized expression (size 83) for cell $\ell = 1, \dots, n_{i,j}$ in batch $i = 1, \dots, I$, ($I = 8$) and medium $j = 0H, 24H, 48HDIFF, 48HREV$ and $\phi(Y_{i,j,\ell})$ its featurization. Then μ is the intercept, $\alpha_1, \dots, \alpha_8$ are the model parameters associated to the batch effect and $\beta_{0H}, \beta_{24H}, \beta_{48HDIFF}, \beta_{48HREV}$ are the model parameters associated to the medium effect. Analyses were performed with the Gaussian kernel with bandwidth calibrated by the median heuristic [6].

Diagnostic plots. In practice, any analysis based on a linear model begins with diagnostic plots that enable the visual inspection of linearity and homoskedasticity hypotheses, as well as to perform model checking [13, 19]. To provide such an operational framework in the feature space, we consider the projection of the mean effects $\widehat{\Phi(\mathbf{Y})} = \mathbf{P}_{\mathbf{X}}\Phi(\mathbf{Y})$ and the residuals of the model $\Phi(\mathbf{Y}) - \widehat{\Phi(\mathbf{Y})}$ onto the eigen-elements of the residual covariance operator $\widehat{\Sigma}_T$ obtained from Eq. (6). Visual inspection on the first direction (\hat{f}_1) confirm that the assumptions of linearity and homoscedasticity in the feature space are met (Fig. 1), also confirmed on other directions (not shown).

Testing with the TKHL statistic. We initiated our analysis by examining technical variability in the experiment, starting with a test for batch effects before assessing the main "medium" effect, which is of biological interest. We tested the null hypothesis of no batch effect, $H_0 : \alpha_1 = \dots = \alpha_8$, using the TKHL method based on the residual covariance operator associated to the one-way model with μ and $\alpha_1, \dots, \alpha_8$ only and the corresponding \mathbf{L}_{α} written in (4). This test revealed a significant difference between batches (p -value = 5.06×10^{-48} for $T = 3$). This initial finding indicates that testing for the medium effect must be conducted after correcting for the batch effect, which, in our setting, is achieved by testing for the medium effect, $H_0 : \beta_{0H} = \beta_{24H} = \beta_{48HDIFF} = \beta_{48HREV}$, using the TKHL method based on the residual covariance operator associated to (6) and a matrix $\mathbf{L}_{\beta|\alpha}$ corresponding to (4) supplemented with columns of zeros on corresponding to $\alpha_1, \dots, \alpha_8$. This test showed highly significant differences (p -value = 4.17×10^{-237} for $T = 3$), indicating that gene expression varies significantly across at least one medium compared to others. To deepen the analysis, we complemented this global test with pairwise comparisons to identify which mediums differ from each other. For this, we considered the pairwise hypotheses $H_0^{j,j'} : \beta_j = \beta_{j'}$, using for instance $\mathbf{L}_{0H,24H} = (0, 1, -1, 0, 0, 0, 0, 0, 0, 0, 0, 0)$. It revealed that the two closest mediums are 0H and 48HREV (p -value = 6.9×10^{-26} for $T = 3$, after Bonferroni adjustment for multiple comparison) and that the most distinct mediums are 0H and 48HDIFF (p -value = 9.95×10^{-198} for $T = 3$, Bonferroni adjusted).

Representation of the linear decomposition. To enhance the visual exploration of our model in the feature space, we propose a representation of the observations based on the linear decomposition suggested by our model. We extend the approach commonly employed in Kernel Fisher Discriminant Analysis (KFDA), which involves inspecting the empirical density of observations projected onto the discriminant axis [20]. Specifically, for the linear model in the feature space, we consider the eigen-directions of the operator $\widehat{\Sigma}_T^{-1}\widehat{H}_L$, used to test the "batch" effect. Although this operator has a rank of $I - 1$, we simplify our analysis by focusing on the projection of observations onto its first direction (Fig. 1) that we call the discriminant axis. We believe this representation provides a powerful visualization of the variability of the data with respect to the linear decomposition of the mean embedding. When applied to single-cell transcriptomic data, this visualization initially demonstrates that despite the statistical significance of the "batch" effect, cells from different batches appear closely positioned on the discriminant axis dedicated to the batch effect, suggesting minimal practical differences between batches. Subsequently, when applied to the operator used for assessing

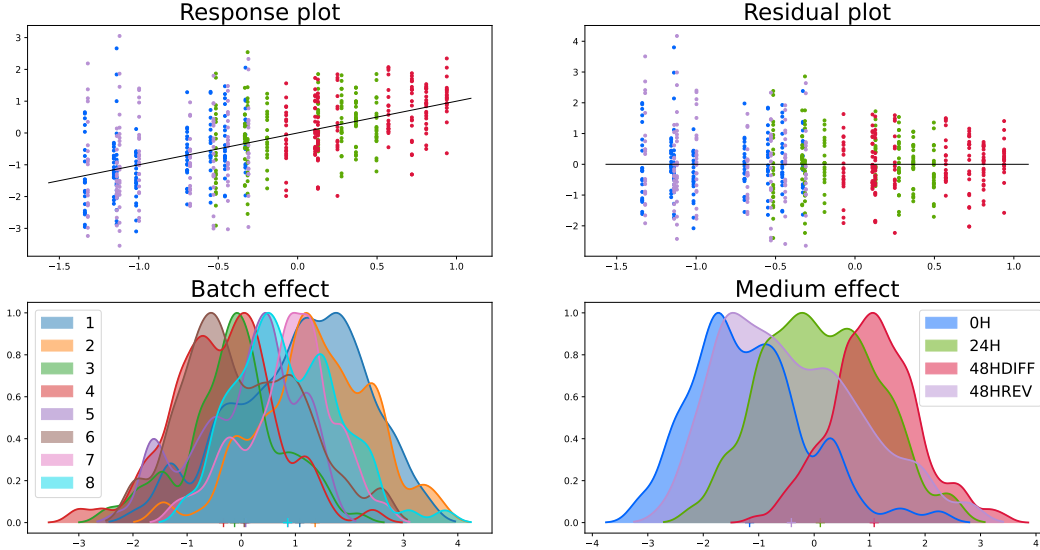


Figure 1: Top: Response plot (left) and residual plot (right) associated to $t = 1$. The identity (left) and horizontal axis (right) are in black. Colors correspond to the medium effect. Bottom : Discriminant axes associated to testing for the batch effect (left) and the medium effect (right).

the "medium" effect, the visualization underscores the proximity of cells in the initial medium (0H, blue) to those returned to the self-renewal medium (48HREV, purple), indicating a reversal in their differentiation. Cells from other mediums (24H, green; 48HDIFF, red) are organized chronologically along the discriminant axis dedicated to the medium effect, which may then be interpreted as a differentiation trajectory of cells over time.

Detecting influential observations with the kernelized Cook distance. We conclude our analysis by proposing a kernelized version of the Cook's distance that is classically used to detect influential observations in linear models [5]. Denoting by $\hat{\Theta}_{(i)}$ the estimated parameters obtained when ignoring the i^{th} observation, we propose a kernelized Cook's distance adapted to any linear combination of $\hat{\Theta}$ to match with the TKHL test we perform, such that:

$$\mathcal{D}_{\mathbf{L}} = \frac{1}{d} \text{tr} \left((\mathbf{L}(\hat{\Theta} - \hat{\Theta}_{(i)}))' \text{Cov}(\mathbf{L}\hat{\Theta})^{-1} (\mathbf{L}(\hat{\Theta} - \hat{\Theta}_{(i)})) \right),$$

where we can show that $\text{Cov}(\mathbf{L}\hat{\Theta}) = ((\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')_{h,h'})_{h,h' \in \{1,\dots,p\}}$. The proofs of the results and the details of the computation of the kernel Cook's distance are developed in Appendix F.3. We conclude our analysis by displaying $\mathcal{D}_{\mathbf{L}|\beta|\alpha}$ according to the position of cells along the discriminant axis (Fig. 2) to identify observations that leverage the medium effect. Our results show that the most influential observations are cells from mediums 0H and 48HREV, positioned at the right margin of the discriminant axis. This positioning may indicate cells with unexpected gene expression patterns in relation to the differentiation process. Since the interpretability of kernel methods relies solely on individual observations rather than input features, identifying outlier observations is crucial in practice to fully exploit the potential of the linear model in the feature space.

6 Discussion and Limitations

Here, we introduce a novel framework that generalizes the standard two-sample kernel test to accommodate any general design, thereby significantly expanding the applicability of kernel mean embedding to practical experimental designs. Our approach combines the benefits of linear models—particularly in terms of the interpretability of the linear decomposition of mean embeddings in the RKHS—with the proven power of non-parametric kernel testing. The statistic we propose constitutes a functional generalization of the Hotelling-Lawley trace test and can also be easily

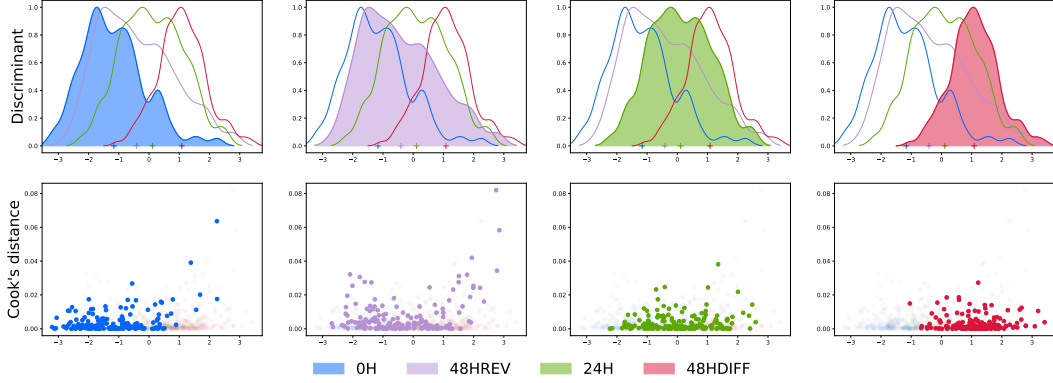


Figure 2: Influence of the observations with respect to the position of their projection on the first generalized discriminant directions associated to the Hotelling-Lawley operator used to test the effect of the medium with $T = 3$.

understood as a ratio between the model-driven variance operator and the residual variance operator. Additionally, classical strategies such as Wilk’s Lambda, Pillai’s Trace, and Roy’s Largest Root, which are all based on the two operators \hat{H}_L and $\hat{\Sigma}$, could also be easily adapted to the RKHS framework. The connection we propose between linear models and kernel testing allows us to enrich kernel testing with well established diagnostic and exploratory tools, extending the applicability and interpretability of kernel testing.

Our theoretical work establishes that the truncated kernel Hotelling-Lawley (TKHL) statistic asymptotically follows a chi-square distribution under the null hypothesis. From a computational standpoint, regularization via truncation proves more tractable than ridge penalization. This advantage is further underscored by our second theorem, which demonstrates that this asymptotic behavior also holds true for the Nyström version of the TKHL, thereby extending its applicability to large datasets where such assumptions are practical. Although permutations offer an alternative strategy for computing quantiles of the null distribution, the complexity of studentized functional contrast testing, which scales as $\mathcal{O}(n^3)$, makes easily computed approximations essential, if not indispensable, for handling large data.

Regarding the selection of landmarks in the Nyström procedure, the generalization of kernel testing to any design presents a new challenge: selecting landmarks that accurately reflect the variability induced by the experimental design. Traditional uniform sampling should be adapted to incorporate more advanced approaches. For example, Determinantal Point Processes can be used to sample landmarks for the Nyström method [15], which could hopefully lead to an increase in the power of the downstream testing procedure.

In future work, we plan to conduct a non-asymptotic analysis of our test procedure, similar to recent minimax power results achieved in ridge-regularized scenarios [10]. This direction could pave the way for theoretically grounding the choice of the truncation hyper-parameter, currently determined empirically. Another promising avenue involves establishing a link between the linear decomposition we propose in the RKHS and the initial distributions compared in the input space. This link is guaranteed in the two-sample case by the injectivity of the feature map and the characteristic property of the kernel. Yet, the relationship between the linear functional effects and their distributional counterparts in the feature space remains to be elucidated.

Finally, our current example employs a Gaussian kernel calibrated by the median heuristic, but our flexible framework could potentially incorporate aggregation strategies to broaden the choice of kernels, as suggested by recent studies [21].

References

- [1] F. Bach. Information theory with kernel methods. *IEEE Transactions on Information Theory*, 69(2):752–775, 2022. Publisher: IEEE.
- [2] R. Bhatia and L. Elsner. The Hoffman-Wielandt inequality in infinite dimensions. *Proceedings of the Indian Academy of Sciences - Mathematical Sciences*, 104(3):483–494, Aug. 1994.
- [3] P. Bonnier, H. Oberhauser, and Z. Szabó. Kernelized cumulants: Beyond kernel mean embeddings. *Advances in Neural Information Processing Systems*, 36, 2024.
- [4] J. B. Conway. *A course in functional analysis*. Number 96 in Graduate texts in mathematics. Springer, New York, 2nd ed edition, 1997.
- [5] J. A. Diaz-Garcia and G. Gonzalez-Farias. Sensitivity analysis in linear regression. 2005.
- [6] D. Garreau, W. Jitkrittum, and M. Kanagawa. Large sample analysis of the median heuristic, Oct. 2018. arXiv:1707.07269 [math, stat].
- [7] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A Kernel Method for the Two-Sample-Problem. In *Advances in Neural Information Processing Systems*, volume 19. MIT Press, 2006.
- [8] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.
- [9] A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf. Measuring Statistical Dependence with Hilbert-Schmidt Norms. In S. Jain, H. U. Simon, and E. Tomita, editors, *Algorithmic Learning Theory*, Lecture Notes in Computer Science, pages 63–77, Berlin, Heidelberg, 2005. Springer.
- [10] O. Hagrass, B. K. Sriperumbudur, and B. Li. Spectral Regularized Kernel Two-Sample Tests, Dec. 2022. arXiv:2212.09201 [cs, math, stat].
- [11] Z. Harchaoui, F. Bach, and E. Moulines. Testing for homogeneity with kernel Fisher discriminant analysis. *Advances in Neural Information Processing Systems*, 20, 2007.
- [12] Z. Harchaoui, F. Vallet, A. Lung-Yut-Fong, and O. Cappe. A regularized kernel-based approach to unsupervised audio segmentation. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1665–1668, Taipei, Taiwan, Apr. 2009. IEEE.
- [13] R. A. Johnson and D. W. Wichern. *Applied multivariate statistical analysis*. 2002. Publisher: Prentice hall Upper Saddle River, NJ.
- [14] V. Koltchinskii and E. Giné. Random matrix approximation of spectra of integral operators. *Bernoulli*, pages 113–167, 2000. Publisher: JSTOR.
- [15] C. Li, S. Jegelka, and S. Sra. Fast dpp sampling for nystrom with application to kernel methods. In *International Conference on Machine Learning*, pages 2061–2070. PMLR, 2016.
- [16] E. Macosko, A. Basu, R. Satija, J. Nemesh, K. Shekhar, M. Goldman, I. Tirosh, A. Bialas, N. Kamitaki, E. Martersteck, J. Trombetta, D. Weitz, J. Sanes, A. Shalek, A. Regev, and S. McCarroll. Highly Parallel Genome-wide Expression Profiling of Individual Cells Using Nanoliter Droplets. *Cell*, 161(5):1202–1214, May 2015.
- [17] H. B. Mann and D. R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60, 1947. Publisher: JSTOR.
- [18] C. McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989. Publisher: Norwich.
- [19] D. J. Olive, D. J. Olive, and Chernyk. *Robust multivariate analysis*. Springer, 2017.
- [20] A. Ozier-Lafontaine, C. Fourneaux, G. Durif, P. Arsenteva, C. Vallot, O. Gandrillon, S. Gonin-Giraud, B. Michel, and F. Picard. Kernel-based testing for single-cell differential analysis. *Genome Biology*, 25(1):114, May 2024.

- [21] A. Schrab, I. Kim, M. Albert, B. Laurent, B. Guedj, and A. Gretton. MMD Aggregated Two-Sample Test, June 2022. arXiv:2110.15073 [cs, math, stat].
- [22] J. Shawe-Taylor. On the Eigenspectrum of the Gram Matrix and the Generalisation Error of Kernel PCA. *IEEE Transactions on Information Theory*, 51(7):2510–2522, 2005.
- [23] N. V. Smirnov. On the estimation of the discrepancy between empirical curves of distribution for two independent samples. *Bull. Math. Univ. Moscou*, 2(2):3–14, 1939.
- [24] L. Song, J. Huang, A. Smola, and K. Fukumizu. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 961–968, 2009.
- [25] K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI’11*, page 804–813, Arlington, Virginia, USA, 2011. AUAI Press.
- [26] G. X. Y. Zheng, J. M. Terry, P. Belgrader, P. Ryvkin, Z. W. Bent, R. Wilson, S. B. Ziraldo, T. D. Wheeler, G. P. McDermott, J. Zhu, M. T. Gregory, J. Shuga, L. Montesclaros, J. G. Underwood, D. A. Masquelier, S. Y. Nishimura, M. Schnall-Levin, P. W. Wyatt, C. M. Hindson, R. Bharadwaj, A. Wong, K. D. Ness, L. W. Beppu, H. J. Deeg, C. McFarland, K. R. Loeb, W. J. Valente, N. G. Ericson, E. A. Stevens, J. P. Radich, T. S. Mikkelsen, B. J. Hindson, and J. H. Bielas. Massively parallel digital transcriptional profiling of single cells. *Nature Communications*, 8:14049, Jan. 2017.
- [27] S. Zreika, C. Fourneau, E. Vallin, L. Modolo, R. Seraphin, A. Moussy, E. Ventre, M. Bouvier, A. Ozier-Lafontaine, A. Bonnaffoux, F. Picard, O. Gandrillon, and S. Gonin-Giraud. Evidence for close molecular proximity between reverting and undifferentiated cells. *BMC Biology*, 20(1):155, July 2022.
- [28] L. Zwald and G. Blanchard. On the Convergence of Eigenspaces in Kernel Principal Component Analysis. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2005.

A Operators on \mathcal{H} and on \mathcal{H}^k

In this section, we provide notation and we define properly all the objects involved in the definition of the TKHL statistic.

A.1 Operators on \mathcal{H}

Let $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\mathcal{H}})$ be a separable Hilbert space. A linear operator C from \mathcal{H} to \mathcal{H} is called Hilbert-Schmidt if $\sum_{i \geq 1} \|Cf_i\|_{\mathcal{H}}^2 < \infty$, where $(f_i)_{i \geq 1}$ is an orthonormal basis of \mathcal{H} , and in this case the sum is independent of the chosen orthonormal basis. The set of Hilbert-Schmidt operators $\text{HS}(\mathcal{H})$ endowed with the inner product

$$\langle C, T \rangle_{\text{HS}(\mathcal{H})} = \sum_{i \geq 1} \langle Cf_i, Tf_i \rangle_{\mathcal{H}} = \sum_{i, j \geq 1} \langle Cf_i, f_j \rangle_{\mathcal{H}} \langle Tf_i, f_j \rangle_{\mathcal{H}}$$

is also an Hilbert space.

A linear operator C on \mathcal{H} is called self-adjoint if $\langle Cg, h \rangle_{\mathcal{H}} = \langle g, Ch \rangle_{\mathcal{H}}$ for any $(g, h) \in \mathcal{H}^2$. A linear operator C on \mathcal{H} is called class-trace if $\sum_{i \geq 1} |\langle Cf_i, f_i \rangle_{\mathcal{H}}| < \infty$. In this case the sum is independent of the chosen orthonormal basis and the trace of C is defined as

$$\text{tr}(C) = \sum_{i \geq 1} \langle Cf_i, f_i \rangle_{\mathcal{H}}.$$

A finite rank operator is trace-class, and a trace-class operator is Hilbert-Schmidt.

For $(f, g) \in \mathcal{H} \setminus 0$, we consider the rank one operator $f \otimes g$ defined by

$$h \in \mathcal{H} \mapsto f \otimes g(h) = \langle g, h \rangle_{\mathcal{H}} f \in \mathcal{H},$$

and we have $\|f \otimes f\|_{\text{HS}(\mathcal{H})} = \|f\|_{\mathcal{H}}^2$ and $\|f \otimes g\|_{\text{HS}(\mathcal{H})} \leq \|f\|_{\mathcal{H}} \|g\|_{\mathcal{H}}$.

A.2 Norm and operators on \mathcal{H}^k

Let $k \in \mathbb{N}^*$. We first introduce a scalar product on \mathcal{H}^k in a natural way by

$$\langle \mathbf{g}, \mathbf{h} \rangle_{\mathcal{H}^k} := \sum_{j=1}^k \langle g_j, h_j \rangle_{\mathcal{H}}$$

for $\mathbf{g} = (g_1, \dots, g_k) \in \mathcal{H}^k$ and $\mathbf{h} = (h_1, \dots, h_k) \in \mathcal{H}^k$. We denote by $\|\cdot\|_{\mathcal{H}^k}$ the associated norm on \mathcal{H}^k .

We are now in position to define rigorously all the operators that appears in the HL statistic $\widehat{\mathcal{F}}_T$. A matrix $\mathbf{X} \in \mathcal{M}_{n,k}(\mathbb{R})$ is also seen as a linear application from \mathcal{H}^k to \mathcal{H}^n : for $\mathbf{g} \in \mathcal{H}^k$, $\mathbf{X}\mathbf{g} \in \mathcal{H}^n$ and

$$(\mathbf{X}\mathbf{g})_i = \sum_{j=1}^k x_{i,j} g_j \quad i = 1, \dots, n. \quad (7)$$

We next define rigorously the adjoint operator that appears in the HL statistic. For $\mathbf{g} \in \mathcal{H}^k$, we consider the associate application between the two Hilbert Spaces $\text{HS}(\mathcal{H})$ and \mathcal{H}^k defined by

$$\mathbf{g} : \mathbf{A} \in \text{HS}(\mathcal{H}) \mapsto \mathbf{g}\mathbf{A} = (\mathbf{A}g_1, \dots, \mathbf{A}g_k)' \in \mathcal{H}^k.$$

The application $u : (\mathbf{A}, \mathbf{h}) \in \text{HS}(\mathcal{H}) \times \mathcal{H}^k \mapsto \langle \mathbf{g}\mathbf{A}, \mathbf{h} \rangle_{\mathcal{H}^k}$ is clearly a sesquilinear form and it is also bounded because

$$\begin{aligned} |u(\mathbf{A}, \mathbf{h})| &\leq \|\mathbf{g}\mathbf{A}\|_{\mathcal{H}^k} \|\mathbf{h}\|_{\mathcal{H}^k} \\ &\leq \sum_{j=1}^k \|\mathbf{A}g_j\|_{\mathcal{H}} \|\mathbf{h}\|_{\mathcal{H}^k} \\ &\leq \sum_{j=1}^k \|\mathbf{A}\| \|g_j\|_{\mathcal{H}} \|\mathbf{h}\|_{\mathcal{H}^k} \\ &\leq \left(\sum_{j=1}^k \|g_j\|_{\mathcal{H}} \right) \|\mathbf{A}\|_{\text{HS}(\mathcal{H})} \|\mathbf{h}\|_{\mathcal{H}^k}. \end{aligned}$$

We can therefore consider the adjoint operator of this application (see for instance Theorem 2.2 in [4]), which is denoted by $\mathbf{g}^* : \mathcal{H}^k \mapsto \text{HS}(\mathcal{H})$ and we check that \mathbf{g}^* is actually the operator:

$$\mathbf{g}^* : \mathbf{h} \in \mathcal{H}^k \mapsto \mathbf{g}^* \mathbf{h} = \sum_{i=1}^k g_i \otimes h_i \in \text{HS}(\mathcal{H}). \quad (8)$$

Indeed, for any $(\mathbf{h}, \mathbf{A}) \in \mathcal{H}^k \times \text{HS}(\mathcal{H})$, on the one hand, by definition of $\mathbf{g}\mathbf{A}$ we have

$$\begin{aligned} \langle \mathbf{g}\mathbf{A}, \mathbf{h} \rangle_{\mathcal{H}^k} &= \sum_{i=1}^k \langle \mathbf{A}g_i, h_i \rangle_{\mathcal{H}} \\ &= \sum_{i=1}^k \sum_{j,\ell} \langle g_i, f_j \rangle_{\mathcal{H}} \langle h_i, f_\ell \rangle_{\mathcal{H}} \langle \mathbf{A}f_j, f_\ell \rangle_{\mathcal{H}}, \end{aligned}$$

where $(f_j)_{j \geq 1}$ is an orthonormal basis of \mathcal{H} so $g_i = \sum_{j \geq 1} \langle g_i, f_j \rangle_{\mathcal{H}} f_j$ and $h_i = \sum_{\ell \geq 1} \langle h_i, f_\ell \rangle_{\mathcal{H}} f_\ell$. On the other hand,

$$\begin{aligned} \langle \mathbf{A}, \sum_{i=1}^k g_i \otimes h_i \rangle_{\text{HS}(\mathcal{H})} &= \sum_{j,\ell} \langle \mathbf{A}f_j, f_\ell \rangle_{\mathcal{H}} \left\langle \sum_{i=1}^k g_i \otimes h_i f_j, f_\ell \right\rangle_{\mathcal{H}} \\ &= \sum_{j,\ell} \sum_{i=1}^k \langle \mathbf{A}f_j, f_\ell \rangle_{\mathcal{H}} \langle g_i f_j \rangle_{\mathcal{H}} \langle h_i, f_\ell \rangle_{\mathcal{H}}. \end{aligned}$$

and we have checked that $\langle \mathbf{g}\mathbf{A}, \mathbf{h} \rangle_{\mathcal{H}^k} = \langle \mathbf{A}, \mathbf{g}^* \mathbf{h} \rangle_{\text{HS}(\mathcal{H})}$.

With above definition, for $\mathbf{g}, \mathbf{h} \in \mathcal{H}^k$ and $\mathbf{M} = (m_{i,j})_{i,j \in \{1, \dots, k\}} \in \mathcal{M}_k(\mathbb{R})$, we then we have that:

$$\mathbf{g}^* \mathbf{M} \mathbf{g} = \sum_{i,j=1}^k m_{i,j} g_i \otimes g_j \in \text{HS}(\mathcal{H}), \quad (9)$$

We also check that the operator $\mathbf{g}\mathbf{h}^* : \mathcal{H}^k \mapsto \mathcal{H}^k$ is exactly the linear operator defined by the matrix (as in (7)):

$$\mathbf{g}\mathbf{h}^* := (\langle g_i, h_j \rangle_{\mathcal{H}})_{i,j \in \{1, \dots, k\}} \in \mathcal{M}_k(\mathbb{R}).$$

For $v = (v_1, \dots, v_k) \in \mathbb{R}^k$ and $\mathbf{g} \in \mathcal{H}^k$, we will abuse the notation of matrix multiplication by defining $v'\mathbf{g}$ for the linear combination of the elements of \mathbf{g} with weights v :

$$v'\mathbf{g} = \sum_{i=1}^k v_i g_i \in \mathcal{H}.$$

For any $v, w \in \mathbb{R}^k$, we then have that

$$\langle v'\mathbf{g}, w'\mathbf{h} \rangle_{\mathcal{H}} = v'(\mathbf{g}\mathbf{h}^*)w.$$

For $\mathbf{g} \in \mathcal{H}^k$, we finally define the linear application

$$h \in \mathcal{H} \mapsto \mathbf{g}h = (\langle g_1, h \rangle_{\mathcal{H}}, \dots, \langle g_k, h \rangle_{\mathcal{H}})' \in \mathbb{R}^k.$$

B Linear model in the feature space \mathcal{H}

B.1 Least squares estimation on the linear model in the feature space

We check that the estimator given in (2) is the least squares estimator of Θ in the model (1) for this norm on \mathcal{H}^n . The Hilbert space \mathcal{H} being separable, let us consider $(g_s)_{s \geq 1}$ an orthonormal basis of \mathcal{H} . For $s \geq 1$ and $h \in \mathcal{H}$, we define $h^s = \langle h, g_s \rangle_{\mathcal{H}}$ such that $h = \sum_{s \geq 1} h^s g_s$. From the definition of the norm $\| \cdot \|_{\mathcal{H}^n}$, we easily find that the least squares problem in \mathcal{H}^n reduces to a collection of

independent and classical univariate least squares problems: for $\Theta = (\Theta_0, \dots, \Theta_{p-1})' \in \mathcal{H}^p$ where $\Theta_j = \sum_{s \geq 1} \theta_j^s g_s$, then

$$\begin{aligned} \|\Phi(\mathbf{Y}) - \mathbf{X}\Theta\|_{\mathcal{H}^n}^2 &= \sum_{i=1 \dots n} \|\phi(Y_j) - x'_i \Theta\|_{\mathcal{H}}^2 \\ &= \sum_{i=1 \dots n} \sum_{s \geq 1} \left| \phi(Y_j)^s - \sum_{j=0 \dots p-1} x_{i,j} \theta_j^s \right|^2 \\ &= \sum_{s \geq 1} \left\{ \sum_{i=1 \dots n} \left| \phi(Y_j)^s - \sum_{j=0 \dots p-1} x_{i,j} \theta_j^s \right|^2 \right\} \\ &= \sum_{s \geq 1} \|\Phi(\mathbf{Y})^s - \mathbf{X}\Theta^s\|_{\mathbb{R}^n}^2. \end{aligned}$$

As \mathbf{X} is assumed to be full rank, the least squares estimator of $\Theta^s = (\theta_0^s, \dots, \theta_{p-1}^s)' \in \mathbb{R}^p$ is given by $\hat{\Theta}^s = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\Phi(\mathbf{Y})^s$. From this, we deduce that for the least square estimation of the linear model (1), we have

$$\hat{\Theta} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\Phi(\mathbf{Y}).$$

Moreover, the predictions in the linear model (1) are defined as $\widehat{\Phi(\mathbf{Y})} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\Phi(\mathbf{Y}) = \mathbf{P}_\mathbf{X}\Phi(\mathbf{Y})$ where $\mathbf{P}_\mathbf{X} = (\pi_{i,j})_{1 \leq i,j \leq n}$ is the (matrix of the) orthogonal projection on $Im(\mathbf{X})$ in $(\mathbb{R}^n, \|\cdot\|_2)$. We also define the residuals as $\hat{\mathbf{E}} = (\hat{e}_1, \dots, \hat{e}_n)$ satisfy $\hat{\mathbf{E}} = \mathbf{P}_\mathbf{X}^\perp \Phi(\mathbf{Y}) = \mathbf{P}_\mathbf{X}^\perp \mathbf{E} \in \mathcal{H}^n$, where $\mathbf{P}_\mathbf{X}^\perp = \mathbf{I}_n - \mathbf{P}_\mathbf{X} = (\pi_{i,j}^\perp)_{1 \leq i,j \leq n}$ is the orthogonal projection on $Im(\mathbf{X})^\perp$.

According to the following lemma, we check that the matrices $\mathbf{P}_\mathbf{X}$ and $\mathbf{P}_\mathbf{X}^\perp$ can be seen as orthogonal matrices in \mathcal{H}^n .

Lemma 1. *Let \mathbf{P} an orthogonal projection in $(\mathbb{R}^n, \langle \cdot, \cdot \rangle_2)$. Then \mathbf{P} is also an orthogonal projection in $(\mathcal{H}^n, \langle \cdot, \cdot \rangle_{\mathcal{H}^n})$. In particular, $\|\mathbf{P}\mathbf{f}\|_{\mathcal{H}^n} \leq \|\mathbf{f}\|_{\mathcal{H}^n}$ for any $\mathbf{f} \in \mathcal{H}^n$.*

Proof. We easily check that for any $\mathbf{f} \in \mathcal{H}^n$, $\mathbf{P}(\mathbf{P}\mathbf{f}) = (\mathbf{P}\mathbf{P})\mathbf{f} = \mathbf{P}\mathbf{f}$. Thus \mathbf{P} is also a projection on \mathcal{H}^n . Next, it can be easily checked that for any $(v, w) \in (\mathbb{R}^n)^2$, and $(\mathbf{f}, \mathbf{g}) \in (\mathcal{H}^n)^2$,

$$\langle (v'v)\mathbf{f}, (w'w)\mathbf{g} \rangle_{\mathcal{H}^n} = \langle v'\mathbf{f}, w'\mathbf{g} \rangle_{\mathcal{H}} \langle v, w \rangle_{\mathbb{R}^n}. \quad (10)$$

and

$$\langle (v'v)\mathbf{f}, \mathbf{g} \rangle_{\mathcal{H}^n} = \langle \mathbf{f}, (v'v)\mathbf{g} \rangle_{\mathcal{H}^n}. \quad (11)$$

since \mathbf{P} is an orthogonal projection in $(\mathbb{R}^n, \|\cdot\|_2)$, we can write $\mathbf{P} = \sum_{i=1}^r v_i v_i'$ where $(v_i)_{i=1 \dots r}$ is an orthonormal basis of $Im(\mathbf{P})$ in \mathbb{R}^n . From (10) and (11), we easily find that $\langle \mathbf{P}\mathbf{f}, \mathbf{g} \rangle_{\mathcal{H}^n} = \langle \mathbf{f}, \mathbf{P}\mathbf{g} \rangle_{\mathcal{H}^n}$. \square

B.2 The KHL and TKHL statistics are well defined

The test operator associated to \mathbf{L} is defined as

$$\hat{H}_\mathbf{L} = (\mathbf{L}\hat{\Theta})^* (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{L}')^{-1} (\mathbf{L}\hat{\Theta}),$$

where $\mathbf{L} = (L_\ell^j)_{\ell \in \{1, \dots, d\}, j \in \{1, \dots, p\}} \in \mathbb{R}^{d \times p}$ is a surjective matrix.

Note that that $\hat{H}_\mathbf{L}$ is a well-defined Hilbert-Schmidt operator. Indeed, $\hat{\Theta} \in \mathcal{H}^n$, and according to (7), $\mathbf{L}\hat{\Theta} \in \mathcal{H}^d$. Next, the matrix $\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{L}'$ is an $d \times d$ invertible matrix because \mathbf{L} is surjective. According to (9), $\hat{H}_\mathbf{L} \in HS(\mathcal{H})$ and can be written as

$$\hat{H}_\mathbf{L} = \sum_{i,j=1}^d c_{i,j} (\mathbf{L}\hat{\Theta})_i \otimes (\mathbf{L}\hat{\Theta})_j,$$

where $C = (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{L}')^{-1}$. It is also a finite rank and class-trace operator.

The residual covariance operator $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \widehat{e}_i \otimes \widehat{e}_i$ is clearly a bounded and finite rank operator, it is thus compact. It is also self-adjoint and thus according to the spectral theorem there exists an orthonormal basis of \mathcal{H} consisting of eigenvectors of $\widehat{\Sigma}$. It can then be written as

$$\widehat{\Sigma} = \sum_{t=1}^r \widehat{\lambda}_t (\widehat{f}_t \otimes \widehat{f}_t)$$

with $\widehat{\lambda}_t \neq 0$ and where r is the rank of $\widehat{\Sigma}$. Moreover $\widehat{\Sigma}$ is obviously positive semi-definite, thus $\widehat{\lambda}_t > 0$ for any non null eigenvalue, and we will assume that they are indexed in decreasing order. We then introduce a generalized inverse for $\widehat{\Sigma}$ as

$$\widehat{\Sigma}^- = \sum_{t=1}^r \frac{1}{\widehat{\lambda}_t} (\widehat{f}_t \otimes \widehat{f}_t),$$

which is also a positive semi-definite bounded and finite rank self-adjoint operator on \mathcal{H} .

The operator $\widehat{\Sigma}^- \widehat{H}_{\mathbf{L}}$ is thus trace-class as $\widehat{\Sigma}^-$ is bounded and $\widehat{H}_{\mathbf{L}}$ is trace-class. We can thus define the kernelized Hotelling-Lawley statistic as

$$\widehat{\mathcal{F}} = \text{tr}(n^{-1} \widehat{\Sigma}^- \widehat{H}_{\mathbf{L}}),$$

where tr is the trace operator on \mathcal{H} . The spectral regularization of parameter $T \leq r$ is defined as:

$$\widehat{\Sigma}_T^{-1} = \sum_{t=1}^T \frac{1}{\widehat{\lambda}_t} (\widehat{f}_t \otimes \widehat{f}_t). \quad (12)$$

It is also a positive semi-definite bounded and finite rank self-adjoint operator on \mathcal{H} . The Truncated kernel Hotelling-Lawley (TKHL) trace statistic can thus be finally defined as

$$\widehat{\mathcal{F}}_T = \text{tr} \left(\frac{1}{n} \widehat{\Sigma}_T^{-1} \widehat{H}_{\mathbf{L}} \right).$$

C Proof of Theorem 1

This section is dedicated to the proof of Theorem 1. The first subsection gives a summary of the proof. We then provide the main tools for the proof in Propositions 1, 2 and 3.

In this section, we consider the linear model (1) in the feature space \mathcal{H} .

C.1 Main proof of Theorem 1

First, we define an alternative truncated Hotelling-Lawley trace based on the true covariance structure shared between the errors Σ_T instead of $\widehat{\Sigma}_T$, such that $\widetilde{\mathcal{F}}_T = \text{tr}(\Sigma_T^{-1} \widehat{H}_{\mathbf{L}})$, where $\Sigma_T^{-1} = \sum_{t=1}^T \lambda_t^{-1} (f_t \otimes f_t)$ is well defined according to Assumption 2. We have :

$$\begin{aligned} |\widehat{\mathcal{F}}_T - \widetilde{\mathcal{F}}_T| &= \left| \text{tr} \left(\frac{1}{n} (\widehat{\Sigma}_T^{-1} - \Sigma_T^{-1}) \widehat{H}_{\mathbf{L}} \right) \right| \\ &= \left| \left\langle (\widehat{\Sigma}_T^{-1} - \Sigma_T^{-1}), \frac{1}{n} \widehat{H}_{\mathbf{L}} \right\rangle_{\text{HS}(\mathcal{H})} \right| \\ &\leq \left\| \widehat{\Sigma}_T^{-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})} \left\| \frac{1}{n} \widehat{H}_{\mathbf{L}} \right\|_{\text{HS}(\mathcal{H})}. \end{aligned}$$

Let $\mathbf{P}_{\mathbf{X}}^\perp = \mathbf{I}_n - \mathbf{P}_{\mathbf{X}} = (\pi_{i,j}^\perp)_{1 \leq i,j \leq n}$ the orthogonal projection on $\text{Im}(\mathbf{X})^\perp$. We define $\bar{\Sigma}_n = \Sigma_T + \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \Sigma_{i,T^c}$. By applying a bounded differences Theorem, under Assumptions 1, 3 and 4, Proposition 2 provides an exponential bound on $\left\| \widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})}$. Following the approach of

[28], we can then we make use of results from operator perturbation theory to derive the exponential bound on $\left\| \widehat{\Sigma}_T^{-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})}$ given in Proposition 3.

We then show in Lemma 2 that $n^{-1} \widehat{H}_L$ is bounded under Assumptions 1 and 3, and we can conclude that $\widehat{\mathcal{F}}_T \xrightarrow{P} \widetilde{\mathcal{F}}_T$ as n tends to infinity. It only remains to derive the asymptotic distribution of $\widetilde{\mathcal{F}}_T$.

Proposition 1 shows that $\widetilde{\mathcal{F}}_T$ corresponds to a standard Hotelling-Lawley Statistic in \mathbb{R}^T . We then invoke a result from the literature (see for instance Theorem 12.8 from [19]) to find that if H_0 is true, then $n \widetilde{\mathcal{F}}_T \xrightarrow{D} \chi_{dT}^2$ as n tends to infinity.

C.2 Statistics for the test operator

Proposition 1. *Under Assumptions 1, 3, 4, 5 and 6, if H_0 is true then we have:*

$$n \widetilde{\mathcal{F}}_T = \text{tr}(\Sigma_T^{-1} \widehat{H}_L) \xrightarrow[n \rightarrow \infty]{D} \chi_{dT}^2.$$

Proof. We recall that the i^{th} equation of the linear model (1) can be written as

$$\phi(Y_i) = x_i' \Theta + e_i,$$

where $\mathbb{E}e_i = 0$. According to Assumption 1, we have $\text{Cov}(e_i) = \Sigma_i = \Sigma_T + \Sigma_{i,T^c}$, and $\Sigma_T = \sum_{t=1}^T \lambda_t f_t \otimes f_t$ where f_1, \dots, f_T is an orthonormal set of eigenfunctions of Σ_T associated to the eigenvalues $\lambda_1, \dots, \lambda_T$. For $t \in \{1, \dots, T\}$ and $h \in \mathcal{H}$, we define $h^t = \langle h, f_t \rangle_{\mathcal{H}}$. For $i \in \{1, \dots, n\}$, let $Z_i = (\phi(Y_i)^1, \dots, \phi(Y_i)^T)$ in \mathbb{R}^T , $\beta_i = (\theta_i^1, \dots, \theta_i^T)$ in \mathbb{R}^T and $\beta = (\beta_0, \dots, \beta_{p-1})' \in \mathcal{M}_{p,T}(\mathbb{R})$. We also define $\varepsilon_i = (e_i^1, \dots, e_i^T)$ in \mathbb{R}^T . The projection of the i^{th} equation of the linear model (1) on $\text{Span}(f_1, \dots, f_T) \subset \mathcal{H}$ is:

$$Z_i = x_i' \beta + \varepsilon_i.$$

We recognize a multivariate linear model in \mathbb{R}^T that has the matrix form:

$$\mathbf{Z} = X\beta + \varepsilon, \tag{13}$$

where $\mathbf{Z} = (Z_1, \dots, Z_n)' \in \mathcal{M}_{n,T}(\mathbb{R})$ and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)' \in \mathcal{M}_{n,T}(\mathbb{R})$. Remark that the errors $\varepsilon_1, \dots, \varepsilon_n$ are independant with $\mathbb{E}\varepsilon_1 = 0$ and same covariance matrix $\mathbf{S}_T = (\sigma_{t,t'})_{t,t' \in \{1, \dots, T\}} \in \mathcal{M}_T(\mathbb{R})$. For $i \in \{1, \dots, n\}$ and $t, t' \in \{1, \dots, T\}$, we have:

$$\begin{aligned} \sigma_{t,t'} &= \mathbb{E}(e_i^t e_i^{t'}) \\ &= \mathbb{E} \langle e_i, f_t \rangle_{\mathcal{H}} \langle e_i, f_{t'} \rangle_{\mathcal{H}} \\ &= \langle f_t, \mathbb{E}(e_i \otimes e_i) f_{t'} \rangle_{\mathcal{H}} \\ &= \langle f_t, \Sigma_i f_{t'} \rangle_{\mathcal{H}}. \end{aligned}$$

By construction, as $t' \in \{1, \dots, T\}$, $f_{t'}$ is an eigenvector of Σ_i associated to the eigenvalue $\lambda_{t'}$, thus according to Assumption 1,

$$\begin{aligned} \sigma_{t,t'} &= \lambda_{t'} \langle f_t, f_{t'} \rangle_{\mathcal{H}} \\ &= \begin{cases} \lambda_t & \text{if } t = t' \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Thus, we have $\mathbf{S}_T = \text{diag}(\lambda_1, \dots, \lambda_T)$ where $\lambda_j > 0$, $j = 1, \dots, T$.

In the multivariate linear model (13), the least squares estimator $\widehat{\beta} = (\widehat{\beta}_0, \dots, \widehat{\beta}_{p-1})' \in \mathcal{M}_{p,T}(\mathbb{R})$ of β is defined by

$$\widehat{\beta} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{Z}.$$

By developing \mathbf{Z} in this equation, it can be easily checked that that for $j \in \{1, \dots, p\}$, $\widehat{\beta}_j = (\widehat{\theta}_j^1, \dots, \widehat{\theta}_j^T)' \in \mathbb{R}^T$. Moreover, the covariance matrix \mathbf{S}_T , which is by definition diagonal can be estimated by $\widehat{\mathbf{S}}_T = \text{diag}(\widehat{\lambda}_1, \dots, \widehat{\lambda}_T)$ where $\widehat{\lambda}_j$ is the empirical variance of the $Z_{i,j}$'s.

In this multivariate linear model (13), the Hotelling-Lawley trace associated to the hypotheses $\tilde{H}_0 : \mathbf{L}\boldsymbol{\beta} = 0_{\mathbb{R}^d}$ versus $\tilde{H}_1 : \mathbf{L}\boldsymbol{\beta} \neq 0_{\mathbb{R}^d}$ is defined by

$$\mathcal{G}_n = \text{trace} \left(\frac{1}{n} \widehat{\mathbf{S}}_T^{-1} \widehat{\mathbf{H}}_{\mathbf{L},T} \right),$$

where $\widehat{\mathbf{H}}_{\mathbf{L},T} = (\mathbf{L}\widehat{\boldsymbol{\beta}})'(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')^{-1}(\mathbf{L}\widehat{\boldsymbol{\beta}}) \in \mathcal{M}_T(\mathbb{R})$. Note that in \mathcal{G}_n we use the trace operator trace defined on $\mathcal{M}_T(\mathbb{R})$, to distinguish it from the trace tr defined on the HS operators on \mathcal{H} .

We now check that the assumptions of Theorem 12.8 from [19] are satisfied (in particular Assumption D1 in this theorem). First, Assumption 1 provides the homoscedasticity assumption in \mathbb{R}^T . Next, Assumptions 3 and 4 are sufficient to guarantee that the errors are bounded (see Lemma 4) and thus the ε_i have obviously a finite fourth moment. Finally, Assumption 5 provides the control of the coefficient on the diagonal of \mathbf{P}_X , and the convergence of $\frac{1}{n}X'X$ is provided by Assumption 6. All this gives that if H_0 is true, then

$$n\mathcal{G}_n \xrightarrow[n \rightarrow \infty]{\mathcal{D}} \chi_{dT}^2.$$

In particular, in the proof of the above result, it is also shown that

$$n\tilde{\mathcal{G}}_n = \text{trace}(\mathbf{S}_T^{-1} \widehat{\mathbf{H}}_{\mathbf{L},T}) \xrightarrow[n \rightarrow \infty]{\mathcal{D}} \chi_{dT}^2.$$

Thus, it remains to check that $\tilde{\mathcal{F}}_T = \tilde{\mathcal{G}}_n$ to complete the proof.

We denote $C = (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')^{-1} = (c_{i,j})_{i,j \in \{1, \dots, p\}} \in \mathcal{M}_p(\mathbb{R})$, and then $\widehat{\mathbf{H}}_{\mathbf{L},T} = (\mathbf{L}\widehat{\boldsymbol{\beta}})'C(\mathbf{L}\widehat{\boldsymbol{\beta}}) = \sum_{i,j=1}^p c_{i,j}(\mathbf{L}\widehat{\boldsymbol{\beta}})_i(\mathbf{L}\widehat{\boldsymbol{\beta}})'_j$. For $i, j \in \{0, \dots, p-1\}$, we have:

$$\mathbf{S}_T^{-1}(\mathbf{L}\widehat{\boldsymbol{\beta}})_i(\mathbf{L}\widehat{\boldsymbol{\beta}})'_j = \begin{pmatrix} \lambda_1^{-1}(\mathbf{L}\widehat{\boldsymbol{\beta}})_i^1(\mathbf{L}\widehat{\boldsymbol{\beta}})'_j^1 & \dots & \lambda_1^{-1}(\mathbf{L}\widehat{\boldsymbol{\beta}})_i^1(\mathbf{L}\widehat{\boldsymbol{\beta}})'_j^T \\ \vdots & \ddots & \vdots \\ \lambda_T^{-1}(\mathbf{L}\widehat{\boldsymbol{\beta}})_i^T(\mathbf{L}\widehat{\boldsymbol{\beta}})'_j^1 & \dots & \lambda_T^{-1}(\mathbf{L}\widehat{\boldsymbol{\beta}})_i^T(\mathbf{L}\widehat{\boldsymbol{\beta}})'_j^T \end{pmatrix}.$$

Then $n\tilde{\mathcal{G}}_n = \text{trace}(\mathbf{S}_T^{-1} \widehat{\mathbf{H}}_{\mathbf{L},T}) = \sum_{i,j=0}^{p-1} c_{i,j} \sum_{t=1}^T \lambda_t^{-1} \hat{\theta}_i^t \hat{\theta}_j^t$. On the other hand, we have that $n\tilde{\mathcal{F}}_T = \text{tr}(\Sigma_T^{-1} \widehat{H}_{\mathbf{L}}) = \langle \widehat{H}_{\mathbf{L}}, \Sigma_T^{-1} \rangle_{\text{HS}(\mathcal{H})}$ because $\Sigma_T^{-1} = \sum_{t=1}^T \hat{\lambda}_t^{-1} \hat{f}_t \otimes \hat{f}_t$ is self-adjoint and finite rank. Thus,

$$\begin{aligned} n\tilde{\mathcal{F}}_T &= \sum_{i,j=0}^{p-1} \sum_{t=1}^T c_{i,j} \lambda_t^{-1} \langle f_t \otimes f_t, (\mathbf{L}\widehat{\boldsymbol{\Theta}})_i \otimes (\mathbf{L}\widehat{\boldsymbol{\Theta}})_j \rangle_{\text{HS}(\mathcal{H})} \\ &= \sum_{i,j=0}^{p-1} \sum_{t=1}^T c_{i,j} \lambda_t^{-1} \langle f_t, (\mathbf{L}\widehat{\boldsymbol{\Theta}})_i \rangle_{\mathcal{H}} \langle f_t, (\mathbf{L}\widehat{\boldsymbol{\Theta}})_j \rangle_{\mathcal{H}} \\ &= \sum_{i,j=0}^{p-1} \sum_{t=1}^T c_{i,j} \lambda_t^{-1} (\mathbf{L}\widehat{\boldsymbol{\Theta}})_i^t (\mathbf{L}\widehat{\boldsymbol{\Theta}})_j^t \\ &= n\tilde{\mathcal{G}}_n. \end{aligned}$$

□

Lemma 2. Under Assumptions 3 and 4, if H_0 is true, then we have:

$$\left\| \frac{1}{n} \widehat{H}_{\mathbf{L}} \right\|_{\text{HS}(\mathcal{H})} \leq M_e^2,$$

where where $M_e = \sqrt{M_k} + pM_x \max_{j \in \{1, \dots, p\}} \|\theta_j\|_{\mathcal{H}}$.

Proof. We inject the expression of $\widehat{\boldsymbol{\Theta}}$ in the expression of $\widehat{H}_{\mathbf{L}}$:

$$\widehat{H}_{\mathbf{L}} = (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\Phi(\mathbf{Y}))^*(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')^{-1}(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\Phi(\mathbf{Y}))$$

According to the linear model (1), we have

$$\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\Phi(\mathbf{Y}) = \mathbf{L}\Theta + \mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{E}.$$

The term $\mathbf{L}\Theta$ is null under H_0 . Then we can write:

$$\widehat{H}_{\mathbf{L}} = \mathbf{E}^*\mathbf{D}\mathbf{E}$$

where $\mathbf{D} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}'(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')^{-1}\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' = (d_{i,j})_{i,j \in \{1, \dots, n\}} \in \mathcal{M}_n(\mathbb{R})$ is an orthogonal projector (of \mathbb{R}^n), as $\mathbf{D}' = \mathbf{D}$ and $\mathbf{D}^2 = \mathbf{D}$. Thus we can write $\widehat{H}_{\mathbf{L}} = \mathbf{E}^*\mathbf{D}'\mathbf{D}\mathbf{E}$. Moreover, it can be easily checked that $(\mathbf{D}\mathbf{E})^* = \mathbf{E}^*\mathbf{D}'$. Thus we have:

$$\begin{aligned} \left\| \widehat{H}_{\mathbf{L}} \right\|_{\text{HS}(\mathcal{H})} &= \|(\mathbf{D}\mathbf{E})^*(\mathbf{D}\mathbf{E})\|_{\text{HS}(\mathcal{H})} \\ &\leq \sum_{i=1}^n \|(\mathbf{D}\mathbf{E})_i \otimes (\mathbf{D}\mathbf{E})_i\|_{\text{HS}(\mathcal{H})} \\ &\leq \sum_{i=1}^n \|(\mathbf{D}\mathbf{E})_i\|_{\mathcal{H}}^2 = \|\mathbf{D}\mathbf{E}\|_{\mathcal{H}^n}^2. \end{aligned}$$

By Lemma 1, we have $\|\mathbf{D}\mathbf{E}\|_{\mathcal{H}^n} \leq \|\mathbf{E}\|_{\mathcal{H}^n}$, thus, according to Lemma 4, we have :

$$\begin{aligned} \left\| \widehat{H}_{\mathbf{L}} \right\|_{\text{HS}(\mathcal{H})} &\leq \|\mathbf{E}\|_{\mathcal{H}^n}^2 \\ &\leq \sum_{i=1}^n \|e_i\|_{\mathcal{H}}^2 \\ &\leq nM_e^2. \end{aligned}$$

Thus we have $\left\| n^{-1}\widehat{H}_{\mathbf{L}} \right\|_{\text{HS}(\mathcal{H})} \leq M_e^2$. □

C.3 Probability bound on the residual covariance

We define $\bar{\Sigma}_n = \Sigma_T + \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^{\perp} \Sigma_{i,T^c}$. Next Proposition is adapted from [22] and [28]. An obvious consequence of Proposition 2 is that $\widehat{\Sigma} - \bar{\Sigma}_n$ tends to 0 in probability as n tends to infinity.

Proposition 2. *If Assumptions 1, 3 and 4 are satisfied, then we have with probability $1 - e^{-\xi}$:*

$$\left\| \widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \leq \frac{p}{n} \|\Sigma_T\|_{\text{HS}(\mathcal{H})} + \frac{M_e^2}{\sqrt{n}} \left(1 + 4\sqrt{\frac{3\xi}{2}} \right) =: \zeta(n, \xi), \quad (14)$$

where $M_e = M_k^{\frac{1}{2}} + M_x \max_{j \in \{1, \dots, p\}} \|\theta_j\|_{\mathcal{H}}$.

Proof. We observe that:

$$\begin{aligned} \left\| \widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} &= \left\| \widehat{\Sigma} - \mathbb{E}\widehat{\Sigma} + \mathbb{E}\widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \\ &\leq \left\| \widehat{\Sigma} - \mathbb{E}\widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})} + \left\| \mathbb{E}\widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})}. \end{aligned}$$

We know from Lemma 3 that:

$$\begin{aligned} \left\| \mathbb{E}\widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} &= \left\| \frac{n-p}{n} \Sigma_T - \Sigma_T \right\|_{\text{HS}(\mathcal{H})} \\ &= \frac{p}{n} \|\Sigma_T\|_{\text{HS}(\mathcal{H})}. \end{aligned}$$

To apply the McDiarmid Inequality (see Theorem 3), we introduce to the function:

$$\begin{aligned} \mathcal{H}^n &\rightarrow \mathbb{R} \\ f : \mathbf{E} &\mapsto \left\| \widehat{\Sigma} - \mathbb{E}(\widehat{\Sigma}) \right\|_{\text{HS}(\mathcal{H})}. \end{aligned}$$

Let $i_0 \in \{1, \dots, n\}$, $\mathbf{E} = (e_1, \dots, e_n)$ and $\tilde{\mathbf{E}}_{i_0} = (\tilde{e}_1, \dots, \tilde{e}_n)$ in \mathcal{H}^n such that $e_{i_0} \neq \tilde{e}_{i_0}$ and $\forall i \neq i_0, e_i = \tilde{e}_i$. Lemma 5 gives us that:

$$|f(\mathbf{E}) - f(\tilde{\mathbf{E}}_{i_0})| \leq \frac{2M_e}{n} \left(M_e + \|(\mathbf{P}_X \mathbf{E})_{i_0}\|_{\mathcal{H}} + \|(\mathbf{P}_X \tilde{\mathbf{E}})_{i_0}\|_{\mathcal{H}} \right) =: c_{i_0}$$

and thus

$$\sum_{i_0=1}^n c_{i_0}^2 \leq \frac{16M_e^4}{n} + \frac{16M_e^2}{n^2} \sum_{i_0} \|(\mathbf{P}_X \mathbf{E})_{i_0}\|_{\mathcal{H}}^2 + \frac{16M_e^2}{n^2} \|(\mathbf{P}_X \tilde{\mathbf{E}})_{i_0}\|_{\mathcal{H}}^2 \quad (15)$$

$$\leq \frac{16M_e^4}{n} + \frac{16M_e^2}{n^2} \|\mathbf{P}_X \mathbf{E}\|_{\mathcal{H}_n}^2 + \frac{16M_e^2}{n^2} \|\mathbf{P}_X \tilde{\mathbf{E}}\|_{\mathcal{H}_n}^2 \quad (16)$$

According to Lemmas 1 and 4, $\|\mathbf{P}_X \mathbf{E}\|_{\mathcal{H}_n} \leq \|\mathbf{E}\|_{\mathcal{H}_n} \leq \sqrt{n}M_e$ and finally $\sum_{i_0=1}^n c_{i_0}^2 \leq \frac{48M_e^4}{n}$.

According to McDiarmid Inequality to f , we have with probability $1 - e^{-\xi}$:

$$\left\| \hat{\Sigma} - \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})} \leq \mathbb{E} \left\| \hat{\Sigma} - \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})} + 4M_e^2 \sqrt{\frac{3\xi}{2n}}.$$

By Lemma 6, we bound the expectation term to obtain that with probability $1 - e^{-\xi}$:

$$\begin{aligned} \left\| \hat{\Sigma} - \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})} &\leq \frac{M_e^2 \sqrt{n-p}}{n} + 4M_e^2 \sqrt{\frac{3\xi}{2n}} \\ &\leq \frac{M_e^2}{\sqrt{n}} \left(1 + 4\sqrt{\frac{3\xi}{2}} \right). \end{aligned}$$

□

Lemma 3. *Under Assumption 1, we have that:*

$$\mathbb{E}(\hat{\Sigma}) = \frac{n-p}{n} \Sigma_T + \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \Sigma_{i,T^c}.$$

Proof. The i -th residual can be written as $\hat{e}_i = \sum_{j=1}^n \pi_{i,j}^\perp e_j$. where $\pi_i^\perp = (\pi_{i,1}^\perp, \dots, \pi_{i,n}^\perp)' \in \mathbb{R}^n$ is the i^{th} column of \mathbf{P}_X^\perp . Since \mathbf{P}_X^\perp is an orthogonal projector, we have

$$\forall j \in \{1, \dots, n\}, \quad \sum_{i=1}^n \pi_{i,j}^{\perp 2} = \pi_{j,j}^\perp. \quad (17)$$

Consequently, the expectation is such that:

$$\begin{aligned} \mathbb{E}(\hat{\Sigma}) &= \frac{1}{n} \mathbb{E} \left(\sum_{i=1}^n \hat{e}_i \otimes \hat{e}_i \right) \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \pi_{i,j}^\perp \pi_{i,k}^\perp \mathbb{E}(e_j \otimes e_k) \\ &= \frac{1}{n} \sum_{i=1}^n \left(\sum_{j=1}^n \pi_{i,j}^{\perp 2} \mathbb{E}(e_j \otimes e_j) + \sum_{j \neq k} \pi_{i,j}^\perp \pi_{i,k}^\perp \mathbb{E}(e_j) \otimes \mathbb{E}(e_k) \right). \end{aligned}$$

Then, for $j \in \{1, \dots, n\}$, $\mathbb{E}(e_j) = 0$. We use Assumption 1 to obtain :

$$\begin{aligned} \mathbb{E}(\hat{\Sigma}) &= \frac{1}{n} \sum_{i,j=1}^n \pi_{i,j}^{\perp 2} (\Sigma_T + \Sigma_{j,T^c}) \\ &= \frac{n-p}{n} \Sigma_T + \frac{1}{n} \sum_{j=1}^n \pi_{j,j}^\perp \Sigma_{j,T^c} \end{aligned}$$

where we apply Lemma 8 and Equation (17) to obtain the last equality. □

Lemma 4. *If Assumptions 3 and 4 are satisfied, then for $i \in \{1, \dots, n\}$,*

$$\|e_i\|_{\mathcal{H}} < M_e := \sqrt{M_k} + pM_x \max_{j \in \{1, \dots, p\}} \|\theta_j\|_{\mathcal{H}}.$$

Proof. We have:

$$\begin{aligned} \|e_i\|_{\mathcal{H}} &= \|\phi(Y_i) - (\mathbf{X}\Theta)_i\|_{\mathcal{H}} \\ &\leq \|\phi(Y_i)\|_{\mathcal{H}} + \|(\mathbf{X}\Theta)_i\|_{\mathcal{H}} \end{aligned}$$

From Assumption 3, we know that $\|\phi(Y_i)\|_{\mathcal{H}} \leq \sqrt{M_k}$, and from the triangular inequality, we have:

$$\begin{aligned} \|(\mathbf{X}\Theta)_i\|_{\mathcal{H}} &= \left\| \sum_{j=0}^{p-1} \theta_j x_i^j \right\|_{\mathcal{H}} \\ &\leq pM_x \max_{j \in \{0, \dots, p-1\}} \|\theta_j\|_{\mathcal{H}} \end{aligned}$$

and the sum of the two bounds gives the result. \square

Lemma 5. *Let \mathbf{E} , $\tilde{\mathbf{E}}_{i_0}$ and f defined as in the proof of Proposition 2. Under Assumptions 3 and 4 we have:*

$$|f(\mathbf{E}) - f(\tilde{\mathbf{E}}_{i_0})| \leq \frac{2M_e}{n} \left(M_e + \|(\mathbf{P}_X \mathbf{E})_{i_0}\|_{\mathcal{H}} + \|(\mathbf{P}_X \tilde{\mathbf{E}})_{i_0}\|_{\mathcal{H}} \right),$$

where M_e is defined in Lemma 4.

Proof. Denote $\tilde{\Sigma} = \frac{1}{n} \sum_{i=1}^n \tilde{e}_i \otimes \tilde{e}_i$, we have:

$$\begin{aligned} |f(\mathbf{E}) - f(\tilde{\mathbf{E}}_{i_0})| &= \left| \left\| \hat{\Sigma} - \mathbb{E}(\hat{\Sigma}) \right\|_{\text{HS}(\mathcal{H})} - \left\| \tilde{\Sigma} - \mathbb{E}(\tilde{\Sigma}) \right\|_{\text{HS}(\mathcal{H})} \right| \\ &\leq \left\| \hat{\Sigma} - \tilde{\Sigma} \right\|_{\text{HS}(\mathcal{H})}. \end{aligned}$$

As $\hat{e}_i = e_i - \pi'_i \mathbf{E}$ where π_i is the i -th column (or row) of \mathbf{P}_X , we have:

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n e_i \otimes e_i - \frac{1}{n} \sum_{i=1}^n e_i \otimes (\pi'_i \mathbf{E}) - \frac{1}{n} \sum_{i=1}^n (\pi'_i \mathbf{E}) \otimes e_i + \frac{1}{n} \sum_{i=1}^n (\pi'_i \mathbf{E}) \otimes (\pi'_i \mathbf{E}).$$

As \mathbf{P}_X is an orthogonal projector, we have that $\mathbf{P}_X = \mathbf{P}'_X = \mathbf{P}^2_X$. Thus for $i, j \in \{1, \dots, n\}$, we have $\pi_{i,j} = \pi_{j,i}$ and $\sum_{k=1}^n \pi_{i,k} \pi_{j,k} = \pi_{i,j}$. Then by developing each term $\pi'_i \mathbf{E} = \sum_{j=1}^n \pi_{i,j} e_j$, we obtain that:

$$\frac{1}{n} \sum_{i=1}^n e_i \otimes (\pi'_i \mathbf{E}) = \frac{1}{n} \sum_{i=1}^n (\pi'_i \mathbf{E}) \otimes e_i = \frac{1}{n} \sum_{i=1}^n (\pi'_i \mathbf{E}) \otimes (\pi'_i \mathbf{E}).$$

It leads to:

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n e_i \otimes e_i - \frac{1}{n} \sum_{i=1}^n e_i \otimes (\pi'_i \mathbf{E}).$$

We replace $\hat{\Sigma}$ and $\tilde{\Sigma}$ by this expression and use the triangular inequality to obtain:

$$|f(\mathbf{E}) - f(\tilde{\mathbf{E}}_{i_0})| \leq \frac{1}{n} \left\| \sum_{i=1}^n e_i \otimes e_i - \tilde{e}_i \otimes \tilde{e}_i \right\|_{\text{HS}(\mathcal{H})} + \underbrace{\frac{1}{n} \left\| \sum_{i=1}^n \tilde{e}_i \otimes (\pi'_i \tilde{\mathbf{E}}) - e_i \otimes (\pi'_i \mathbf{E}) \right\|_{\text{HS}(\mathcal{H})}}_{=A}.$$

The first term is such that:

$$\begin{aligned} \frac{1}{n} \left\| \sum_{i=1}^n e_i \otimes e_i - \tilde{e}_i \otimes \tilde{e}_i \right\|_{\text{HS}(\mathcal{H})} &= \frac{1}{n} \|e_{i_0} \otimes e_{i_0} - \tilde{e}_{i_0} \otimes \tilde{e}_{i_0}\|_{\text{HS}(\mathcal{H})} \\ &\leq \frac{2M_e^2}{n}. \end{aligned} \tag{18}$$

The second term can be decomposed as:

$$\begin{aligned}
A &= \frac{1}{n} \left\| \sum_{i=1, i \neq i_0}^n e_i \otimes \left(\pi'_i(\tilde{\mathbf{E}} - \mathbf{E}) \right) + \tilde{e}_{i_0} \otimes \left(\pi'_{i_0} \tilde{\mathbf{E}} \right) - e_{i_0} \otimes \left(\pi'_{i_0} \tilde{\mathbf{E}} \right) \right\|_{\text{HS}(\mathcal{H})} \\
&= \frac{1}{n} \left\| \sum_{i=1}^n e_i \otimes \left(\pi'_i(\tilde{\mathbf{E}} - \mathbf{E}) \right) + (\tilde{e}_{i_0} - e_{i_0}) \otimes \left(\pi'_{i_0} \tilde{\mathbf{E}} \right) \right\|_{\text{HS}(\mathcal{H})} \\
&\leq \underbrace{\frac{1}{n} \left\| \sum_{i=1}^n e_i \otimes \left(\pi'_i(\tilde{\mathbf{E}} - \mathbf{E}) \right) \right\|_{\text{HS}(\mathcal{H})}}_{A_1} + \underbrace{\frac{1}{n} \left\| (\tilde{e}_{i_0} - e_{i_0}) \otimes \left(\pi'_{i_0} \tilde{\mathbf{E}} \right) \right\|_{\text{HS}(\mathcal{H})}}_{A_2}.
\end{aligned}$$

We first bound A_1 . Note that par definition of $\tilde{\mathbf{E}}$, we have $\pi'_i(\tilde{\mathbf{E}} - \mathbf{E}) = \pi_{i, i_0}(\tilde{e}_{i_0} - e_{i_0})$ and thus $\sum_{i=1}^n e_i \otimes \left(\pi'_i(\tilde{\mathbf{E}} - \mathbf{E}) \right) = \sum_{i=1}^n e_i \otimes (\pi_{i, i_0}(\tilde{e}_{i_0} - e_{i_0})) = (\sum_{i=1}^n \pi_{i, i_0} e_i) \otimes (\tilde{e}_{i_0} - e_{i_0}) = (\mathbf{P}_X \mathbf{E})_{i_0} \otimes (\tilde{e}_{i_0} - e_{i_0})$. From this we obtain that

$$\begin{aligned}
A_1 &\leq \frac{1}{n} \left\| (\mathbf{P}_X \mathbf{E})_{i_0} \otimes (\tilde{e}_{i_0} - e_{i_0}) \right\|_{\text{HS}(\mathcal{H})} \\
&\leq \frac{2M_e}{n} \left\| (\mathbf{P}_X \mathbf{E})_{i_0} \right\|_{\mathcal{H}}
\end{aligned}$$

Regarding A_2 , by noting that $\pi'_{i_0} \tilde{\mathbf{E}} = (\mathbf{P}_X \tilde{\mathbf{E}})_{i_0}$, we get that $A_2 \leq \frac{2M_e}{n} \left\| (\mathbf{P}_X \tilde{\mathbf{E}})_{i_0} \right\|_{\mathcal{H}}$ and thus $A \leq \frac{2M_e}{n} \left(\left\| (\mathbf{P}_X \mathbf{E})_{i_0} \right\|_{\mathcal{H}} + \left\| (\mathbf{P}_X \tilde{\mathbf{E}})_{i_0} \right\|_{\mathcal{H}} \right)$. We conclude the proof by combining this bound with (18). \square

Lemma 6. *Under Assumptions 3 and 4, we have:*

$$\mathbb{E} \left\| \hat{\Sigma} - \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{M_e^2 \sqrt{n-p}}{n},$$

where M_e is defined in Lemma 4.

Proof. Note that:

$$\begin{aligned}
\mathbb{E} \left\| \hat{\Sigma} - \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 &= \mathbb{E} \left\| \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 - 2 \mathbb{E} \left\langle \hat{\Sigma}, \mathbb{E} \hat{\Sigma} \right\rangle_{\text{HS}(\mathcal{H})} + \left\| \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 \\
&= \mathbb{E} \left\| \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 - 2 \left\langle \mathbb{E} \hat{\Sigma}, \mathbb{E} \hat{\Sigma} \right\rangle_{\text{HS}(\mathcal{H})} + \left\| \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 \\
&= \mathbb{E} \left\| \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 - \left\| \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2.
\end{aligned}$$

By Jensen's inequality, we have that:

$$\begin{aligned}
\mathbb{E} \left\| \hat{\Sigma} - \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})} &\leq \left[\mathbb{E} \left\| \hat{\Sigma} - \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 \right]^{\frac{1}{2}} \\
&\leq \left[\mathbb{E} \left\| \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 - \left\| \mathbb{E} \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 \right]^{\frac{1}{2}}.
\end{aligned}$$

Remember that $\mathbf{P}_X = (\pi_{i,j})_{1 \leq i, j \leq n}$ is the orthogonal projection on $\text{Im}(\mathbf{X})$. We can develop $\hat{\Sigma}$ to obtain that:

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n (1 - \pi_{i,i}) e_i \otimes e_i - \frac{1}{n} \sum_{i,j=1, i \neq j}^n \pi_{i,j} e_i \otimes e_j. \quad (19)$$

Thus:

$$\left\| \hat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 = A - 2B + C,$$

where

$$\begin{aligned}
A &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n (1 - \pi_{i,i})(1 - \pi_{j,j}) \langle e_i \otimes e_i, e_j \otimes e_j \rangle_{\text{HS}(\mathcal{H})}, \\
B &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1, k \neq j}^n (1 - \pi_{i,i}) \pi_{j,k} \langle e_i \otimes e_i, e_j \otimes e_k \rangle_{\text{HS}(\mathcal{H})}, \\
C &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1, j \neq i}^n \sum_{k=1}^n \sum_{l=1, l \neq k}^n \pi_{i,j} \pi_{k,l} \langle e_i \otimes e_j, e_k \otimes e_l \rangle_{\text{HS}(\mathcal{H})}.
\end{aligned}$$

We now compute the expectation of each term:

$$\begin{aligned}
\mathbb{E}(A) &= \frac{1}{n^2} \sum_{i=1}^n (1 - \pi_{i,i})^2 \mathbb{E} \left(\langle e_i \otimes e_i, e_i \otimes e_i \rangle_{\text{HS}(\mathcal{H})} \right) \\
&\quad + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1, j \neq i}^n (1 - \pi_{i,i})(1 - \pi_{j,j}) \langle \mathbb{E}(e_i \otimes e_i), \mathbb{E}(e_j \otimes e_j) \rangle_{\text{HS}(\mathcal{H})} \\
&= \frac{1}{n^2} \sum_{i=1}^n (1 - \pi_{i,i})^2 \mathbb{E} \left(\|e_i\|_{\mathcal{H}}^4 \right) \\
&\quad + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1, j \neq i}^n (1 - \pi_{i,i})(1 - \pi_{j,j}) \langle \Sigma_i, \Sigma_j \rangle_{\text{HS}(\mathcal{H})}.
\end{aligned}$$

We can directly see that every term of B contains at least one indice j or k different from the others, involving that the expectation of each term of B is null, which give that $\mathbb{E}(B) = 0$. The same happens for every term of C where indices k and l are different from indices i and j . Consequently we have:

$$\begin{aligned}
\mathbb{E}(C) &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1, j \neq i}^n \pi_{i,j}^2 \left(\mathbb{E} \langle e_i \otimes e_j, e_i \otimes e_j \rangle_{\text{HS}(\mathcal{H})} + \mathbb{E} \langle e_i \otimes e_j, e_j \otimes e_i \rangle_{\text{HS}(\mathcal{H})} \right) \\
&= \frac{2}{n^2} \sum_{i=1}^n \sum_{j=1, j \neq i}^n \pi_{i,j}^2 \mathbb{E} \|e_i\|_{\mathcal{H}}^2 \mathbb{E} \|e_j\|_{\mathcal{H}}^2.
\end{aligned}$$

According to Equation (19), $\mathbb{E}\widehat{\Sigma}$ is such that:

$$\begin{aligned}
\mathbb{E}\widehat{\Sigma} &= \frac{1}{n} \sum_{i=1}^n (1 - \pi_{i,i}) \mathbb{E}(e_i \otimes e_i) + \frac{1}{n} \sum_{i,j=1, i \neq j}^n \sum_{j=1, j \neq i}^n \pi_{i,j} \mathbb{E} e_i \otimes \mathbb{E} e_j \\
&= \frac{1}{n} \sum_{i=1}^n (1 - \pi_{i,i}) \Sigma_i.
\end{aligned}$$

Then:

$$\begin{aligned}
\left\| \mathbb{E}\widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n (1 - \pi_{i,i})(1 - \pi_{j,j}) \langle \Sigma_i, \Sigma_j \rangle_{\text{HS}(\mathcal{H})} \\
&= \frac{1}{n^2} \sum_{i=1}^n (1 - \pi_{i,i})^2 \|\Sigma_i\|_{\text{HS}(\mathcal{H})}^2 + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1, j \neq i}^n (1 - \pi_{i,i})(1 - \pi_{j,j}) \langle \Sigma_i, \Sigma_j \rangle_{\text{HS}(\mathcal{H})}.
\end{aligned}$$

Now we can sum $\mathbb{E}(A)$, $\mathbb{E}(C)$ and $\left\| \mathbb{E}\widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2$:

$$\begin{aligned}
\mathbb{E} \left\| \widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 - \left\| \mathbb{E}\widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 &= \mathbb{E}(A) + \mathbb{E}(C) - \left\| \mathbb{E}\widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 \\
&= \frac{1}{n^2} \sum_{i=1}^n (1 - \pi_{i,i})^2 \left(\mathbb{E} \|e_i\|_{\mathcal{H}}^4 - \|\Sigma_i\|_{\text{HS}(\mathcal{H})}^2 \right) \\
&\quad + \frac{2}{n^2} \sum_{i,j=1, i \neq j}^n \pi_{i,j}^2 \mathbb{E} \|e_i\|_{\mathcal{H}}^2 \mathbb{E} \|e_j\|_{\mathcal{H}}^2
\end{aligned}$$

Note that under Assumptions 3 and 4, Lemma 4 gives that for $i \in \{1, \dots, n\}$, $\mathbb{E} \|e_i\|_{\mathcal{H}}^4 - \|\Sigma_i\|_{\text{HS}(\mathcal{H})}^2 \leq \mathbb{E} \|e_i\|_{\mathcal{H}}^4 \leq M_e^4$ and that $\mathbb{E} \|e_i\|_{\mathcal{H}}^2 \leq M_e^2$. Thus, we have :

$$\begin{aligned} \mathbb{E} \left\| \widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 - \left\| \mathbb{E} \widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})}^2 &\leq \frac{1}{n^2} \sum_{i=1}^n (1 - \pi_{i,i})^2 M_e^4 + \frac{2}{n^2} \sum_{i,j=1, i \neq j}^n \pi_{i,j}^2 M_e^4 \\ &\leq \frac{M_e^4}{n^2} (n - p + 2p) \\ &\leq \frac{M_e^4 (n + p)}{n^2} \end{aligned}$$

by applying Lemma 8 to both terms. It finally gives that

$$\mathbb{E} \left\| \widehat{\Sigma} - \mathbb{E} \widehat{\Sigma} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{M_e^2 \sqrt{n-p}}{n}.$$

□

C.4 Probability bound on the inverse covariance operator

Under Assumption 1, we define $\bar{\Sigma}_n = \Sigma_T + \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \Sigma_{i,T^c}$. Under Assumptions 1 and 2, and according to Lemma 7, $\bar{\Sigma}_n$ is symmetric positive operator and admits a spectral decomposition $\Sigma_T = \sum_{t=1}^T \lambda_t f_t \otimes f_t$ where the eigenvalues $\lambda_1 > \dots > \lambda_T$ and eigenfunctions f_1, \dots, f_T correspond to those of Σ_T .

Proposition 3. *Assume that Assumptions 1, 2, 3 and 4 are verified. There exists N_T such that for $n \geq N_T$, we have with probability $1 - e^{-\xi}$:*

$$\left\| \widehat{\Sigma}_T^{-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{\zeta(n, \xi)}{\lambda_T - \zeta(n, \xi)} \left(\sum_{t=1}^T \frac{2}{\tilde{\delta}_t} + \frac{T}{\lambda_T} \right),$$

where $\zeta(n, \xi)$ is given in Proposition 2.

Remark 1. *In the Proposition, for $t < T$, all the spectral gaps $\tilde{\delta}_t$ do not depend on n . Regarding the spectral gap at T , note that $\tilde{\delta}_T = \delta_T \wedge \delta_{T-1} \geq (\lambda_T - \mu_T)/2 \wedge (\lambda_{T-1} - \lambda_T)/2 > 0$ (see Remark 2). Consequently the exponential bound of the proposition implies the convergence in probability.*

Proof. For $t \in \{1, \dots, n-1\}$, denote $\delta_t = (\lambda_t - \lambda_{t+1})/2$ and $\tilde{\delta}_t = \min(\delta_t, \delta_{t-1})$ for the spectral gaps of Σ . Let $\Pi_{f_t} = f_t \otimes f_t$ and $\Pi_{\hat{f}_t} = \hat{f}_t \otimes \hat{f}_t$ the orthogonal projector on the axis spanned by f_t (resp. \hat{f}_t).

According to Proposition 2, we have that with probability $1 - e^{-\xi}$:

$$\left\| \widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \leq \zeta(n, \xi).$$

Moreover, according to the Hoffman-Wielandt Inequality in infinite dimensional setting [2] (see also the proof of Theorem 3 from [28]), we have

$$\max_{t=1, \dots, n} |\hat{\lambda}_t - \lambda_t| \leq \left\| \widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})}. \quad (20)$$

According Lemma 7, the eigenvalues $\lambda_1, \dots, \lambda_T$ of $\bar{\Sigma}_n$ are simple. In particular the spectral gap at T is strictly positive: $\lambda_T - \lambda_{T+1} \geq \lambda_T - \mu_{T+1} > 0$, see also Remark 2.

It then can be easily checked that there exists N_T such for $n \geq N_T$, the event

$$A = \left\{ \left\| \widehat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \leq \zeta(n, \xi) \wedge \min_{t=1 \dots T} \delta_t \right\} \cap \left\{ \hat{\lambda}_T > 0 \right\} \cap \left\{ \hat{\lambda}_1 > \dots > \hat{\lambda}_T > \hat{\lambda}_{T+1} \right\}$$

is satisfied with probability at least $1 - e^{-\xi}$. Moreover, N_T only depends on λ_T , on $\delta_1, \dots, \delta_T$, and on the constants involved in $\zeta(n, \xi)$, that is p , M_e and $\|\Sigma_T\|_{\text{HS}(\mathcal{H})}$.

In particular, Theorem 4 applies on the event A : for any $t \in \{1, \dots, T\}$,

$$\left\| \Pi_{\hat{f}_t} - \Pi_{f_t} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{2 \left\| \hat{\Sigma} - \Sigma \right\|_{\text{HS}(\mathcal{H})}}{\tilde{\delta}_t}. \quad (21)$$

On the event A , the rank r of $\hat{\Sigma}_T$ is larger than T and thus $\hat{\Sigma}_T^{-1} = \sum_{t=1}^T \hat{\lambda}_t^{-1} \hat{f}_t \otimes \hat{f}_t$. We can use the spectral decomposition of the inverse operators to obtain that on A :

$$\begin{aligned} \left\| \hat{\Sigma}_T^{-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})} &= \left\| \sum_{t=1}^T \hat{\lambda}_t^{-1} \hat{f}_t \otimes \hat{f}_t - \lambda_t^{-1} f_t \otimes f_t \right\|_{\text{HS}(\mathcal{H})} \\ &= \left\| \sum_{t=1}^T \hat{\lambda}_t^{-1} \Pi_{\hat{f}_t} - \lambda_t^{-1} \Pi_{f_t} \right\|_{\text{HS}(\mathcal{H})} \\ &\leq \sum_{t=1}^T \hat{\lambda}_t^{-1} \left\| \Pi_{\hat{f}_t} - \Pi_{f_t} \right\|_{\text{HS}(\mathcal{H})} + \sum_{t=1}^T |\hat{\lambda}_t^{-1} - \lambda_t^{-1}| \left\| \Pi_{f_t} \right\|_{\text{HS}(\mathcal{H})}. \end{aligned} \quad (22)$$

For the first term of (22), on the event A we find from (21) that for $n \geq \tilde{N}_T$

$$\begin{aligned} \sum_{t=1}^T \hat{\lambda}_t^{-1} \left\| \Pi_{\hat{f}_t} - \Pi_{f_t} \right\|_{\text{HS}(\mathcal{H})} &\leq \frac{\zeta(n, \xi)}{\hat{\lambda}_T} \sum_{t=1}^T \frac{2}{\tilde{\delta}_t} \\ &\leq \frac{\zeta(n, \xi)}{\lambda_T - \zeta(n, \xi)} \sum_{t=1}^T \frac{2}{\tilde{\delta}_t}. \end{aligned}$$

Regarding the second term in the upper bound (22), on A and for $n \geq \tilde{N}_T$, we have that $\frac{1}{\lambda_T} \leq \frac{1}{\lambda_T - \zeta(n, \xi)}$ and thus for all $t \in \{1, \dots, T\}$

$$\begin{aligned} |\hat{\lambda}_t^{-1} - \lambda_t^{-1}| &= \hat{\lambda}_t^{-1} \lambda_t^{-1} |\hat{\lambda}_t - \lambda_t| \\ &\leq \lambda_T^{-1} \hat{\lambda}_T^{-1} \zeta(n, \xi), \\ &\leq \frac{1}{\lambda_T} \frac{\zeta(n, \xi)}{\lambda_T - \zeta(n, \xi)}. \end{aligned}$$

Also note that $\left\| \Pi_{f_t} \right\|_{\text{HS}(\mathcal{H})} = \|f_t\|_{\mathcal{H}}^2 = 1$. Thus, on the event A , we have that

$$\sum_{t=1}^T |\hat{\lambda}_t^{-1} - \lambda_t^{-1}| \left\| \Pi_{f_t} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{T}{\lambda_T} \frac{\zeta(n, \xi)}{\lambda_T - \zeta(n, \xi)}.$$

Thus, for $n \geq N_T$, we have with probability $1 - e^{-\xi}$:

$$\left\| \hat{\Sigma}_T^{-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{\zeta(n, \xi)}{\lambda_T - \zeta(n, \xi)} \left(\sum_{t=1}^T \frac{2}{\tilde{\delta}_t} + \frac{T}{\lambda_T} \right).$$

□

Lemma 7. *Under Assumptions 1 and 2, the operator $\bar{\Sigma}_n = \Sigma_T + \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \Sigma_{i,T^c}$ is a symmetric and positive operator that admits a spectral decomposition $\bar{\Sigma}_n = \sum_{t=1}^t \lambda_t f_t \otimes f_t$ where the eigenvalues are in descending order, with the first ones $\lambda_1 > \dots > \lambda_T$ and the corresponding eigenfunctions f_1, \dots, f_T being the same as those of Σ_T . Moreover, the spectral gap at T is strictly positive: $\lambda_T - \lambda_{T+1} \geq \lambda_T - \mu_{T+1} > 0$.*

Remark 2. *Note that λ_t depends on n for $t \geq T + 1$ and it does not for $t \leq T$. Thus, the spectral gap $\delta_t = \lambda_t - \lambda_{t+1}$ does not depend on n for $t \leq T - 1$. Moreover, the last spectral gap $\lambda_T - \lambda_{T+1}$ is lower bounded by $\lambda_T - \mu_{T+1}$ for all $n \geq 1$.*

Proof. The operator $\bar{\Sigma}_n$ is symmetric and positive as a linear combination of symmetric and positive operators, and because $\pi_{i,i}^\perp = \sum_{j=1}^n \pi_{i,j}^{\perp 2} \geq 0$ (see Lemma 8). We directly check that for $t \leq T$, the eigenfunction f_t for the eigenvalue λ_t of Σ_T is also an eigenfunction of $\bar{\Sigma}_n$ for the eigenvalue λ_t . Moreover, for any $h \in \text{Span}(f_1, \dots, f_T)^\perp$, we have $\|\bar{\Sigma}_n h\|_{\mathcal{H}} \leq \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \|\Sigma_{i,T^c} h\|_{\mathcal{H}} \leq \|h\|_{\mathcal{H}} \max_{i=1 \dots n} \lambda_{i,T+1} \left(\frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \right) < \frac{n-p}{n} \mu_{T+1} \|h\|_{\mathcal{H}}$. Thus $0 \leq \lambda_t < \mu_{T+1} < \lambda_T$ for any $t > T + 1$. In particular $\lambda_T - \lambda_{T+1} \geq \lambda_T - \mu_{T+1} > 0$. \square

D Proof of Theorem 2: KFDA with Nyström

This section provides the proof for the consistency of the Nyström TKHL statistic, see Theorem 2.

D.1 Main proof of Theorem 2

Let $\mathbf{P}_X^\perp = \mathbf{I}_n - \mathbf{P}_X = (\pi_{i,j}^\perp)_{1 \leq i,j \leq n}$ the orthogonal projection on $\text{Im}(\mathbf{X})^\perp$. Under Assumptions 1' and 2', as for the non-Nyström case, we consider

$$\begin{aligned} \bar{\Sigma}_n &:= \Sigma_T + \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \Sigma_{i,T^c} \\ &= \Sigma_m + \frac{1}{n} \sum_{i=1}^n \pi_{i,i}^\perp \Sigma_{i,m^c}, \end{aligned}$$

where the second equality is valid because $T \leq m$. According to Lemma 7, we have that $\bar{\Sigma}_n$ is a symmetric and semi-definite positive operator that admits a spectral decomposition $\bar{\Sigma}_n = \sum_{t \geq 1} \lambda_t f_t \otimes f_t$ where the eigenvalues are in descending order, with the first ones $\lambda_1 > \dots > \lambda_m$ and the corresponding eigenfunctions f_1, \dots, f_m being the same as those of Σ_m .

Let I be a n -uple of indices sampled uniformly in $\{1, \dots, n\}$ to define the q landmarks for Nyström. Let also $\mathbf{Z} = (\phi(Y_i))_{i \in I} = (Z_1, \dots, Z_q)$ be the q landmarks and \mathbf{X}^I the associated design matrix extracted from \mathbf{X} . Let $\hat{\Sigma}^I$ be the landmark residual covariance operator induced by the linear model (1) based on \mathbf{Z}, \mathbf{X}^I . For this model we define (as in the non-Nyström case) the operator $\bar{\Sigma}_I = \Sigma_m + \frac{1}{q} \sum_{i=1}^q \pi_{i,i}^{I,\perp} \Sigma_{i,m^c}$, where $\mathbf{P}_X^{I,\perp} = (\pi_{i,j}^{I,\perp})$ is the orthogonal projection onto $\text{Im}(\mathbf{X}^I)^\perp$ in \mathbb{R}^q . According to Lemma 7 (adapted to this model), we have that $\bar{\Sigma}_I$ is a symmetric and semi-definite positive operator. It admits a spectral decomposition which first m eigenvalues $\lambda_1 > \dots > \lambda_m$ and corresponding eigenfunctions f_1, \dots, f_m are the same as those of Σ_m .

For $s \leq m - 1$, we denote by $\delta_t = (\lambda_t - \lambda_{t+1})/2$ the spectral gaps of Σ_m (and of $\bar{\Sigma}_I$).

We define the m first unit eigenfunctions $(\hat{a}_1, \dots, \hat{a}_m)$ of the landmark residual covariance operator $\hat{\Sigma}^I$ as our Nyström anchors. The Nyström residual covariance operator $\hat{\Sigma}^a$ is finally defined as $\hat{\Sigma}^a = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^a \otimes \hat{e}_i^a$ where \hat{e}_i^a is the orthogonal projection of \hat{e}_i onto $\text{Span}(\hat{a}_1, \dots, \hat{a}_m)$.

According to Proposition 4, when q is large enough, we have with probability $1 - 2e^{-\xi}$:

$$\left\| \hat{\Sigma}^a - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \leq \frac{1}{2} \zeta(n, \xi) + 2m M_e^2 \zeta(q, \xi) \sum_{s=1}^m \frac{1}{\delta_s}.$$

From this, we obtain (see Proposition 5) that with probability $1 - 2e^{-\xi}$,

$$\left\| \hat{\Sigma}_T^{a,-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{\eta(n, q, \zeta)}{\lambda_T - \eta(n, q, \zeta)} \left(\sum_{t=1}^T \frac{2}{\delta_t} + \frac{T}{\lambda_T} \right)$$

with $\eta(n, q, \zeta) = \frac{C_1}{\sqrt{n}} (1 + \sqrt{\zeta}) + \frac{C_2}{\sqrt{q}} (1 + \sqrt{\zeta}) \sum_{s=1}^m \frac{1}{\delta_s}$, where C_1 and C_2 are constants that depend on $(\lambda_s)_{s \leq m}, \mu_{m+1}, p$ and M_e . Thus, we have $\left\| \hat{\Sigma}_T^{a,-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{C'}{\sqrt{q}} (1 + \sqrt{\zeta})$ where C' is a constant which does not depend on q (and n). Thus $\hat{\Sigma}_T^{a,-1}$ converges in probability to Σ_T^{-1} as q tends to infinity.

As the landmarks are not used to approximate \widehat{H}_L , we conclude the proof of Theorem 2 by following the lines of the proof of Theorem 1: we use the convergence in probability of $\widehat{\Sigma}_T^a$ (instead of these of $\widehat{\Sigma}_T^{-1}$) together with Proposition 1.

D.2 Probability bound on the Nyström residual covariance

Proposition 4. *Under Assumptions 1', 2', 3 and 4, when q is large enough, we have with probability $1 - 2e^{-\xi}$:*

$$\left\| \widehat{\Sigma}^a - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \leq \frac{1}{2} \zeta(n, \xi) + 2mM_e^2 \zeta(q, \xi) \sum_{s=1}^m \frac{1}{\widetilde{\delta}_s} := \eta(n, q, \zeta)$$

where ζ is defined in Proposition 2, and where $\widetilde{\delta}_s = \delta_s \wedge \delta_{s-1}$ for $s < m$ and $\delta_m := (\lambda_m - \mu_{m+1})/2$.

Proof. For $s \in \{1, \dots, m\}$, we denote by Π_{f_s} and $\Pi_{\widehat{a}_s}$ the orthogonal projections onto the one-dimensional axis supported by f_s and \widehat{a}_s respectively. Then we have $\Pi_{f_s} = f_s \otimes f_s$, $\Pi_{\widehat{a}_s} = \widehat{a}_s \otimes \widehat{a}_s$ and by definition of $\widehat{\Sigma}^a$:

$$\widehat{\Sigma}^a = \frac{1}{n} \sum_{i=1}^n \left(\sum_{s=1}^m \Pi_{\widehat{a}_s} \widehat{e}_i \right) \otimes \left(\sum_{s'=1}^m \Pi_{\widehat{a}_{s'}} \widehat{e}_i \right).$$

We also define :

$$\widetilde{\Sigma}^a = \frac{1}{n} \sum_{i=1}^n \left(\sum_{s=1}^m \Pi_{f_s} \widehat{e}_i \right) \otimes \left(\sum_{s'=1}^m \Pi_{f_{s'}} \widehat{e}_i \right),$$

and we have

$$\left\| \widehat{\Sigma}^a - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \leq \left\| \widehat{\Sigma}^a - \widetilde{\Sigma}^a \right\|_{\text{HS}(\mathcal{H})} + \left\| \widetilde{\Sigma}^a - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})}. \quad (23)$$

First, for $i \in \{1, \dots, n\}$ and $s \in \{1, \dots, m\}$, we have :

$$(\Pi_{\widehat{a}_s} \widehat{e}_i) \otimes (\Pi_{\widehat{a}_{s'}} \widehat{e}_i) - (\Pi_{f_s} \widehat{e}_i) \otimes (\Pi_{f_{s'}} \widehat{e}_i) = ((\Pi_{\widehat{a}_s} - \Pi_{f_s}) \widehat{e}_i) \otimes (\Pi_{\widehat{a}_{s'}} \widehat{e}_i) + (\Pi_{f_s} \widehat{e}_i) \otimes ((\Pi_{\widehat{a}_{s'}} - \Pi_{f_{s'}}) \widehat{e}_i)$$

According to the triangular inequality, we have :

$$\begin{aligned} \left\| \widehat{\Sigma}^a - \widetilde{\Sigma}^a \right\|_{\text{HS}(\mathcal{H})} &\leq \frac{1}{n} \sum_{i=1}^n \sum_{s, s'=1}^m \left\| ((\Pi_{\widehat{a}_s} - \Pi_{f_s}) \widehat{e}_i) \otimes (\Pi_{\widehat{a}_{s'}} \widehat{e}_i) \right\|_{\text{HS}(\mathcal{H})} \\ &\quad + \left\| (\Pi_{f_s} \widehat{e}_i) \otimes ((\Pi_{\widehat{a}_{s'}} - \Pi_{f_{s'}}) \widehat{e}_i) \right\|_{\text{HS}(\mathcal{H})} \\ &\leq \frac{1}{n} \sum_{i=1}^n \sum_{s, s'=1}^m \left\| \Pi_{\widehat{a}_s} - \Pi_{f_s} \right\|_{\text{HS}(\mathcal{H})} \left\| \Pi_{\widehat{a}_{s'}} \right\|_{\text{HS}(\mathcal{H})} \|\widehat{e}_i\|_{\mathcal{H}}^2 \\ &\quad + \left\| \Pi_{\widehat{a}_{s'}} - \Pi_{f_{s'}} \right\|_{\text{HS}(\mathcal{H})} \left\| \Pi_{f_s} \right\|_{\text{HS}(\mathcal{H})} \|\widehat{e}_i\|_{\mathcal{H}}^2 \\ &\leq 2mM_e^2 \sum_{s=1}^m \left\| \Pi_{\widehat{a}_s} - \Pi_{f_s} \right\|_{\text{HS}(\mathcal{H})}. \end{aligned}$$

because $\left\| \Pi_{f_s} \right\|_{\text{HS}(\mathcal{H})} = \left\| \Pi_{\widehat{a}_s} \right\|_{\text{HS}(\mathcal{H})} = 1$ and $\|\widehat{e}_i\| \leq M_e$, see Lemma 4.

We now apply Proposition 2 to $\widehat{\Sigma}^I$ with the sample \mathbf{Z} which associated design matrix design is \mathbf{X}^I . Assumption 1' is still satisfied for \mathbf{Z} , which corresponds to Assumption 1 for a truncate parameter m (instead of T). Assumptions 3 and 4 are also obviously valid. Thus, conditionally to the Nyström sample I , with probability larger than $1 - e^{-\xi}$,

$$\left\| \widehat{\Sigma}^I - \bar{\Sigma}_I \right\|_{\text{HS}(\mathcal{H})} \leq \frac{p}{q} \left\| \Sigma_m \right\|_{\text{HS}(\mathcal{H})} + \frac{M_e^2}{\sqrt{q}} \left(1 + 4\sqrt{\frac{3\xi}{2}} \right) =: \zeta(q, \xi).$$

Note that this upper bound does not depends on the Nyström sample I and thus the probability bound is also valid after intergrating according to the distribution of I .

The m first eigenvalues of $\bar{\Sigma}_I$ being simple, following the first lines of the proof of Proposition 3, there exists Q_m such for $q \geq Q_m$, the event

$$\tilde{A} = \left\{ \|\hat{\Sigma}^I - \bar{\Sigma}_I\|_{\text{HS}(\mathcal{H})} \leq \zeta(q, \xi) \wedge \min_{t=1 \dots T} \delta_t \right\} \cap \left\{ \hat{\lambda}_T > 0 \right\} \cap \left\{ \hat{\lambda}_1 > \dots > \hat{\lambda}_T > \hat{\lambda}_{T+1} \right\}$$

is satisfied with probability at least $1 - e^{-\xi}$. Regarding the last spectral gap, according to Lemma 7, the last spectral gap $\lambda_m - \lambda_{m+1}$ is lower bounded by $\lambda_m - \mu_{m+1}$ for all $q \geq 1$, see Remark 2.

Following the first lines of the proof of Proposition 3, the eigenvalues of $\bar{\Sigma}_I$ being simple, we find that Theorem 4 can be applied on an event of probability at least $1 - e^{-\xi}$ and for q large enough: for $s \leq m$,

$$\|\Pi_{\hat{a}_s} - \Pi_{f_s}\|_{\text{HS}(\mathcal{H})} \leq \frac{2 \left\| \bar{\Sigma}_I - \hat{\Sigma}^I \right\|_{\text{HS}(\mathcal{H})}}{\tilde{\delta}_s}$$

where $\tilde{\delta}_s = \delta_s \wedge \delta_{s-1}$, and where we take $\delta_m := (\lambda_m - \mu_{m+1})/2$. Thus, for q large enough, we have that on the event \tilde{A} :

$$\left\| \hat{\Sigma}^a - \tilde{\Sigma}^a \right\|_{\text{HS}(\mathcal{H})} \leq 4mM_e^2 \zeta(q, \xi) \sum_{s=1}^m \frac{1}{\tilde{\delta}_s}. \quad (24)$$

Regarding the second term in the upper bound (23), we have $\Pi_{f_s} = f_s \otimes f_s$, which leads to :

$$\begin{aligned} \tilde{\Sigma}^a &= \frac{1}{n} \sum_{i=1}^n \sum_{s,s'=1}^m (\Pi_{f_s} \hat{e}_i) \otimes (\Pi_{f_{s'}} \hat{e}_i) \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{s,s'=1}^m \langle f_s, \hat{e}_i \rangle_{\mathcal{H}} \langle f_{s'}, \hat{e}_i \rangle_{\mathcal{H}} f_s \otimes f_{s'} \\ &= \sum_{s,s'=1}^m \left\langle f_s, \hat{\Sigma} f_{s'} \right\rangle_{\mathcal{H}} f_s \otimes f_{s'}. \end{aligned}$$

We also have that for $\bar{\Sigma}_n$:

$$\begin{aligned} \sum_{s,s'=1}^m \langle f_s, \bar{\Sigma}_n f_{s'} \rangle_{\mathcal{H}} f_s \otimes f_{s'} &= \sum_{s,s'=1}^m \lambda_{s'} \langle f_s, f_{s'} \rangle_{\mathcal{H}} f_s \otimes f_{s'} \\ &= \sum_{s=1}^m \lambda_s f_s \otimes f_s \\ &= \Sigma_m. \end{aligned}$$

Finally, we obtain:

$$\begin{aligned} \left\| \tilde{\Sigma}^a - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} &= \sum_{s,s'=1}^m \left\langle f_s, (\hat{\Sigma} - \bar{\Sigma}_n) f_{s'} \right\rangle_{\mathcal{H}} f_s \otimes f_{s'} \\ &\leq \sum_{s,s'=1}^m \left\| \hat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})} \|f_s\|_{\mathcal{H}}^2 \|f_{s'}\|_{\mathcal{H}}^2 \\ &\leq m^2 \left\| \hat{\Sigma} - \bar{\Sigma}_n \right\|_{\text{HS}(\mathcal{H})}. \end{aligned}$$

We then apply Proposition 2 (for all the data) to control this term and combine it with (24) to conclude the proof of the proposition. \square

D.3 Probability bound on the inverse Nyström residual covariance

Proposition 5. *Assume that Assumptions 1', 2' 3 and 4 are verified. When q and n are large enough, with probability $1 - 2e^{-\xi}$:*

$$\left\| \hat{\Sigma}_T^{-1} - \Sigma_T^{-1} \right\|_{\text{HS}(\mathcal{H})} \leq \frac{\eta(n, q, \zeta)}{\lambda_T - \eta(n, q, \zeta)} \left(\sum_{t=1}^T \frac{2}{\tilde{\delta}_t} + \frac{T}{\lambda_T} \right),$$

where $\eta(n, q, \zeta)$ is given in Proposition 4.

Proof. We follow exactly the lines of Proposition 3, but starting from Proposition 4 instead of Proposition 2. \square

E Auxiliary results

E.1 McDiarmid Inequality

Theorem 3 (McDiarmid inequality [18]). *If Y_1, \dots, Y_n are i.i.d. random variables in a measurable space \mathcal{Y} and the function*

$$\begin{aligned} \mathcal{Y}^n &\longrightarrow \mathbb{R} \\ f : y_1, \dots, y_n &\longmapsto f(y_1, \dots, y_n) \end{aligned}$$

is such that for all $i \in \{1, \dots, n\}$, we have:

$$\sup_{y_1, \dots, y_n, y_{i'} \in \mathcal{Y}} |f(y_1, \dots, y_i, \dots, y_n) - f(y_1, \dots, y_{i'}, \dots, y_n)| \leq c_i.$$

Then we have with probability lower than $e^{-\xi}$ that:

$$f(y_1, \dots, y_n) - \mathbb{E}(f(y_1, \dots, y_n)) \geq \sqrt{\frac{\xi}{2} \sum_{i=1}^n c_i^2},$$

and we also have with probability lower than $e^{-\xi}$ that:

$$\mathbb{E}(f(y_1, \dots, y_n)) - f(y_1, \dots, y_n) \geq \sqrt{\frac{\xi}{2} \sum_{i=1}^n c_i^2}.$$

E.2 Perturbation bounds

In this section, we use some results from operator perturbations theory to obtain an exponential bound on a truncated inverse of $\widehat{\Sigma}$, that is used in the proof of Theorem 1. We apply the following result from [28] and [14] rewritten in the particular case where only the T first eigenvalues are assumed to be simple.

Theorem 4 (Theorem 2 from [28], Lemma 5.2 from [14]). *Let $A \in \text{HS}(\mathcal{H})$ be a symmetric positive operator such that its T first eigenvalues are positive and simple: $\lambda_1 > \lambda_2 > \dots > \lambda_T > \lambda_{T+1} \geq \lambda_{T+2} \geq \dots \geq 0$. The spectral gaps of A at dimension t are denoted $\delta_t = (\lambda_t - \lambda_{t+1})/2$ and $\widetilde{\delta}_t = \delta_t \wedge \delta_{t-1}$. Let $B \in \text{HS}(\mathcal{H})$ a symmetric operator such that $A + B$ is a positive operator such that the T first eigenvalues are also positive and simple. If for some $t \leq T$ we have $\|B\|_{\text{HS}(\mathcal{H})} \leq \widetilde{\delta}_t/2$, then*

$$\|\Pi_t(A) - \Pi_t(A + B)\|_{\text{HS}(\mathcal{H})} \leq \frac{2\|B\|_{\text{HS}(\mathcal{H})}}{\widetilde{\delta}_t},$$

where $\Pi_t(A)$ denotes the orthogonal projector onto the one-dimensional subspace of \mathcal{H} spanned by the t^{th} eigenfunction of A .

E.3 Results on Orthogonal Projectors

Lemma 8. *Let $\mathbf{P}_{\mathbf{X}} = (\pi_{i,j})_{i,j \in \{1, \dots, n\}} \in \mathcal{M}_n(\mathbb{R})$ be the matrix of an orthogonal projector of rank p , then we have:*

$$\begin{aligned} \sum_{i=1}^n \sum_{j=1}^n \pi_{i,j}^2 &= p, \\ \sum_{i=1}^n \pi_{i,i}^2 &\leq p, \\ \sum_{i=1}^n (1 - \pi_{i,i})^2 &\leq n - p. \end{aligned}$$

Proof. As $\mathbf{P}'_{\mathbf{X}}\mathbf{P}'_{\mathbf{X}} = \mathbf{P}_{\mathbf{X}}$, we directly have $\sum_{j=1}^n \pi_{i,j}^2 = \pi_{i,i}$ and by computing the trace we find that:

$$\sum_{i=1}^n \sum_{j=1}^n \pi_{i,j}^2 = p.$$

Then:

$$\begin{aligned} \sum_{i=1}^n \pi_{i,i}^2 &\leq \sum_{i=1}^n \sum_{j=1}^n \pi_{i,j}^2 \\ &\leq p, \\ \sum_{i=1}^n (1 - \pi_{i,i})^2 &= \sum_{i=1}^n (1 - 2\pi_{i,i} + \pi_{i,i}^2) \\ &\leq n - p. \end{aligned}$$

□

F Kernel tricks

In this section, we apply kernel tricks to compute explicit expressions for all the quantities presented in the paper.

F.1 Kernel trick for the truncated kernel Hotelling-Lawley statistic

F.1.1 Diagonalization of the residual covariance operator.

We start with the determination of the eigenfunction and associated eigenvalues of the residual covariance operator $\widehat{\Sigma}$. Observe that we have $\widehat{\Sigma} = n^{-1}\widehat{\mathbf{E}}^*\widehat{\mathbf{E}} = n^{-1}\sum_{i=1}^n \widehat{e}_i \otimes \widehat{e}_i \in \text{HS}(\mathcal{H})$, where $\widehat{\mathbf{E}} = \mathbf{P}_{\mathbf{X}}^{\perp}\Phi(\mathbf{Y})$. We then define $\mathbf{K}_e = n^{-1}\widehat{\mathbf{E}}\widehat{\mathbf{E}}^* = \left(n^{-1}\langle \widehat{e}_i, \widehat{e}_j \rangle_{\mathcal{H}}\right)_{i,j \in \{1, \dots, n\}} \in \mathcal{M}_n(\mathbb{R})$. This matrix can be numerically computed and diagonalized, as we have :

$$\mathbf{K}_e = \frac{1}{n}\mathbf{P}_{\mathbf{X}}^{\perp}\mathbf{K}_{\mathbf{Y}}\mathbf{P}_{\mathbf{X}}^{\perp},$$

where $\mathbf{K}_{\mathbf{Y}} = (k(Y_i, Y_j))_{i,j \in \{1, \dots, n\}}$ is the gram matrix of \mathbf{Y} with respect to $k(\cdot, \cdot)$. The following proposition highlights the link between $\widehat{\Sigma}$ and \mathbf{K}_e .

Proposition 6. *The Hilbert-Schmidt operator $\widehat{\Sigma}$ has the same spectrum than the matrix \mathbf{K}_e . Moreover, if u is a unit eigenvector of \mathbf{K}_e associated to the eigenvalue $\widehat{\lambda}$, then a unit eigenfunction \widehat{f} of $\widehat{\Sigma}$ associated to $\widehat{\lambda}$ may be obtained by :*

$$\widehat{f} = \frac{1}{\sqrt{n\widehat{\lambda}}}u'\widehat{\mathbf{E}} \quad (25)$$

Proof. Observe that $n^{-1}\widehat{\mathbf{E}}\widehat{\mathbf{E}}^*\widehat{\mathbf{E}} \in \mathcal{H}^n$ can be read as $n^{-1}(\widehat{\mathbf{E}}\widehat{\mathbf{E}}^*)\widehat{\mathbf{E}} = \mathbf{K}_e\widehat{\mathbf{E}}$ as well as $n^{-1}\widehat{\mathbf{E}}(\widehat{\mathbf{E}}^*\widehat{\mathbf{E}}) = \widehat{\mathbf{E}}\widehat{\Sigma}$, where $\widehat{\mathbf{E}}$ is interpreted as an application defined as in (8). Thus we have the relation

$$\mathbf{K}_e\widehat{\mathbf{E}} = \widehat{\mathbf{E}}\widehat{\Sigma} \quad (26)$$

. Then, if \widehat{f} is a unit eigenfunction of $\widehat{\Sigma}$ associated to the eigenvalue $\widehat{\lambda}$, we have:

$$\begin{aligned} \widehat{\Sigma}\widehat{f} &= \widehat{\lambda}\widehat{f} \\ \Leftrightarrow \widehat{\mathbf{E}}\widehat{\Sigma}\widehat{f} &= \widehat{\lambda}\widehat{\mathbf{E}}\widehat{f} \\ \Leftrightarrow \mathbf{K}_e(\widehat{\mathbf{E}}\widehat{f}) &= \widehat{\lambda}(\widehat{\mathbf{E}}\widehat{f}) \end{aligned}$$

Thus the vector $\widehat{\mathbf{E}}\widehat{f} \in \mathbb{R}^n$ is an eigenvector of \mathbf{K}_e associated to the eigenvalue $\widehat{\lambda}$, and the spectrum of $\widehat{\Sigma}$ is included in the spectrum of \mathbf{K}_e . On the other hand, let u be a unit eigenvector of \mathbf{K}_e associated to the eigenvalue $\widehat{\lambda}$. We have:

$$\begin{aligned} u' \mathbf{K}_e &= \widehat{\lambda} u' \\ \Leftrightarrow u' \mathbf{K}_e \widehat{\mathbf{E}} &= \widehat{\lambda} u' \widehat{\mathbf{E}} \\ \Leftrightarrow u' (\widehat{\mathbf{E}} \widehat{\Sigma}) &= \widehat{\lambda} u' \widehat{\mathbf{E}} \\ \Leftrightarrow \widehat{\Sigma} (u' \widehat{\mathbf{E}}) &= \widehat{\lambda} (u' \widehat{\mathbf{E}}), \end{aligned}$$

where the last equivalence is due to $u' (\widehat{\mathbf{E}} \widehat{\Sigma}) = u' (\widehat{\Sigma} \widehat{e}_i)_{i=1, \dots, n} = \widehat{\Sigma} \sum_{i=1}^n u_i \widehat{e}_i = \widehat{\Sigma} (u' \widehat{\mathbf{E}})$.

Thus the vector $u' \widehat{\mathbf{E}} \in \mathcal{H}$ is an eigenvector of $\widehat{\Sigma}$ associated to the eigenvalue $\widehat{\lambda}$, and the spectrum of \mathbf{K}_e is included in the spectrum of $\widehat{\Sigma}$. We conclude that $\widehat{\Sigma}$ and \mathbf{K}_e share the same spectrum.

Let u be a unit eigenvector of \mathbf{K}_e associated to the eigenvalue $\widehat{\lambda}$. We have that:

$$\begin{aligned} \|u' \widehat{\mathbf{E}}\|_{\mathcal{H}}^2 &= \left\langle \sum_{i=1}^n u_i \widehat{e}_i, \sum_{j=1}^n u_j \widehat{e}_j \right\rangle_{\mathcal{H}} \\ &= \sum_{i=1}^n \sum_{j=1}^n u_i \langle \widehat{e}_i, \widehat{e}_j \rangle_{\mathcal{H}} u_j \\ &= n u' \mathbf{K}_e u \\ &= n \widehat{\lambda} u' u \\ &= n \widehat{\lambda} \end{aligned}$$

Thus, $(n \widehat{\lambda})^{-\frac{1}{2}} u' \widehat{\mathbf{E}} \in \mathcal{H}$ is a unit eigenfunction of $\widehat{\Sigma}$. \square

F.1.2 Computation of the truncated kernel Hotelling-Lawley trace statistic

Let $\mathbf{U} = (u_1, \dots, u_T) \in \mathcal{M}_{n, T}(\mathbb{R})$ and $\widehat{\Lambda} = \text{diag}(\widehat{\lambda}_1, \dots, \widehat{\lambda}_T) \in \mathcal{M}_T(\mathbb{R})$ be the matrix containing the T first eigenvectors and eigenvalues of \mathbf{K}_e respectively, and let $\mathbf{D} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{L}'(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{L}')^{-1} \mathbf{L}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}' = (d_{i,j})_{i,j \in \{1, \dots, n\}} \in \mathcal{M}_n(\mathbb{R})$. \mathbf{D} is an orthogonal projector, as $\mathbf{D}' = \mathbf{D}$ and $\mathbf{D}^2 = \mathbf{D}$ and we have $\widehat{H}_{\mathbf{L}} = \Phi(\mathbf{Y})^* \mathbf{D} \Phi(\mathbf{Y}) = \sum_{i,j=1}^n d_{i,j} \phi(Y_i) \otimes \phi(Y_j)$. Then :

$$\begin{aligned} \widehat{F}_T &= \text{tr}(\widehat{\Sigma}_T^{-1} \widehat{H}_{\mathbf{L}}) \\ &= \left\langle \widehat{\Sigma}_T^{-1}, \widehat{H}_{\mathbf{L}} \right\rangle_{\text{HS}(\mathcal{H})} \\ &= \sum_{i,j=1}^n \sum_{t=1}^T d_{i,j} \widehat{\lambda}_t^{-1} \left\langle \widehat{f}_t \otimes \widehat{f}_t, \phi(Y_i) \otimes \phi(Y_j) \right\rangle_{\text{HS}(\mathcal{H})} \\ &= \sum_{i,j=1}^n \sum_{t=1}^T d_{i,j} \widehat{\lambda}_t^{-1} \left\langle \widehat{f}_t, \phi(Y_i) \right\rangle_{\mathcal{H}} \left\langle \widehat{f}_t, \phi(Y_j) \right\rangle_{\mathcal{H}} \\ &= \text{trace}(\mathbf{K}_T \mathbf{D} \mathbf{K}_T'), \end{aligned} \tag{27}$$

where $\mathbf{K}_T = n^{-\frac{1}{2}} \widehat{\Lambda}^{-1} \mathbf{U}' \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}} = \left(\widehat{\lambda}_t^{-\frac{1}{2}} \left\langle \widehat{f}_t, \phi(Y_i) \right\rangle_{\mathcal{H}} \right)_{t \in \{1, \dots, T\}, i \in \{1, \dots, n\}} \in \mathcal{M}_{T, n}(\mathbb{R})$. The following proposition shows the relation between $\mathbf{K}_T \mathbf{D} \mathbf{K}_T'$ and $\widehat{\Sigma}_T^{-1} \widehat{H}_{\mathbf{L}}$. Ψ

Proposition 7. *The Hilbert-Schmidt operator $\widehat{\Sigma}_T^{-1} \widehat{H}_{\mathbf{L}}$ has the same spectrum than the matrix $\mathbf{K}_T \mathbf{D} \mathbf{K}_T'$. Moreover, if v is a unit eigenvector of $\mathbf{K}_T \mathbf{D} \mathbf{K}_T'$ associated to the eigenvalue $\widehat{\lambda}$, then an unit eigenfunction \widehat{g} of $\widehat{\Sigma}_T^{-1} \widehat{H}_{\mathbf{L}}$ associated to $\widehat{\lambda}$ may be obtained by :*

$$\widehat{g} = \frac{1}{\sqrt{\sum_{t=1}^T \widehat{\lambda}_t^{-1} v_t}} \Psi^* v, \tag{28}$$

where $\Psi = n^{-\frac{1}{2}}\widehat{\Lambda}^{-1}\mathbf{U}'\mathbf{P}_{\mathbf{X}}^{\perp}\Phi(\mathbf{Y})$

Proof. Note that we have $\mathbf{K}_T = \Psi\Phi(\mathbf{Y})^*$ and

$$\begin{aligned}\widehat{\Sigma}_T^{-1} &= \frac{1}{n}\Phi(\mathbf{Y})^*\mathbf{P}_{\mathbf{X}}^{\perp}\mathbf{U}\widehat{\Lambda}^{-2}\mathbf{U}'\mathbf{P}_{\mathbf{X}}^{\perp}\Phi(\mathbf{Y}) \\ &= \Psi^*\Psi.\end{aligned}$$

Thus, $\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}} = (\Psi^*\Psi)(\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y}))$ and $\mathbf{K}_T\mathbf{D}\mathbf{K}'_T = \Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\Psi^*$. Observe that similarly to (26), we have that:

$$\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}\Psi^* = \Psi^*\mathbf{K}_T\mathbf{D}\mathbf{K}'_T,$$

that :

$$\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}} = \mathbf{K}_T\mathbf{D}\mathbf{K}'_T\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y}),$$

and that :

$$\begin{aligned}\Psi\Psi^* &= \widehat{\Lambda}^{-\frac{1}{2}}\left(\left\langle\widehat{f}_t, \widehat{f}_{t'}\right\rangle_{\mathcal{H}}\right)_{t, t' \in \{1, \dots, T\}}\widehat{\Lambda}^{-\frac{1}{2}} \\ &= \widehat{\Lambda}^{-1}\end{aligned}$$

Then, if \widehat{g} is an unit eigenfunction of $\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}$ associated to the eigenvalue ξ , we have:

$$\begin{aligned}\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}\widehat{g} &= \xi\widehat{g} \\ \Leftrightarrow \Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}\widehat{g} &= \xi\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\widehat{g} \\ \Leftrightarrow \mathbf{K}_T\mathbf{D}\mathbf{K}'_T\left(\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\widehat{g}\right) &= \xi\left(\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\widehat{g}\right)\end{aligned}$$

Thus the vector $\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\widehat{g} \in \mathbb{R}^n$ is an eigenvector of $\mathbf{K}_T\mathbf{D}\mathbf{K}'_T$ associated to the eigenvalue ξ , and the spectrum of $\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}$ is included in the spectrum of $\mathbf{K}_T\mathbf{D}\mathbf{K}'_T$. On the other hand, let v be an unit eigenvector of $\mathbf{K}_T\mathbf{D}\mathbf{K}'_T$ associated to the eigenvalue ξ . We have:

$$\begin{aligned}v'\mathbf{K}_T\mathbf{D}\mathbf{K}'_T &= \xi v' \\ \Leftrightarrow v'\mathbf{K}_T\mathbf{D}\mathbf{K}'_T\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y}) &= \xi v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y}) \\ \Leftrightarrow v'\left(\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}\right) &= \xi\left(v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\right) \\ \Leftrightarrow \widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}\left(v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\right) &= \xi\left(v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\right)\end{aligned}$$

See the proof of Proposition 6 for the details of the last equivalence. Thus the vector $v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y}) \in \mathcal{H}$ is an eigenvector of $\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}$ associated to the eigenvalue ξ , and the spectrum of $\mathbf{K}_T\mathbf{D}\mathbf{K}'_T$ is included in the spectrum of $\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}$. We conclude that $\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}$ and $\mathbf{K}_T\mathbf{D}\mathbf{K}'_T$ share the same spectrum. Let v be an unit eigenvector of $\mathbf{K}_T\mathbf{D}\mathbf{K}'_T$ associated to the eigenvalue ξ . We have that:

$$\begin{aligned}\|v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\|_{\mathcal{H}}^2 &= v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y})\Psi^*v \\ &= v'\mathbf{K}_T\mathbf{D}\mathbf{K}_{\mathbf{Y}}\mathbf{D}\mathbf{K}'_Tv\end{aligned}$$

Thus, $(v'\mathbf{K}_T\mathbf{D}\mathbf{K}_{\mathbf{Y}}\mathbf{D}\mathbf{K}'_T\widehat{g})^{-\frac{1}{2}}v'\Psi\Phi(\mathbf{Y})^*\mathbf{D}\Phi(\mathbf{Y}) \in \mathcal{H}$ is an unit eigenfunction of $\widehat{\Sigma}_T^{-1}\widehat{H}_{\mathbf{L}}$. □

F2 Kernel trick for the Nyström statistic

F2.1 Computation of the landmarks and anchors

To compute the Nystrom approximation of the truncated kernel Hotelling-Lawley trace statistic (5), e need a kernel trick to determine the eigenfunctions and eigenvalues of $\widehat{\Sigma}^a$.

Let $\mathbf{Z} = (Z_1, \dots, Z_q)$ be the $q < n$ landmarks sampled from \mathbf{Y} . We denote $\varphi : \{1, \dots, q\} \rightarrow \{1, \dots, n\}$ such that $\mathbf{Z} = (Y_{\varphi(1)}, \dots, Y_{\varphi(q)})$ and the associated explanatory variables are $\mathbf{X}^I = (x_{\varphi(1)}, \dots, x_{\varphi(q)})' \in \mathcal{M}_{q,p}(\mathbb{R})$. Let $P_{\mathbf{Z}} \in \mathcal{M}_{q,n}(\mathbb{R})$ such that the i^{th} column of $P_{\mathbf{Z}}'$ is the vector of \mathbb{R}^n full of zero except in $\varphi(i)$ where it is equal to 1, then we have $\mathbf{Z} = P_{\mathbf{Z}}\mathbf{Y}$ and $\mathbf{X}^I = P_{\mathbf{Z}}\mathbf{X}$. Let $\Phi(\mathbf{Z}) = (\phi(Z_1), \dots, \phi(Z_n))$ be the embeddings of the landmarks in \mathcal{H} . We define the Gram matrices $\mathbf{K}_{\mathbf{Z}} = (\langle \phi(Z_i), \phi(Z_j) \rangle_{\mathcal{H}})_{i,j \in \{1, \dots, q\}} \in \mathcal{M}_q(\mathbb{R})$, $\mathbf{K}_{\mathbf{Y}, \mathbf{Z}} = (\langle \phi(Y_i), \phi(Z_j) \rangle_{\mathcal{H}})_{i \in \{1, \dots, n\}, j \in \{1, \dots, q\}} \in \mathcal{M}_{n,q}(\mathbb{R})$ and $\mathbf{K}_{\mathbf{Z}, \mathbf{Y}} = \mathbf{K}_{\mathbf{Y}, \mathbf{Z}}'$. We have $\Phi(\mathbf{Z}) = P_{\mathbf{Z}}\Phi(\mathbf{Y})$, $\mathbf{K}_{\mathbf{Z}} = \Phi(\mathbf{Z})\Phi(\mathbf{Z})^*$, $\mathbf{K}_{\mathbf{Y}, \mathbf{Z}} = \Phi(\mathbf{Y})\Phi(\mathbf{Z})^*$ and $\mathbf{K}_{\mathbf{Z}, \mathbf{Y}} = \Phi(\mathbf{Z})\Phi(\mathbf{Y})^*$. We assume that \mathbf{X}^I is full rank and denote $\mathbf{P}_{\mathbf{X}^I} = \mathbf{X}^I(\mathbf{X}^{I'}\mathbf{X}^I)^{-1}\mathbf{X}^{I'}$ and $\mathbf{P}_{\mathbf{X}^I}^{\perp} = (\mathbf{I}_q - \mathbf{P}_{\mathbf{X}^I})$, then we have the vector of landmark residuals $\widehat{E}^{\mathbf{Z}} = (\widehat{e}_1^{\mathbf{Z}}, \dots, \widehat{e}_q^{\mathbf{Z}})$ in \mathcal{H}^n such that :

$$\widehat{E}^{\mathbf{Z}} = \mathbf{P}_{\mathbf{X}^I}^{\perp} \Phi(\mathbf{Z}) \quad (29)$$

The landmark residual covariance operator is such that $\widehat{\Sigma}^I = q^{-1} \widehat{E}^{\mathbf{Z}} \widehat{E}^{\mathbf{Z}*}$. The m anchors $\mathbf{a} = (\widehat{a}_1, \dots, \widehat{a}_m)'$ are defined as the orthonormal set of eigenfunctions of $\widehat{\Sigma}^I$ associated to its highest and non-increasing set of eigenvalues encoded in the diagonal matrix $\widehat{\Lambda}^{\mathbf{Z}} = \text{Diag}(\widehat{\lambda}_1^{\mathbf{Z}}, \dots, \widehat{\lambda}_m^{\mathbf{Z}}) \in \mathcal{M}_m(\mathbb{R})$. We also define $\mathbf{K}_e^{\mathbf{Z}} = q^{-1} \widehat{E}^{\mathbf{Z}} \widehat{E}^{\mathbf{Z}*} = q^{-1} \mathbf{P}_{\mathbf{X}^I}^{\perp} \mathbf{K}_{\mathbf{Z}} \mathbf{P}_{\mathbf{X}^I}^{\perp}$. The following proposition highlights the link between $\widehat{\Sigma}^I$ and $\mathbf{K}_e^{\mathbf{Z}}$:

Proposition 8. *The Hilbert-Schmidt operator $\widehat{\Sigma}^I$ has the same spectrum than the matrix $\mathbf{K}_e^{\mathbf{Z}}$. Moreover, if $u^{\mathbf{Z}}$ is a unit eigenvector of $\mathbf{K}_e^{\mathbf{Z}}$ associated to the eigenvalue $\widehat{\lambda}^{\mathbf{Z}}$, then a unit eigenfunction $\widehat{f}^{\mathbf{Z}}$ of $\widehat{\Sigma}^I$ associated to $\widehat{\lambda}^{\mathbf{Z}}$ may be obtained by :*

$$\widehat{f}^{\mathbf{Z}} = \frac{1}{\sqrt{q \widehat{\lambda}^{\mathbf{Z}}}} u^{\mathbf{Z}'} \widehat{E}^{\mathbf{Z}} \quad (30)$$

Proof. The proof is exactly the same than the proof of Proposition 6. \square

We denote $\mathbf{U}^{\mathbf{Z}} = (u_1^{\mathbf{Z}}, \dots, u_m^{\mathbf{Z}}) \in \mathcal{M}_{q,m}(\mathbb{R})$ the matrix where the m columns form an orthonormal set of eigenvectors of $\mathbf{K}_e^{\mathbf{Z}}$ associated to the eigenvalues $\widehat{\lambda}_1^{\mathbf{Z}}, \dots, \widehat{\lambda}_m^{\mathbf{Z}}$, then for $i \in \{1, \dots, m\}$, we have :

$$\mathbf{a} = \frac{1}{\sqrt{q}} \widehat{\Lambda}^{\mathbf{Z}}^{-\frac{1}{2}} \mathbf{U}^{\mathbf{Z}'} \widehat{E}^{\mathbf{Z}}. \quad (31)$$

F.2.2 Diagonalisation of the Nyström residual covariance operator

We now derive a kernel trick to determine the orthonormal eigenfunctions $\widehat{f}_1^{\mathbf{a}}, \dots, \widehat{f}_m^{\mathbf{a}}$ of $\widehat{\Sigma}^{\mathbf{a}}$ associated to the eigenvalues $\widehat{\lambda}_1^{\mathbf{a}}, \dots, \widehat{\lambda}_m^{\mathbf{a}}$. Let $\widehat{E}^{\mathbf{a}} = (\widehat{e}_1^{\mathbf{a}}, \dots, \widehat{e}_n^{\mathbf{a}})$ be the vector of \mathcal{H}^n containing the projections of the residuals onto $\text{Span}(\widehat{a}_1, \dots, \widehat{a}_m)$. We have $\widehat{e}_i^{\mathbf{a}} = \sum_{s=1}^m \langle \widehat{a}_s, \widehat{e}_i \rangle_{\mathcal{H}} \widehat{a}_s$ and $\widehat{\Sigma}^{\mathbf{a}} = \frac{1}{n} \sum_{i=1}^n \widehat{e}_i^{\mathbf{a}} \otimes \widehat{e}_i^{\mathbf{a}}$. We define $\mathbf{K}_e^{\mathbf{a}} \in \mathcal{M}_m(\mathbb{R})$ with general term $\langle \widehat{a}_s, \widehat{\Sigma} \widehat{a}_{s'} \rangle_{\mathcal{H}}$. The following proposition highlights the link between $\widehat{\Sigma}^{\mathbf{a}}$ and $\mathbf{K}_e^{\mathbf{a}}$.

Proposition 9. *The Hilbert-Schmidt operator $\widehat{\Sigma}^{\mathbf{a}}$ has the same spectrum than the matrix $\mathbf{K}_e^{\mathbf{a}}$. Moreover, if $u^{\mathbf{a}}$ is a unit eigenvector of $\mathbf{K}_e^{\mathbf{a}}$ associated to the eigenvalue $\widehat{\lambda}^{\mathbf{a}}$, then a unit eigenfunction $\widehat{f}^{\mathbf{a}}$ of $\widehat{\Sigma}^{\mathbf{a}}$ associated to $\widehat{\lambda}^{\mathbf{a}}$ may be obtained by :*

$$\widehat{f}^{\mathbf{a}} = u^{\mathbf{a}'} \mathbf{a} \quad (32)$$

Proof. Observe that:

$$\begin{aligned} \widehat{\Sigma}^{\mathbf{a}} &= \frac{1}{n} \sum_{i=1}^n \left(\sum_{s=1}^m \langle \widehat{a}_s, \widehat{e}_i \rangle_{\mathcal{H}} \widehat{a}_s \right) \otimes \left(\sum_{s'=1}^m \langle \widehat{a}_{s'}, \widehat{e}_i \rangle_{\mathcal{H}} \widehat{a}_{s'} \right) \\ &= \frac{1}{n} \sum_{s,s'=1}^m \left\langle \widehat{a}_s, \left(\sum_{i=1}^n \widehat{e}_i \otimes \widehat{e}_i \right) \widehat{a}_{s'} \right\rangle_{\mathcal{H}} \widehat{a}_s \otimes \widehat{a}_{s'} \\ &= \widehat{\mathbf{a}}^* \mathbf{K}_e^{\mathbf{a}} \widehat{\mathbf{a}}. \end{aligned}$$

As the anchors $(\widehat{a}_1, \dots, \widehat{a}_m)$ form an orthonormal basis of \mathcal{H}_a , we have $\mathbf{a}\mathbf{a}^* = (\langle \widehat{a}_i, \widehat{a}_j \rangle_{\mathcal{H}})_{i,j \in \{1, \dots, m\}} = \mathbf{I}_m \in \mathcal{M}_m(\mathbb{R})$. Thanks to this relation, we directly have that $\mathbf{a}\widehat{\Sigma}^a = \mathbf{K}_e^a \mathbf{a}$. Let u^a be the eigenvector of \mathbf{K}_e^a associated to the non-zero eigenvalue $\widehat{\lambda}^a$. We have $u^{a'}(\mathbf{K}_e^a \mathbf{a}) = u^a(\mathbf{a}\widehat{\Sigma}^a) = \widehat{\Sigma}^a(u^a \mathbf{a})$, where the last equality is obtained similarly to the proof of Proposition 6. We also have $(u^{a'} \mathbf{K}_e^a) \mathbf{a} = \widehat{\lambda}^a u^{a'} \mathbf{a}$, thus $u^{a'} \mathbf{a}$ is an eigenfunction of $\widehat{\Sigma}^a$ associated to the eigenvalue $\widehat{\lambda}^a$, thus the non-zero eigenvalues of \mathbf{K}_e^a are included in the set of non-zero eigenvalues of $\widehat{\Sigma}^a$. Now let \widehat{f}^a be an eigenfunction of $\widehat{\Sigma}^a$ associated to the non-zero eigenvalue $\widehat{\lambda}^a$. As $\widehat{\Sigma}^a$ is self-adjoint, for $i \in \{1, \dots, n\}$ we have $\left((\mathbf{a}\widehat{\Sigma}^a)\widehat{f}^a \right)_i = \langle \widehat{\Sigma}^a \widehat{a}_i, \widehat{f}^a \rangle_{\mathcal{H}} = \langle \widehat{a}_i, \widehat{\Sigma}^a \widehat{f}^a \rangle_{\mathcal{H}}$. Thus, we have $\mathbf{K}_e^a \mathbf{a}\widehat{f}^a = \mathbf{a}\widehat{\Sigma}^a \widehat{f}^a = \widehat{\lambda}^a \mathbf{a}\widehat{f}^a$. Then $\mathbf{a}\widehat{f}^a \in \mathbb{R}^m$ is an eigenvector of \mathbf{K}_e^a associated to the eigenvalue $\widehat{\lambda}^a$, thus the non-zero eigenvalues of $\widehat{\Sigma}^a$ are included in the set of non-zero eigenvalues of \mathbf{K}_e^a . Thus, $\widehat{\Sigma}^a$ and \mathbf{K}_e^a share the same spectrum. As we have $\|u^{a'} \mathbf{a}\|_{\mathcal{H}} = 1$, then $\widehat{f}^a = u^{a'} \mathbf{a}$. \square

In practice, the full expression of \mathbf{K}_e^a is such that :

$$\mathbf{K}_e^a = \frac{1}{nq} \widehat{\Lambda}^{\mathbf{Z}}^{-\frac{1}{2}} \mathbf{U}^{\mathbf{Z}'} \mathbf{P}_{\mathbf{X}'}^{\perp} \mathbf{K}_{\mathbf{Z}, \mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}, \mathbf{Z}} \mathbf{P}_{\mathbf{X}'}^{\perp} \mathbf{U}^{\mathbf{Z}} \widehat{\Lambda}^{\mathbf{Z}}^{-\frac{1}{2}}$$

F.2.3 Computation of the Nystrom statistic

Let $\mathbf{U}^a = (u_1^a, \dots, u_T^a) \in \mathcal{M}_{m, T}(\mathbb{R})$ and $\widehat{\Lambda}^a = \text{diag}(\widehat{\lambda}_1^a, \dots, \widehat{\lambda}_T^a) \in \mathcal{M}_T(\mathbb{R})$ be the matrix containing the T first eigenvectors and eigenvalues of \mathbf{K}_e^a respectively. The Nystrom statistic $\widehat{\mathcal{F}}_T^{\text{NY}}$ can be written with respect to \mathbf{U}^a , $\widehat{\Lambda}^a$ and \mathbf{D} . With the same development than in Equation (27), we have :

$$\widehat{\mathcal{F}}_T^{\text{NY}} = \text{trace}(\mathbf{K}_T^a \mathbf{D} \mathbf{K}_T^a),$$

where $\mathbf{K}_T^a = (q\widehat{\Lambda}^a)^{-\frac{1}{2}} \mathbf{U}^{a'} \widehat{\Lambda}^{\mathbf{Z}}^{-\frac{1}{2}} \mathbf{U}^{\mathbf{Z}'} \mathbf{P}_{\mathbf{X}'}^{\perp} \mathbf{K}_{\mathbf{Z}, \mathbf{Y}} = \left(\widehat{\lambda}_t^a^{-\frac{1}{2}} \left\langle \widehat{f}_t^a, \phi(Y_i) \right\rangle_{\mathcal{H}} \right)_{t \in \{1, \dots, T\}, i \in \{1, \dots, n\}} \in \mathcal{M}_{T, n}(\mathbb{R})$.

F.3 Kernel trick for the diagnostics and kernel Cook's distance

F.3.1 Computation of the Diagnostic Plots

For $t \in \{1, \dots, T\}$, the t^{th} kernel response plot and t^{th} kernel residual plot represent the values of the t^{th} column of the matrix $\frac{1}{\sqrt{n}} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{U} \widehat{\Lambda}^{-\frac{1}{2}} = \left(\left\langle \phi(Y_i), \widehat{f}_t \right\rangle_{\mathcal{H}} \right)_{i \in \{1, \dots, n\}, t \in \{1, \dots, T\}} \in \mathcal{M}_{n, T}(\mathbb{R})$ and the values of the t^{th} column of the matrix $\frac{1}{\sqrt{n}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{U} \widehat{\Lambda}^{-\frac{1}{2}} = \left(\left\langle \widehat{e}_i, \widehat{f}_t \right\rangle_{\mathcal{H}} \right)_{i \in \{1, \dots, n\}, t \in \{1, \dots, T\}}$ respectively, both with respect to the values of the t^{th} column of the matrix $\frac{1}{\sqrt{n}} \mathbf{P}_{\mathbf{X}} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{U} \widehat{\Lambda}^{-\frac{1}{2}} = \left(\left\langle \widehat{\phi}(Y_i), \widehat{f}_t \right\rangle_{\mathcal{H}} \right)_{i \in \{1, \dots, n\}, t \in \{1, \dots, T\}} \in \mathcal{M}_{n, T}(\mathbb{R})$.

F.3.2 Computation of the kernel Cook's distance

In order to generalize the Cook's distance to our linear model (1) in the feature space \mathcal{H} , we need to define the expectation and covariance of vectors of \mathcal{H}^d in order to compute $\mathbb{E}(\mathbf{L}\widehat{\Theta})$ and $\text{Cov}(\mathbf{L}\widehat{\Theta})$. Let $\mathbf{g} = (g_1, \dots, g_d)$ a random vector of \mathcal{H}^d . We define the expectation of \mathbf{g} as $\mathbb{E}(\mathbf{g}) = (\mathbb{E}(g_1), \dots, \mathbb{E}(g_d)) \in \mathcal{H}^d$. We define the covariance of \mathbf{g} as a $d \times d$ matrix where the coordinates are elements of $\text{HS}(\mathcal{H})$, such that for $i, j \in \{1, \dots, d\}$, we have $\text{Cov}(\mathbf{g})_{i, j} = \text{Cov}(g_i, g_j) = \mathbb{E}((g_i - \mathbb{E}(g_i)) \otimes (g_j - \mathbb{E}(g_j)))$. Now we compute the expectation $\mathbb{E}(\mathbf{L}\widehat{\Theta})$ and covariance $\text{Cov}(\mathbf{L}\widehat{\Theta})$ of $\mathbf{L}\widehat{\Theta}$. Note that the combination of Equations (1) and (2) gives $\mathbf{L}\widehat{\Theta} = \mathbf{L}\Theta + \mathbf{L}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{E}$. Let $j \in \{1, \dots, d\}$, we have $(\mathbf{L}\widehat{\Theta})_j = (\mathbf{L}\Theta)_j + w'_j \mathbf{E}$, where $w_j \in \mathbb{R}^n$ is the j^{th} column of $\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{L}'$. Then we have $\mathbb{E}((\mathbf{L}\widehat{\Theta})_j) = (\mathbf{L}\Theta)_j$. Let $i, j \in \{1, \dots, d\}$, we

have:

$$\begin{aligned}
\text{Cov} \left((\mathbf{L}\widehat{\Theta})_i, (\mathbf{L}\widehat{\Theta})_j \right) &= \mathbb{E} \left(((\mathbf{L}\widehat{\Theta})_i - (\mathbf{L}\Theta)_i) \otimes ((\mathbf{L}\widehat{\Theta})_j - (\mathbf{L}\Theta)_j) \right) \\
&= \mathbb{E} \left((w'_i \mathbf{E}) \otimes (w'_j \mathbf{E}) \right) \\
&= \sum_{k,l=1}^n w_{i,k} w_{j,l} \mathbb{E} (e_k \otimes e_l) \\
&= \sum_{k=1}^n w_{i,k} w_{j,k} \Sigma_k \\
&= \sum_{k=1}^n w_{i,k} w_{j,k} \Sigma_T + \sum_{k=1}^n w_{i,k} w_{j,k} \Sigma_{k,T^c} \\
&= (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')_{i,j} \Sigma_T + \sum_{k=1}^n w_{i,k} w_{j,k} \Sigma_{k,T^c}.
\end{aligned}$$

We choose to focus on the truncated covariance $\text{Cov}_T(\mathbf{L}\widehat{\Theta})$ of general term $(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')_{i,j} \Sigma_T$ approximated through $\widehat{\text{Cov}}_T(\mathbf{L}\widehat{\Theta})$ of general term $(\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')_{i,j} \widehat{\Sigma}_T$. Let $i \in \{1, \dots, n\}$, the RKHS version of the expression expression of $\mathbf{L}(\widehat{\Theta} - \widehat{\Theta}_{(i)})$ is such that (see [5] for the multivariate version):

$$\mathbf{L}(\widehat{\Theta} - \widehat{\Theta}_{(i)}) = \begin{pmatrix} \frac{w_{1,i}}{1-\pi_{i,i}} \widehat{e}_i \\ \vdots \\ \frac{w_{d,i}}{1-\pi_{i,i}} \widehat{e}_i \end{pmatrix} \in \mathcal{H}^d$$

Then we have that:

$$\begin{aligned}
\mathcal{D}_{\mathbf{L}} &= \frac{1}{d} \text{tr} \left(\sum_{j,k=1}^d \frac{w_{j,i}}{1-\pi_{i,i}} \frac{w_{k,i}}{1-\pi_{i,i}} (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')_{j,k}^{-1} \widehat{e}_i \otimes \widehat{\Sigma}_T \widehat{e}_i \right) \\
&= \frac{1}{d} \sum_{j,k=1}^d \frac{w_{j,i}}{1-\pi_{i,i}} \frac{w_{k,i}}{1-\pi_{i,i}} (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')_{j,k}^{-1} \text{tr} \left(\widehat{e}_i \otimes \widehat{\Sigma}_T \widehat{e}_i \right)
\end{aligned}$$

We remark that:

$$\sum_{j,k=1}^d \frac{w_{j,i}}{1-\pi_{i,i}} \frac{w_{k,i}}{1-\pi_{i,i}} (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')_{j,k}^{-1} = \frac{w'_i (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')^{-1} w_i}{(1-\pi_{i,i})^2}$$

And we have:

$$\text{tr} \left(\widehat{e}_i \otimes \widehat{\Sigma}_T \widehat{e}_i \right) = \sum_{t=1}^T \widehat{\lambda}_t^{-1} \left\langle \widehat{f}_t, \widehat{e}_i \right\rangle_{\mathcal{H}}^2.$$

We recognize the inner products contained in the matrix $\frac{1}{\sqrt{n}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{U} \widehat{\Lambda}^{-\frac{1}{2}}$, thus we have:

$$\text{tr} \left(\widehat{e}_i \otimes \widehat{\Sigma}_T \widehat{e}_i \right) = \frac{1}{n} (\mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{U} \widehat{\Lambda}^{-2} \mathbf{U}' \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp})_{i,i}$$

Finally, we have an expression for the Cook distance $\mathcal{D}_{\mathbf{L}}$ associated to the i^{th} observation:

$$\frac{w'_i (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{L}')^{-1} w_i}{dn(1-\pi_{i,i})^2} (\mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{U} \widehat{\Lambda}^{-2} \mathbf{U}' \mathbf{P}_{\mathbf{X}}^{\perp} \mathbf{K}_{\mathbf{Y}} \mathbf{P}_{\mathbf{X}}^{\perp})_{i,i}.$$

F.4 Computational costs

Computational cost of the TKHL statistic We compute the projection matrix $\mathbf{P}_{\mathbf{X}}^{\perp}$ for $O(p^3)$ operations. We compute and diagonalize the matrix \mathbf{K}_e for $O(n^3)$ operations. Then we compute the matrix \mathbf{K}_T and the TKHL statistic for $O(T^2n + n^2T)$ operations. As $n \gg p$ and $n \gg T$, the global computational cost of the algorithm is $O(n^3)$.

Computational cost of the TKHL statistic We compute the projection matrix $\mathbf{P}_{\mathbf{X}_I}^\perp$ for $O(p^3)$ operations. We compute and diagonalize the matrix \mathbf{K}_e^Z for $O(q^3)$ operations. Then we compute and diagonalize the matrix \mathbf{K}_e^a for $O(m^3 + m^2q + mq^2 + mqn + mn^2)$. Then we compute the matrix \mathbf{K}_T^a and the Nystrom TKHL statistic for $O(T^2m + Tm^2 + Tmq + Tq^2 + Tqn + Tn^2)$ operations. Typically, we have $n \gg q \geq m \geq T$ and p is small, then the global computational cost of the algorithm is $O(n^2(m + T) + q^3 + m^3)$, which is lower than $O(n^3)$.