

A Kernel Two-Sample Test for Functional Data

George Wynne

*Department of Mathematics
Imperial College London
London, SW7 2BU, UK*

G.WYNNE18@IC.AC.UK

Andrew B. Duncan

*Department of Mathematics
Imperial College London
London, SW7 2BU, UK*

A.DUNCAN@IC.AC.UK

Editor: Kenji Fukumizu

Abstract

We propose a nonparametric two-sample test procedure based on Maximum Mean Discrepancy (MMD) for testing the hypothesis that two samples of functions have the same underlying distribution, using kernels defined on function spaces. This construction is motivated by a scaling analysis of the efficiency of MMD-based tests for datasets of increasing dimension. Theoretical properties of kernels on function spaces and their associated MMD are established and employed to ascertain the efficacy of the newly proposed test, as well as to assess the effects of using functional reconstructions based on discretised function samples. The theoretical results are demonstrated over a range of synthetic and real world datasets.

Keywords: Kernel Methods, Functional Data Analysis, Testing

1. Introduction

Nonparametric two-sample tests for equality of distributions are widely studied in statistics, driven by applications in goodness-of-fit tests, anomaly and change-point detection and clustering. Classical examples of such tests include the Kolmogorov-Smirnov test (Kolmogorov-Smirnov et al., 1933; Smirnov, 1948; Schmid, 1958) and Wald-Wolfowitz runs test (Wald and Wolfowitz, 1940) with subsequent multivariate extensions (Friedman and Rafsky, 1979).

Due to advances in the ability to collect large amounts of real time or spatially distributed data there is a need to develop statistical methods appropriate for functional data, where each data sample is a discretised function. Such data has been studied for decades in the Functional Data Analysis (FDA) literature (Horváth and Kokoszka, 2012; Hsing and Eubank, 2015) particularly in the context of analysing populations of time series, or in statistical shape analysis (Mardia and Dryden, 1989). More recently, due to this modern abundance of functional data, increased study has been made in the machine learning literature for algorithms suited to such data (Berrendero et al., 2020; Chevyrev and Oberhauser, 2018; Kadri et al., 2016; Carmeli et al., 2010; Zhang et al., 2012).

In this paper we consider the case where the two probability distributions being compared are supported over a real, separable Hilbert space, for example $L^2(\mathcal{D})$ with $\mathcal{D} \subset \mathbb{R}^d$

and we have discretised observations of the function samples. If the samples consist of evaluations of the functions over a common mesh of points, then well-known methods for nonparametric two-sample testing for vector data can be used directly. Aside from the practical issue that observations are often made on irregular meshes for each different sample there is also the issue of degrading performance of classical tests as mesh size increases, meaning the observed vectors are high dimensional. As is typical with nonparametric two-sample tests, the testing power will degenerate rapidly with increasing data dimension. We therefore seek to better understand how to develop testing methods which are not strongly affected by the mesh resolution, exploiting intrinsic statistical properties of the underlying functional probability distributions.

In the past two decades kernels have seen a surge of use in statistical applications (Muandet et al., 2017; Gretton et al., 2012a; Sutherland et al., 2016; Borgwardt et al., 2006). In particular, kernel based two-sample testing (Gretton et al., 2012a, 2007) has become increasingly popular. These approaches are based on a distance on the space of probability measures known as *Maximum Mean Discrepancy*. Given two probability distributions P and Q , a kernel k is employed to construct a mapping known as the *mean embedding*, of the two distributions into an infinite dimensional Reproducing Kernel Hilbert Space (RKHS). The MMD between P and Q , denoted $\text{MMD}_k(P, Q)$ is given by the RKHS norm of the difference between the two embeddings, and defines a pseudo-metric on the space of probability measures. This becomes a metric if k is *characteristic*, see Section 3. By the kernel trick, MMD simplifies to a closed form, up to expectations, with respect to P and Q , which can be estimated unbiasedly using Monte Carlo simulations.

A major advantage of kernel two-sample tests is that they can be constructed on any input space which admits a well-defined kernel, including Riemannian manifolds (Pelletier, 2005), as well as discrete structures such as graphs (Shawe-Taylor and Cristianini, 2004) and strings (Gärtner, 2003). The flexibility in the choice of kernel is one of the strengths of MMD-based testing, where *a priori* knowledge of the structure of the underlying distributions can be encoded within the kernel to improve the sensitivity or specificity of the corresponding test. The particular choice of kernel strongly influences the efficiency of the test, however a general recipe for constructing a good kernel is still an open problem. On Euclidean spaces, radial basis function kernels are often used, meaning kernels of the form $k(x, y) = \phi(\gamma^{-1}\|x - y\|_2)$, where $\|\cdot\|_2$ is the Euclidean norm, $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a function and $\gamma > 0$ is the *bandwidth*. Numerous kernels used in practice belong to this class of kernels, including the Gaussian kernel $\phi(r) = e^{-r^2/2}$, the Laplace kernel $\phi(r) = e^{-r}$ and others including the rational quadratic kernel, the Matern kernel and the multiquadric kernel. The problem of selecting the bandwidth parameter to maximise test efficiency over a particular input space has been widely studied. One commonly used strategy is the *median heuristic* where the bandwidth is chosen to be the median of the inter-sample distance. Despite its popularity, there is only limited understanding of the median heuristic, with some notable exceptions. In Ramdas et al. (2015b); Reddi et al. (2015) the authors investigate the diminishing power of the kernel two-sample test using a Gauss kernel for distributions of white Gaussian random vectors with increasing dimension, demonstrating that under appropriate alternatives, the power of the test will decay with a rate dependent on the relative scaling of γ with respect to dimension. Related to kernel based tests are energy distance tests

(Székely, 2003; Székely and Rizzo, 2004), the relationship was made clear in Sejdinovic et al. (2013).

There has been relatively little work on understanding the theoretical properties of kernels on function spaces. A Gauss type kernel on $L^2([0, 1])$ was briefly considered in Christmann and Steinwart (2010, Example 3). Recently, in Chevyrev and Oberhauser (2018) a kernel was defined on the Banach space of paths on $[0, 1]$ of unbounded 1-variation, using a novel approach based on path signatures, demonstrating that this is a characteristic kernel over the space of such paths. The associated MMD has been employed as a loss function to train models generating stochastic processes (Kidger et al., 2019). Furthermore, in Nelsen and Stuart (2020) the authors propose an extension of the random Fourier feature kernel of Rahimi and Recht (2008) to the setting of an infinite dimensional Banach space, with the objective of regression between Banach spaces. This paper will build on aspects of these works, but with a specific emphasis on two-sample testing for functional data.

Two-sample testing in function spaces has received much attention in FDA and is studied in a variety of contexts. Broadly speaking there are two classes of methods. The first approach seeks to initially reduce the problem to a finite dimensional problem through a projection onto a finite orthonormal basis within the function space, typically using principal components, and then makes use of standard multivariate two-sample tests (Benko et al., 2009; Lopes et al., 2011). The second approach poses a two-sample test directly on function space (Aue et al., 2018; Bucchia and Wendler, 2017; Horváth et al., 2014; Pomann et al., 2016; Cabana et al., 2017). Many of these works construct the test on the Hilbert space $L^2(\mathcal{D})$ using the $L^2(\mathcal{D})$ norm as the testing statistic. A priori, it is not obvious why this norm will be well suited to the testing problem, in general. Investigation into the impact of the choice of distance in distanced based tests for functional data has been studied in the literature (Chen et al., 2014; Chakraborty and Zhang, 2019; Zhu et al., 2019) and a distance other than $L^2(\mathcal{D})$ for the functional data was advocated. This motivates the investigation into kernels which involve distances other than $L^2(\mathcal{D})$ in their formulation. In many works, the two-sample tests are designed to handle a specific class of discrepancy, such as a shift in mean (Horváth et al., 2012; Zhang et al., 2010), or a shift in covariance structure (Panaretos et al., 2010; Ferraty and Vieu, 2003; Fremdt et al., 2012).

This paper has two main aims. First, to naturally generalise the finite dimensional theory of kernels to real, separable Hilbert spaces to establish kernels that are characteristic, identify their RKHS and establish topological properties of the associated MMD. In particular the proof of characteristicity builds upon the spectral methods introduced in Sriperumbudur et al. (2010) and the weak convergence results build upon Simon-Gabriel and Schölkopf (2018). Second, we apply such kernels to the two-sample testing problem and analyse the power of the tests as well as the statistical impact of performing the tests using data reconstructed from discrete functional observations.

The specific contributions are as follows.

1. For Gaussian processes, we identify a scaling of the Gauss kernel bandwidth with mesh-size which results in testing power which is asymptotically independent of mesh-size, under mean-shift alternatives. In the scaling of vanishing mesh-size we demonstrate that the associated kernel converges to a kernel over functions.

2. Motivated by this, we construct a family of kernels defined on real, separable Hilbert spaces and identify sufficient conditions for the kernels to be characteristic, when MMD metrises the weak topology and provide an explicit construction of the reproducing kernel Hilbert space for a Gauss type kernel.
3. Using these kernels we investigate the statistical effect of using reconstructed functional data in the two-sample test.
4. We numerically validate our theory and compare the kernel based test with established two-sample tests from the functional data analysis literature.

The remainder of the paper is as follows. Section 2 covers preliminaries of modelling random functional data such as the Karhunen-Loève expansion and Gaussian measures. Section 3 recalls some important properties of kernels and their associated reproducing kernel Hilbert spaces, defines maximum mean discrepancy and the kernel two-sample test. Section 4 outlines the scaling of test power that occurs when an increasingly finer observation mesh is used for functional data. Section 5 defines a broad class of kernels and offers an integral feature map interpretation as well as outlining when the kernels are characteristic, meaning the two-sample test is valid. Section 6 highlights the statistical impact of fitting curves to discretised functions before performing the test. A relationship between MMD and weak convergence is highlighted and closed form expressions for the MMD and mean-embeddings when the distributions are Gaussian processes are given. Section 7 provides multiple examples of choices for the kernel hyper parameters and principled methods of constructing them. Section 8 contains multiple numerical experiments validating the theory in the paper, a simulation is performed to validate the scaling arguments of Section 4 and synthetic and real data sets are used to compare the performance of the kernel based test against existing functional two-sample tests. Concluding remarks and thoughts about future work are provided in Section 9.

2. Hilbert Space Modelling of Functional Data

In this paper we shall follow the Hilbert space approach to functional data analysis and use this section to outline the required preliminaries (Cuevas, 2014; Hsing and Eubank, 2015). Before discussing random functions we establish notation for families of operators that will be used extensively. Let \mathcal{X} be a real, separable Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{X}}$ then $L(\mathcal{X})$ denotes the set of bounded linear maps from \mathcal{X} to itself, $L^+(\mathcal{X})$ denotes the subset of $L(\mathcal{X})$ of operators that are self-adjoint (also known as symmetric) and non-negative, meaning $\langle Tx, y \rangle_{\mathcal{X}} \geq 0 \forall x, y \in \mathcal{X}$. The subset of $L^+(\mathcal{X})$ of trace class operators is denoted $L_1^+(\mathcal{X})$ and by the spectral theorem (Steinwart and Christmann, 2008, Theorem A.5.13) such operators can be diagonalised. This means for every $T \in L_1^+(\mathcal{X})$ there exists an orthonormal basis of eigenfunctions $\{e_n\}_{n=1}^{\infty}$ in \mathcal{X} such that $Tx = \sum_{n=1}^{\infty} \lambda_n \langle x, e_n \rangle_{\mathcal{X}} e_n$, where $\{\lambda_n\}_{n=1}^{\infty}$ are non-negative eigenvalues and the trace satisfies $\text{Tr}(T) = \sum_{n=1}^{\infty} \lambda_n < \infty$. When the eigenvalues are square summable the operator is called Hilbert-Schmidt and the Hilbert-Schmidt norm is $\|T\|_{HS}^2 = \sum_{n=1}^{\infty} \lambda_n^2$.

We now outline the Karhunen-Loève expansion of stochastic processes. Let $x(\cdot)$ be a mean zero stochastic process in $\mathcal{X} = L^2([0, 1])$, note the following will hold for a stochastic

process taking values in any real, separable Hilbert space but we focus on $L^2([0, 1])$ since it is the most common setting for functional data. Suppose that the pointwise covariance function $\mathbb{E}[x(s)x(t)] = k(s, t)$ is continuous. Define the covariance operator $C_k: \mathcal{X} \rightarrow \mathcal{X}$ associated with X by $C_k y(t) = \int_0^1 k(s, t)y(s)ds$. Then $C_k \in L_1^+(\mathcal{X})$ and denote the spectral decomposition $C_k y = \sum_{n=1}^{\infty} \lambda_n \langle y, e_n \rangle_{\mathcal{X}} e_n$. The Karhunen-Loève (KL) expansion (Sullivan, 2015, Theorem 11.4) provides a characterisation of the law of the process $x(\cdot)$ in terms of an infinite-series expansion. More specifically, we can write $x(\cdot) \sim \sum_{n=1}^{\infty} \lambda_n^{1/2} \eta_n e_n(\cdot)$, where $\{\eta_n\}_{n=1}^{\infty}$ are unit-variance uncorrelated random variables. Additionally, Mercer's theorem (Steinwart and Scovel, 2012) provides an expansion of the covariance as $k(s, t) = \sum_{n=1}^{\infty} \lambda_n e_n(s)e_n(t)$ where the convergence is uniform.

An important case of random functions are Gaussian processes (Rasmussen and Williams, 2006). Given a kernel k , see Section 3, and a function m we say x is a Gaussian process with mean function m and covariance function k if for every finite collection of points $\{s_n\}_{n=1}^N$ the random vector $(x(s_1), \dots, x(s_N))$ is a multivariate Gaussian random variable with mean vector $(m(s_1), \dots, m(s_N))$ and covariance matrix $k(s_n, s_m)_{n,m=1}^N$. The mean function and covariance function completely determines the Gaussian process. We write $x \sim \mathcal{GP}(m, k)$ to denote the Gaussian process with mean function m and covariance function k . If $x \sim \mathcal{GP}(0, k)$ then in the Karhunen-Loève representation $\eta_n \sim \mathcal{N}(0, 1)$ and the η_n are all independent.

Gaussian processes that take values in \mathcal{X} can be associated with Gaussian measures on \mathcal{X} . Gaussian measures are natural generalisations of Gaussian distributions on \mathbb{R}^d to infinite dimensional spaces, which are defined by a mean element and covariance operator rather than a mean vector and covariance matrix, for an introduction see Da Prato (2006, Chapter 1). Specifically $x \sim \mathcal{GP}(m, k)$ can be associated with the Gaussian measure N_{m, C_k} with mean m and covariance operator C_k , the covariance operator associated with k as outlined above. Similarly given any $m \in \mathcal{X}$ and $C \in L_1^+(\mathcal{X})$ then there exists a Gaussian measure $N_{m, C}$ with mean m and covariance operator C (Da Prato, 2006, Theorem 1.12). In fact, the Gaussian measure $N_{m, C}$ is characterised as the unique probability measure on \mathcal{X} with Fourier transform $\widehat{N}_{m, C}(y) = \exp(i\langle m, y \rangle_{\mathcal{X}} - \frac{1}{2}\langle Cy, y \rangle_{\mathcal{X}})$. Finally, if C is injective then a Gaussian measure with covariance operator C is called non-degenerate and has full support on \mathcal{X} (Da Prato, 2006, Proposition 1.25).

3. Reproducing Kernel Hilbert Spaces and Maximum Mean Discrepancy

This section will outline what a kernel and a reproducing kernel Hilbert space is with examples and associated references. Subsection 3.1 defines kernels and RKHS, Subsection 3.2 defines MMD and the corresponding estimators and Subsection 3.3 outlines the testing procedure.

3.1 Kernels and Reproducing Kernel Hilbert Spaces

Given a nonempty set \mathcal{X} a kernel is a function $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ which is symmetric, meaning $k(x, y) = k(y, x)$, for all $x, y \in \mathcal{X}$, and positive definite, that is, the matrix $\{k(x_n, x_m); n, m \in \{1, \dots, N\}\}$ is positive semi-definite, for all $\{x_n\}_{n=1}^N \subset \mathcal{X}$ and for $N \in \mathbb{N}$. For each kernel k there is an associated Hilbert space of functions over \mathcal{X}

known as the reproducing kernel Hilbert space (RKHS) denoted $\mathcal{H}_k(\mathcal{X})$ (Berlinet and Thomas-Agnan, 2004; Steinwart and Christmann, 2008; Fasshauer and McCourt, 2014). RKHSs have found numerous applications in function approximation and inference for decades, multiple detailed surveys exist in the literature (Saitoh and Sawano, 2016; Paulsen and Raghupathi, 2016). The RKHS associated with k satisfies the following two properties i). $k(\cdot, x) \in \mathcal{H}(\mathcal{X})$ for all $x \in \mathcal{X}$ ii). $\langle f, k(\cdot, x) \rangle_{\mathcal{H}(\mathcal{X})} = f(x)$ for all $x \in \mathcal{X}$ and $f \in \mathcal{H}(\mathcal{X})$. The latter is known as the reproducing property. The RKHS is constructed from the kernel in a natural way. The linear span of a kernel k with one input fixed $\mathcal{H}_0(\mathcal{X}) = \left\{ \sum_{n=1}^N a_n k(\cdot, x_n) : N \in \mathbb{N}, \{a_n\}_{n=1}^N \subset \mathbb{R}, \{x_n\}_{n=1}^N \subset \mathcal{X} \right\}$ is a pre-Hilbert space equipped with the following inner product $\langle f, g \rangle_{\mathcal{H}_0(\mathcal{X})} = \sum_{n=1}^N \sum_{m=1}^M a_n b_m k(x_n, y_m)$ where $f = \sum_{n=1}^N a_n k(\cdot, x_n)$ and $g = \sum_{m=1}^M b_m k(\cdot, y_m)$. The RKHS $\mathcal{H}_k(\mathcal{X})$ of k is then obtained from $\mathcal{H}_0(\mathcal{X})$ through completion. More specifically $\mathcal{H}_k(\mathcal{X})$ is the set of functions which are pointwise limits of Cauchy sequences in $\mathcal{H}_0(\mathcal{X})$ (Berlinet and Thomas-Agnan, 2004, Theorem 3). The relationship between kernels and RKHS is one-to-one, for every kernel the RKHS is unique and for every Hilbert space of functions such that there exists a function k satisfying the two properties above it may be concluded that the k is unique and a kernel. This result is known as the Aronszajn theorem (Berlinet and Thomas-Agnan, 2004, Theorem 3).

A kernel k on $\mathcal{X} \subseteq \mathbb{R}^d$ is said to be *translation invariant* if it can be written as $k(x, y) = \phi(x - y)$ for some ϕ . Bochner's theorem, Theorem 23 in the Appendix, tells us that if k is continuous and translation invariant then there exists a Borel measure on \mathcal{X} such that $\hat{\mu}_k(x - y) = k(x, y)$ where $\hat{\mu}$ is the characteristic function of μ defined as $\hat{\mu}(x) = \int_{\mathcal{X}} e^{i\langle x, y \rangle} d\mu(y)$. We call μ_k the spectral measure of k . The spectral measure is an important tool in the analysis of kernel methods and shall become important later when discussing the two-sample problem.

3.2 Maximum Mean Discrepancy

Given a kernel k and associated RKHS $\mathcal{H}_k(\mathcal{X})$ let \mathcal{P} be the set of Borel probability measures on \mathcal{X} and assuming k is measurable define $\mathcal{P}_k \subset \mathcal{P}$ as the set of all $P \in \mathcal{P}_k$ such that $\int k(x, x)^{\frac{1}{2}} dP(x) < \infty$. Note that $\mathcal{P}_k = \mathcal{P}$ if and only if k is bounded (Sriperumbudur et al., 2010, Proposition 2) which is very common in practice and shall be the case for all kernels considered in this paper. For $P, Q \in \mathcal{P}_k$ we define the *Maximum Mean Discrepancy* denoted $\text{MMD}_k(P, Q)$ as follows $\text{MMD}_k(P, Q) = \sup_{\|f\|_{\mathcal{H}_k(\mathcal{X})} \leq 1} \left| \int f dP - \int f dQ \right|$. This is an *integral probability metric* (Müller, 1997; Sriperumbudur et al., 2010) and without further assumptions defines a pseudo-metric on \mathcal{P}_k , which permits the possibility that $\text{MMD}_k(P, Q) = 0$ but $P \neq Q$.

We introduce the *mean embedding* $\Phi_k P$ of $P \in \mathcal{P}_k$ into $\mathcal{H}_k(\mathcal{X})$ defined by $\Phi_k P = \int k(\cdot, x) dP(x)$. This can be viewed as the mean in $\mathcal{H}_k(\mathcal{X})$ of the function $k(x, \cdot)$ with respect to P in the sense of a Bochner integral (Hsing and Eubank, 2015, Section 2.6).

Following Sriperumbudur et al. (2010, Section 2) this allows us to write

$$\begin{aligned} \text{MMD}_k(P, Q)^2 &= \left(\sup_{\|f\|_{\mathcal{H}_k(\mathcal{X})} \leq 1} \left| \int f dP - \int f dQ \right| \right)^2 = \left(\sup_{\|f\|_{\mathcal{H}_k(\mathcal{X})} \leq 1} |\langle \Phi_k P - \Phi_k Q, f \rangle| \right)^2 \\ &= \|\Phi_k P - \Phi_k Q\|_{\mathcal{H}_k(\mathcal{X})}^2. \end{aligned} \quad (1)$$

The crucial observation which motivates the use of MMD as an effective measure of discrepancy is that the supremum can be eliminated using the reproducing property of the inner product (Sriperumbudur et al., 2010, Section 2). This yields the following closed form representation

$$\begin{aligned} \text{MMD}_k(P, Q)^2 &= \int \int k(x, x') dP(x) dP(x') + \int \int k(y, y') dQ(y) dQ(y') \\ &\quad - 2 \int \int k(x, y) dP(x) dQ(y). \end{aligned} \quad (2)$$

It is clear that MMD_k is a metric over \mathcal{P}_k if and only if the map $\Phi_k: \mathcal{P}_k \rightarrow \mathcal{H}_k(\mathcal{X})$ is injective. A kernel is *characteristic* if the map Φ_k is injective over all of \mathcal{P} . Various works have provided sufficient conditions for a kernel over finite dimensional spaces to be characteristic (Sriperumbudur et al., 2010, 2011; Simon-Gabriel and Schölkopf, 2018).

Given independent samples $X_n = \{x_i\}_{i=1}^n$ from P and $Y_m = \{y_i\}_{i=1}^m$ from Q we wish to estimate $\text{MMD}_k(P, Q)^2$. A number of estimators have been proposed. For clarity of presentation we shall assume that $m = n$, but stress that all of the following can be generalised to situations where the two data-sets are unbalanced. Given samples X_n and Y_n , the following U-statistic is an unbiased estimator of $\text{MMD}_k^2(P, Q)^2$

$$\widehat{\text{MMD}}_k(X_n, Y_n)^2 := \frac{1}{n(n-1)} \sum_{i \neq j}^n h(z_i, z_j), \quad (3)$$

where $z_i = (x_i, y_i)$ and $h(z_i, z_j) = k(x_i, x_j) + k(y_i, y_j) - k(x_i, y_j) - k(x_j, y_i)$. This estimator can be evaluated in $O(n^2)$ time. An unbiased linear time estimator proposed in Gretton et al. (2012a) is given by

$$\widehat{\text{MMD}}_{k,\text{lin}}(X_n, Y_n)^2 := \frac{2}{n} \sum_{i=1}^{n/2} h(z_{2i-1}, z_{2i}), \quad (4)$$

where it is assumed that n is even. While the cost for computing $\widehat{\text{MMD}}_{k,\text{lin}}(X_n, Y_n)^2$ is only $O(n)$ this comes at the cost of reduced efficiency, i.e. $\text{Var}(\widehat{\text{MMD}}_k(X_n, Y_n)^2) < \text{Var}(\widehat{\text{MMD}}_{k,\text{lin}}(X_n, Y_n)^2)$, see for example Sutherland (2019). Various probabilistic bounds have been derived on the error between the estimator and $\text{MMD}_k(P, Q)^2$ (Gretton et al., 2012a, Theorem 10, Theorem 15).

3.3 The Kernel Two-Sample Test

Given independent samples $X_n = \{x_i\}_{i=1}^n$ from P and $Y_n = \{y_i\}_{i=1}^n$ from Q we seek to test the hypothesis $H_0: P = Q$ against the alternative hypothesis $H_1: P \neq Q$ without

making any distributional assumptions. The *kernel two-sample test* of Gretton et al. (2012a) employs an estimator of MMD as the test statistic. Indeed, fixing a characteristic kernel k , we reject H_0 if $\widehat{\text{MMD}}_k(X_n, Y_n)^2 > c_\alpha$, where c_α is a threshold selected to ensure a false-positive rate of α . While we do not have a closed-form expression for c_α , it can be estimated using a permutation bootstrap. The theoretical justification of this is provided by Arcones and Giné (1992). More specifically, we randomly shuffle $X_n \cup Y_n$, split it into two data sets X'_n and Y'_n , from which $\widehat{\text{MMD}}_k(X'_n, Y'_n)^2$ is calculated. This is repeated numerous times so that an estimator of the threshold \hat{c}_α is then obtained as the $(1 - \alpha)$ -th quantile of the resulting empirical distribution. The same test procedure may be performed using the linear time MMD estimator as the test statistic.

The power of a test is a measure of its ability to correctly reject the null hypothesis. More specifically, fixing α the false-positive rate, and obtaining an estimator \hat{c}_α of the threshold, we define the power of the test at α to be $\mathbb{P}(n\widehat{\text{MMD}}_k(X_n, Y_n)^2 \geq \hat{c}_\alpha)$. Invoking the central limit theorem for U-statistics (Serfling, 1980) we can quantify the decrease in variance of the unbiased MMD estimators, asymptotically as $n \rightarrow \infty$.

Theorem 1 (Gretton et al., 2012a, Corollary 16) *Suppose that $\mathbb{E}_{x \sim P, y \sim Q}[h^2(x, y)] < \infty$. Then under the alternative hypothesis $P \neq Q$, the estimator $\widehat{\text{MMD}}_k(X_n, Y_n)^2$ converges in distribution to a Gaussian as follows*

$$\sqrt{n} \left(\widehat{\text{MMD}}_k(X_n, Y_n)^2 - \text{MMD}_k^2(P, Q) \right) \xrightarrow{D} \mathcal{N}(0, 4\xi_1), \quad n \rightarrow \infty,$$

where $\xi_1 = \text{Var}_z [\mathbb{E}_{z'}[h(z, z')]]$. An analogous result holds for the linear-time estimator, with $\xi_2 = \text{Var}_{z, z'} [h(z, z')]$ instead of ξ_1 .

In particular, under the conditions of Theorem 1, for large n , the power of the test will satisfy the following asymptotic result

$$\mathbb{P} \left(n\widehat{\text{MMD}}_k(X_n, Y_n)^2 > \hat{c}_\alpha \right) \approx \Phi \left(\sqrt{n} \frac{\text{MMD}_k(P, Q)^2}{2\sqrt{\xi_1}} - \frac{c_\alpha}{2\sqrt{n\xi_1}} \right), \quad (5)$$

where Φ is the CDF for a standard Gaussian distribution and $\xi_1 = \text{Var}_z [\mathbb{E}_{z'}[h(z, z')]]$. The analogous result for the linear-time estimator holds with $\xi_2 = \text{Var}_{z, z'} [h(z, z')]$ instead of ξ_1 (Reddi et al., 2015; Liu et al., 2020). This suggests that the test power can be maximised by maximising $\text{MMD}_k(P, Q)^2/\sqrt{\xi_1}$ which can be seen as a signal-to-noise-ratio (Liu et al., 2020). It is evident from previous works that the properties of the kernel will have a very significant impact on the power of the test, and methods have been proposed for increasing test power by optimising the kernel parameters using the signal-to-noise-ratio as an objective (Sutherland et al., 2016; Reddi et al., 2015; Liu et al., 2020).

4. Resolution Independent Tests for Gaussian Processes

To motivate the construction of kernel two-sample tests for random functions, in this section we will consider the case where the samples X_n and Y_n are independent realisations of two Gaussian processes, observed along a regular mesh $\Xi_N = \{t_1, \dots, t_N\}$ of N points in \mathcal{D} where $\mathcal{D} \subset \mathbb{R}^d$ is some compact set. Therefore N will be the dimension of the

observed vectors. To develop ideas, we shall focus on a mean-shift alternative, where the underlying Gaussian processes are given by $\mathcal{GP}(0, k_0)$ and $\mathcal{GP}(m, k_0)$ respectively, where k_0 is a covariance function, and $m \in L^2(\mathcal{D})$ is the mean function. We use the subscript on k_0 to distinguish it from the kernel k we use to perform the test. We will use the linear time test due to easier calculations. This reduces to a multivariate two-sample hypothesis test problem on \mathbb{R}^N , with samples $X_n = \{x_i\}_{i=1}^n$ from $P = \mathcal{N}(0, \Sigma)$ and $Y_n = \{y_i\}_{i=1}^n$ from $Q = \mathcal{N}(m_N, \Sigma)$, where $\Sigma_{i,j} = k_0(t_i, t_j)$ for $i, j = 1, \dots, N$ and $m_N = (m(t_1), \dots, m(t_N))^T$.

We consider applying a two-sample kernel test as detailed in Section 3, with a Gaussian kernel $k(x, y) = \exp(-\frac{1}{2}\gamma_N^{-2}\|x - y\|_2^2)$ on \mathbb{R}^N where γ_N may depend on N . The large N limit was studied in Ramdas et al. (2015b) but not in the context of functional data. This motivates the question whether there is a scaling of γ_N with respect to N which, employing the structure of the underlying random functions, guarantees that the statistical power remains independent of the mesh size N . To better understand the influence of bandwidth on power, we use the signal-to-noise ratio as a convenient proxy, and study its behaviour in the large N limit. We say the mesh Ξ_N satisfies the Riemann scaling property if $\frac{1}{N}\|m_N\|_2^2 = \frac{1}{N}\sum_{i=1}^N m(t_i)^2 \rightarrow \int_{\mathcal{D}} m(t)^2 dt = \|m\|_{L^2(\mathcal{D})}^2$ as $N \rightarrow \infty$ for all $m \in L^2(\mathcal{D})$, this will be used in the next result to characterise the signal-to-noise ratio from the previous subsection.

Proposition 2 *Let P, Q be as above with Ξ_N satisfying the Riemann scaling property and $\gamma_N = \Omega(N^\alpha)$ with $\alpha > 1/2$ then if $k_0(s, t) = \delta_{st}$*

$$\frac{\text{MMD}_k(P, Q)^2}{\sqrt{\xi_2}} \sim \frac{\sqrt{N}\|m\|_{L^2(\mathcal{D})}^2}{2\sqrt{1 + \|m\|_{L^2(\mathcal{D})}^2}}, \quad (6)$$

and if k_0 is continuous and bounded then

$$\frac{\text{MMD}_k(P, Q)^2}{\sqrt{\xi_2}} \sim \frac{\|m\|_{L^2(\mathcal{D})}^2}{2\sqrt{\|C_{k_0}\|_{HS}^2 + \|C_{k_0}^{1/2}m\|_{L^2(\mathcal{D})}^2}}, \quad (7)$$

where \sim means asymptotically equal in the sense that the ratio of the left and right hand side converges to one as $N \rightarrow \infty$.

The proof of this result is in the Appendix and generalises Ramdas et al. (2015a) by considering non-identity Σ . Included in the proof are expressions for rates of $\text{MMD}_k(P, Q)^2$ and ξ_2 . The way this ratio increases with N , the number of observation points, in the white noise case makes sense since each observation is revealing new information about the signal as the randomness is independent at each observation. On the other hand the non-identity covariance matrix means the randomness is not independent at each observation and thus new information is not obtained at each observation point. Indeed the stronger the correlations, meaning the slower the decay of the eigenvalues of the covariance operator C_{k_0} , the smaller this ratio shall be since the Hilbert-Schmidt norm in the denominator will be larger.

It is important to note that the ratio in the right hand sides of (7) and (6) are independent of the choice of α once $\alpha > 1/2$ meaning that once greater than 1/2 this parameter

will be ineffective for obtaining greater testing power. The next subsection discusses how $\alpha = 1/2$ provides a scaling that results in kernels defined directly over function spaces, facilitating other methods to gain better test power.

4.1 Kernel Scaling

Proposition 2 does not include the case $\gamma_N = \Theta(N^{1/2})$ however it can be shown that the ratio does not degenerate in this case, see Theorem 18 and Theorem 20. In fact, the two different scales of the ratio, when Σ is the identity matrix or a kernel matrix, still occur. This is numerically verified in Section 8.

Suppose $\gamma_N = \gamma_0 N^{1/2}$ for some $\gamma_0 \in \mathbb{R}$ and one uses a kernel of the form $k(x, y) = f(\gamma_N^{-2} \|x - y\|_2^2)$ over \mathbb{R}^N for some continuous f . Suppose now though that our inputs shall be x_N, y_N , discretisations of functions $x, y \in L^2(\mathcal{D})$ observed on a mesh Ξ_N that satisfies the Riemann scaling property. Then as the mesh gets finer we observe the following scaling

$$k(x_N, y_N) = f(\gamma_N^{-2} \|x_N - y_N\|_2^2) \xrightarrow{N \rightarrow \infty} f(\gamma_0^{-2} \|x - y\|_{L^2(\mathcal{D})}^2).$$

Therefore the kernel, as the discretisation resolution increases, will converge to a kernel over $L^2(\mathcal{D})$ where the Euclidean norm is replaced with the $L^2(\mathcal{D})$ norm. For example the Gauss kernel would become $\exp(-\gamma_0^{-2} \|x - y\|_{L^2(\mathcal{D})}^2)$.

This scaling coincidentally is similar to the scaling of the widely used *median heuristic* defined as

$$\gamma^2 = \text{Median}\{\|a - b\|_2^2 : a, b \in \{x_i\}_{i=1}^n \cup \{y_i\}_{i=1}^m, a \neq b\}, \quad (8)$$

where $\{x_i\}_{i=1}^n$ are the samples from P , $\{y_i\}_{i=1}^m$ samples from Q . It was not designed with scaling in mind however in Reddi et al. (2015) it was noted that it results in a $\gamma^2 = \Theta(N)$ scaling for the mean shift, identity matrix case. The next lemma makes this more precise by relating the median of the squared distance to its expectation.

Lemma 3 *Let $P = \mathcal{N}(\mu_1, \Sigma_1)$ and $Q = \mathcal{N}(\mu_2, \Sigma_2)$ be independent Gaussian distributions on \mathbb{R}^N then $\mathbb{E}_{x \sim P, y \sim Q}[\|x - y\|_2^2] = \text{Tr}(\Sigma_1 + \Sigma_2) + \|\mu_1 - \mu_2\|_2^2$ and*

$$\left| \frac{\text{Median}_{x \sim P, y \sim Q}[\|x - y\|_2^2]}{\mathbb{E}_{x \sim P, y \sim Q}[\|x - y\|_2^2]} - 1 \right| \leq \sqrt{2} \left(1 - \frac{\|\mu_1 - \mu_2\|_2^4}{(\text{Tr}(\Sigma_1 + \Sigma_2) + \|\mu_1 - \mu_2\|_2^2)^2} \right)^{\frac{1}{2}},$$

in particular if P, Q are discretisations of Gaussian processes $\mathcal{GP}(m_1, k_1), \mathcal{GP}(m_2, k_2)$ on a mesh Ξ_N of N points satisfying the Riemann scaling property over some compact $\mathcal{D} \subset \mathbb{R}^d$ with $m_1, m_2 \in L^2(\mathcal{D})$ and k_1, k_2 continuous then $\mathbb{E}_{x \sim P, y \sim Q}[\|x - y\|_2^2] \sim N(\text{Tr}(C_{k_1} + C_{k_2}) + \|m_1 - m_2\|_{L^2(\mathcal{D})}^2)$ and as $N \rightarrow \infty$ the right hand side of the above inequality converges to

$$\sqrt{2} \left(1 - \frac{\|m_1 - m_2\|_{L^2(\mathcal{D})}^4}{(\text{Tr}(C_{k_1} + C_{k_2}) + \|m_1 - m_2\|_{L^2(\mathcal{D})}^2)^2} \right)^{\frac{1}{2}}.$$

The above lemma does not show that the median heuristic results in $\gamma_N = \gamma_0 N^{1/2}$ but relates it to the expected squared distance which does scale directly as $\gamma_0 N^{1/2}$. Therefore investigating the properties of such a scaling is natural.

Since $L^2(\mathcal{D})$ is a real, separable Hilbert space when using kernels defined directly over $L^2(\mathcal{D})$ in later sections we can leverage the theory of probability measures on such Hilbert spaces to deduce results about the testing performance of such kernels. In fact, we shall move past $L^2(\mathcal{D})$ and obtain results for kernels over arbitrary real, separable Hilbert spaces. Note that a different scaling of γ_N would not result in such a scaling of the norm to $L^2(\mathcal{D})$ so such theory cannot be applied.

5. Kernels and RKHS on Function Spaces

For the rest of the paper, unless specified otherwise, for example in Theorem 9, the spaces \mathcal{X}, \mathcal{Y} will be real, separable Hilbert spaces with inner products and norms $\langle \cdot, \cdot \rangle_{\mathcal{X}}, \langle \cdot, \cdot \rangle_{\mathcal{Y}}, \|\cdot\|_{\mathcal{X}}, \|\cdot\|_{\mathcal{Y}}$. We adopt the notation in Section 2 for various families of operators.

5.1 The Squared-Exponential T kernel

Motivated by the scaling discussions in Section 4 we define a kernel that acts directly on a Hilbert space.

Definition 4 For $T: \mathcal{X} \rightarrow \mathcal{Y}$ the squared-exponential T kernel (SE- T) is defined as

$$k_T(x, y) = e^{-\frac{1}{2}\|T(x)-T(y)\|_{\mathcal{Y}}^2}.$$

We use the name squared-exponential instead of Gauss because the SE- T kernel is not always the Fourier transform of a Gaussian distribution whereas the Gauss kernel on \mathbb{R}^d is, which is a key distinction and is relevant for our proofs. Lemma 25 in the Appendix assures us this function is a kernel. This definition allows us to adapt results about the Gauss kernel on \mathbb{R}^d to the SE- T kernel since it is the natural infinite dimensional generalisation. For example the following theorem characterises the RKHS of the SE- T kernel for a certain choice of T , as was done in the finite dimensional case in Minh (2009). Before we state the result we introduce the infinite dimensional generalisation of a multi-index, define Γ to be the set of summable sequences indexed by \mathbb{N} taking values in $\mathbb{N} \cup \{0\}$ and for $\gamma \in \Gamma$ set $|\gamma| = \sum_{n=1}^{\infty} \gamma_n$, so $\gamma \in \Gamma$ if and only if $\gamma_n = 0$ for all but finitely many $n \in \mathbb{N}$ meaning Γ is a countable set. We set $\Gamma_n = \{\gamma \in \Gamma: |\gamma| = n\}$ and the notation $\sum_{|\gamma| \geq 0}$ shall mean $\sum_{n=0}^{\infty} \sum_{\gamma \in \Gamma_n}$ which is a countable sum.

Theorem 5 Let $T \in L^+(\mathcal{X})$ be of the form $Tx = \sum_{n=1}^{\infty} \lambda_n^{1/2} \langle x, e_n \rangle_{\mathcal{X}} e_n$ with convergence in \mathcal{X} for some orthonormal basis $\{e_n\}_{n=1}^{\infty}$ and bounded positive coefficients $\{\lambda_n\}_{n=1}^{\infty}$ then the RKHS of the SE- T kernel is

$$\mathcal{H}_{k_T}(\mathcal{X}) = \left\{ F(x) = e^{-\frac{1}{2}\|Tx\|_{\mathcal{X}}^2} \sum_{|\gamma| \geq 0} w_{\gamma} x^{\gamma}: \sum_{|\gamma| \geq 0} \frac{\gamma!}{\lambda^{\gamma}} w_{\gamma}^2 < \infty \right\},$$

where $x^{\gamma} = \prod_{n=1}^{\infty} x_n^{\gamma_n}$, $x_n = \langle x, e_n \rangle_{\mathcal{X}}$, $\lambda^{\gamma} = \prod_{n=1}^{\infty} \lambda_n^{\gamma_n}$ and $\gamma! = \prod_{n=1}^{\infty} \gamma_n!$ and $\mathcal{H}_{k_T}(\mathcal{X})$ is equipped with the inner product $\langle F, G \rangle_{\mathcal{H}_{k_T}(\mathcal{X})} = \sum_{|\gamma| \geq 0} \frac{\gamma!}{\lambda^{\gamma}} w_{\gamma} v_{\gamma}$ where $F(x) = e^{-\frac{1}{2}\|Tx\|_{\mathcal{X}}^2} \sum_{|\gamma| \geq 0} w_{\gamma} x^{\gamma}$, $G(x) = e^{-\frac{1}{2}\|Tx\|_{\mathcal{X}}^2} \sum_{|\gamma| \geq 0} v_{\gamma} x^{\gamma}$.

Remark 6 *In the proof of Theorem 5 an orthonormal basis of $\mathcal{H}_{k_T}(\mathcal{X})$ is given which resembles the infinite dimensional Hermite polynomials which are used throughout infinite dimensional analysis and probability theory, for example see Da Prato and Zabczyk (2002, Chapter 10) and Nourdin and Peccati (2009, Chapter 2). In particular they are used to define Sobolev spaces for functions over a real, separable Hilbert space (Da Prato and Zabczyk, 2002, Theorem 9.2.12) which raises the interesting and, as far as we are aware, open question of how $\mathcal{H}_{k_T}(\mathcal{X})$ relates to such Sobolev spaces for different choices of T .*

For the two-sample test to be valid we need the kernel to be characteristic meaning the mean-embedding is injective over \mathcal{P} , so the test can tell the difference between any two probability measures. To understand the problem better we again leverage results regarding the Gauss kernel on \mathbb{R}^d , in particular the proof in Sriperumbudur et al. (2010, Theorem 9) that the Gauss kernel on \mathbb{R}^d is characteristic. This uses the fact that the Gauss kernel on \mathbb{R}^d is the Fourier transform of a Gaussian distribution on \mathbb{R}^d whose full support implies the kernel is characteristic. By choosing T such that the SE- T kernel is the Fourier transform of a Gaussian measure on \mathcal{X} that has full support we can use the same argument.

Theorem 7 *Let $T \in L_1^+(\mathcal{X})$ then the SE- T kernel is characteristic if and only if T is injective.*

This is dissatisfyingly limiting since $T \in L_1^+(\mathcal{X})$ is a restrictive assumption, for example it does not include $T = I$ the identity operator. We shall employ a limit argument to reduce the requirements on T . To this end we define admissible maps.

Definition 8 *A map $T: \mathcal{X} \rightarrow \mathcal{Y}$ is called admissible if it is Borel measurable, continuous and injective.*

The next result provides a broad family of kernels which are characteristic. It applies for \mathcal{X} being more general than a real, separable Hilbert space. A Polish space is a separable, completely metrizable topological space. Multiple examples of admissible T are given in Section 7 and are examined numerically in Section 8.

Theorem 9 *Let \mathcal{X} be a Polish space, \mathcal{Y} a real, separable Hilbert space and T an admissible map then the SE- T kernel is characteristic.*

Theorem 9 generalises Theorem 7. A critical result used in the proof is the Minlos-Sazanov theorem, detailed as Theorem 24 in the Appendix, which is an infinite dimensional version of Bochner's theorem. The result allows us to identify spectral properties of the SE- T kernel which are used to deduce characteristicity.

5.2 Integral Kernel Formulation

Let $k_0: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a kernel, $C \in L_1^+(\mathcal{X})$ and N_C the corresponding mean zero Gaussian measure on \mathcal{X} and define $k_{C,k_0}: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ as follows

$$k_{C,k_0}(x, y) := \int_{\mathcal{X}} k_0(\langle x, h \rangle_{\mathcal{X}}, \langle y, h \rangle_{\mathcal{X}}) dN_C(h).$$

Consider the particular case where $k_0(s, t) = \langle \Phi(s), \Phi(t) \rangle_{\mathcal{F}}$, where $\Phi : \mathbb{R} \rightarrow \mathcal{F}$ is a continuous feature map, mapping into a Hilbert space $(\mathcal{F}, \langle \cdot, \cdot \rangle_{\mathcal{F}})$, which will typically be \mathbb{R}^F for some $F \in \mathbb{N}$. In this case the functions $x \rightarrow \Phi(\langle x, h \rangle_{\mathcal{X}})$ can be viewed as \mathcal{F} -valued random features for each $h \in \mathcal{X}$ randomly sampled from N_C , and k_{C, k_0} is very similar to the random feature kernels considered in Nelsen and Stuart (2020) and Bach (2017). Following these previous works, we may completely characterise the RKHS of this kernel, the result involves $L_{N_C}^2(\mathcal{X}; \mathcal{F})$ which is the space of equivalence classes of functions from \mathcal{X} to \mathcal{F} that are square integrable in the \mathcal{F} norm with respect to N_C and $L^2(\mathcal{X}) := L^2(\mathcal{X}; \mathcal{F})$.

Proposition 10 *Suppose that $\psi(x, h) = \Phi(\langle x, h \rangle_{\mathcal{X}})$ satisfies $\psi \in L_{N_C \times N_C}^2(\mathcal{X} \times \mathcal{X}; \mathcal{F})$ then the RKHS defined by the kernel k_{C, k_0} is given by*

$$\mathcal{H}_{k_{C, k_0}}(\mathcal{X}) = \left\{ \int \langle v(h), \psi(\cdot, h) \rangle_{\mathcal{F}} dN_C(h) : v \in L_{N_C}^2(\mathcal{X}; \mathcal{F}) \right\} \subset L_{N_C}^2(\mathcal{X}).$$

The proof of this result is an immediate generalization of the real-valued case given in Nelsen and Stuart (2020). Using the spectral representation of translation invariant kernels we can provide conditions for k_{C, k_0} to be a characteristic kernel.

Proposition 11 *If k_0 is a kernel over $\mathbb{R} \times \mathbb{R}$ then k_{C, k_0} is a kernel over $\mathcal{X} \times \mathcal{X}$. If C is injective and k_0 is also continuous and translation invariant with spectral measure μ such that there exists an interval $(a, b) \subset \mathbb{R}$ with $\mu(U) > 0$ for every open subset $U \subset (a, b)$ then k_{C, k_0} is characteristic.*

For certain choices of T the SE- T kernel falls into a family of integral kernels. Indeed, if $k_0(x, y) = \cos(x - y)$ then k_{C, k_0} is the SE- $C^{\frac{1}{2}}$ kernel

$$k_{C, k_0}(x, y) = \widehat{N}_C(x - y) = e^{-\frac{1}{2}\|x-y\|_C^2} = e^{-\frac{1}{2}\sum_{n=1}^{\infty} \lambda_n(x_n - y_n)^2},$$

where $\|x - y\|_C^2 = \langle C(x - y), x - y \rangle_{\mathcal{X}}$, $\{\lambda_n\}_{n=1}^{\infty}$ are the eigenvalues of C and $x_n = \langle x, e_n \rangle$ are the coefficients with respect to the eigenfunction basis $\{e_n\}_{n=1}^{\infty}$ of C .

Secondly, let $\gamma > 0$ and assume C is non-degenerate and set k_0 to be the complex exponential of γ multiplied the by white noise mapping associated with C , see Da Prato and Zabczyk (2002, Section 1.2.4), then k_{C, k_0} is the SE- γI kernel

$$k_{C, k_0}(x, y) = k_{\gamma I}(x, y) = e^{-\frac{\gamma}{2}\|x-y\|_{\mathcal{X}}^2}, \quad (9)$$

Note that $k_{\gamma I}$ is not the Fourier transform of any Gaussian measure on \mathcal{X} (Maniglia and Rhandi, 2004, Proposition 1.2.11) which shows how the integral kernel framework is more general than only using the Fourier transform of Gaussian measures to obtain kernels, as was done in Theorem 7.

The integral framework can yield non-SE type kernels. Let N_1 be the measure associated with the Gaussian distribution $\mathcal{N}(0, 1)$ on \mathbb{R} , C be non-degenerate and $k_0(x, y) = \widehat{N}_1(x - y)$ then we have

$$k_{C, k_0}(x, y) = \int_{\mathcal{X}} \int_{\mathbb{R}} e^{iz\langle h, x-y \rangle} dN_1(z) dN_C(h) = (\|x - y\|_C^2 + 1)^{-\frac{1}{2}}. \quad (10)$$

Definition 12 For $T: \mathcal{X} \rightarrow \mathcal{Y}$ the inverse multi-quadric T kernel (IMQ- T) is defined as $k_T(x, y) = (\|T(x) - T(y)\|_{\mathcal{Y}}^2 + 1)^{-1/2}$.

By using Proposition 11 we immediately obtain that if $T \in L_1^+(\mathcal{X})$ and T is non-degenerate then the IMQ- T kernel is characteristic. But by the same limiting argument as Theorem 9 and the integral kernel formulation of IMQ- T we obtain a more general result.

Corollary 13 Under the same conditions as Theorem 9 the IMQ- T kernel is characteristic.

6. MMD on Function Spaces

In Section 5 we derived kernels directly over function spaces that were characteristic, meaning that the MMD induced by them is a metric on $\mathcal{P}(\mathcal{X})$. Therefore a two-sample test based on such kernels may be constructed, as detailed in Section 3, using the same form of U-statistic estimators and bootstrap technique as the finite dimensional scenario. This section will explore properties of the test. Subsection 6.1 will investigate the effect of performing the test on reconstructions of the random function based on observed data. Subsection 6.2 will provide explicit calculations for MMD when P, Q are Gaussian processes. Subsection 6.3 discusses the topology on $\mathcal{P}(\mathcal{X})$ induced by MMD and how it relates to the weak topology.

6.1 Influence of Function Reconstruction on MMD Estimator

In practice, rather than having access to the full realisation of random functions the data available will be some finite-dimensional representation of the functions, for example through discretisation over a mesh, or as a projection onto a finite dimensional basis of \mathcal{X} . Therefore to compute the kernel a user may need to approximate the true underlying functions from this finite dimensional representation. We wish to ensure the effectiveness of the tests using reconstructed data.

We formalise the notion of discretisation and reconstruction as follows. Assume that we observe $\{\mathcal{I}x_i\}_{i=1}^n$ where $\{x_i\}_{i=1}^n$ are the random samples from P and $\mathcal{I}: \mathcal{X} \rightarrow \mathbb{R}^N$ is a discretisation map. For example, \mathcal{I} could be point evaluation at some some fixed $t_1, t_2, \dots, t_N \in \mathcal{D}$ i.e. $\mathcal{I}x_i = (x_i(t_1), \dots, x_i(t_N))^{\top}$. Noisy point evaluation can also be considered in this framework. Then a reconstruction map $\mathcal{R}: \mathbb{R}^N \rightarrow \mathcal{X}$ is employed so that $\mathcal{R}\mathcal{I}X_n = \{\mathcal{R}\mathcal{I}x_i\}_{i=1}^n$ is used to perform the test, analogously for $\mathcal{R}\mathcal{I}Y_n$. For example, \mathcal{R} could be a kernel smoother or a spline interpolation operator. In practice one might have a different number of observations for each function, the following results can be adapted to this case straightforwardly.

Proposition 14 Assume k is a kernel on \mathcal{X} satisfying $|k(x, y) - k(u, v)| \leq L\|x - y - (u - v)\|_{\mathcal{X}}$ for all $u, v, x, y \in \mathcal{X}$ for some $L > 0$ and let $P, Q \in \mathcal{P}$ with $X_n = \{x_i\}_{i=1}^n, Y_n = \{y_i\}_{i=1}^n$ i.i.d. samples from P, Q respectively with reconstructed data $\mathcal{R}\mathcal{I}X_n = \{\mathcal{R}\mathcal{I}x_i\}_{i=1}^n, \mathcal{R}\mathcal{I}Y_n = \{\mathcal{R}\mathcal{I}y_i\}_{i=1}^n$ then

$$\left| \widehat{\text{MMD}}_k(X_n, Y_n)^2 - \widehat{\text{MMD}}_k(\mathcal{R}\mathcal{I}X_n, \mathcal{R}\mathcal{I}Y_n)^2 \right| \leq \frac{4L}{n} \sum_{i=1}^n \|\mathcal{R}\mathcal{I}x_i - x_i\|_{\mathcal{X}} + \|\mathcal{R}\mathcal{I}y_i - y_i\|_{\mathcal{X}}.$$

Corollary 15 *If k_T is the SE-T or IMQ-T kernel then the above bound holds with $\|T(\mathcal{R}\mathcal{I}x_i) - T(x_i)\|_{\mathcal{Y}}$, $\|T(\mathcal{R}\mathcal{I}y_i) - T(y_i)\|_{\mathcal{Y}}$ instead of $\|\mathcal{R}\mathcal{I}x_i - x_i\|_{\mathcal{X}}$, $\|\mathcal{R}\mathcal{I}y_i - y_i\|_{\mathcal{X}}$ with $L = \frac{1}{\sqrt{e}}$ and $L = \frac{2}{3\sqrt{3}}$ respectively.*

An analogous result can be derived for the linear time estimator with the same proof technique. While Proposition 14 provides a statement on the approximation of $\widehat{\text{MMD}}_k(\mathcal{R}\mathcal{I}X_n, \mathcal{R}\mathcal{I}Y_n)^2$ we are primarily concerned with its statistical properties. Asymptotically, the test efficiency is characterised via the Gaussian approximation in Theorem 1, specifically through the asymptotic variance in (5). The following result provides conditions under which a similar central limit theorem holds for the estimator based on reconstructed data, with the same asymptotic variance. It imposes conditions on the number of discretisation points per function sample N , the error of the approximations and the number of function samples n .

Theorem 16 *Let k satisfy the condition in Proposition 14 and let $X_n = \{x_i\}_{i=1}^n$ and $Y_n = \{y_i\}_{i=1}^n$ be i.i.d. samples from P and Q respectively with $P \neq Q$, and associated reconstructions $\mathcal{R}\mathcal{I}X_n$ and $\mathcal{R}\mathcal{I}Y_n$ based on $N(n)$ dimensional discretisations $\mathcal{I}X_n$ and $\mathcal{I}Y_n$ where $N(n) \rightarrow \infty$ as $n \rightarrow \infty$. If $n^{\frac{1}{2}}\mathbb{E}_{x \sim P}[\|x - \mathcal{R}\mathcal{I}x\|_{\mathcal{X}}] \rightarrow 0$ and $n^{\frac{1}{2}}\mathbb{E}_{y \sim P}[\|y - \mathcal{R}\mathcal{I}y\|_{\mathcal{X}}] \rightarrow 0$ as $n \rightarrow \infty$, then for $\xi = 4\text{Var}_z[\mathbb{E}_{z'}[h(z, z')]]$*

$$n^{\frac{1}{2}}(\widehat{\text{MMD}}_{k_T}(\mathcal{R}\mathcal{I}X_n, \mathcal{R}\mathcal{I}Y_n)^2 - \text{MMD}_{k_T}(P, Q)^2) \xrightarrow{d} \mathcal{N}(0, \xi).$$

A similar result can be derived for the linear time estimator by using the linear time estimator version of Proposition 14. The discretisation map, number of discretisations per function sample and the reconstruction map need to combine to satisfy the convergence assumption. For example if a weaker reconstruction map is used then more observations per function sample will be needed to compensate for this. Additionally if the discretisation map offers less information about the underlying function, for example it provides observations that are noise corrupted, then more observations per function sample are needed.

We now discuss three settings in which these assumptions hold, relevant to different applications. We shall assume that k satisfies the conditions of Proposition 14 and that $T = I$.

6.1.1 LINEAR INTERPOLATION OF REGULARLY SAMPLED DATA

Let $\mathcal{X} = L^2([0, 1])$ and $\Xi_{N(n)} = \{t_i\}_{i=1}^{N(n)}$ be a mesh of evaluation points where $t_{i+1} - t_i = N(n)^{-1}$ for all i and define $\mathcal{I}x = (x(t_1), \dots, x(t_{N(n)}))^{\top} \in \mathbb{R}^{N(n)}$. Let \mathcal{R} be the piecewise linear interpolant defined as

$$(\mathcal{R}\mathcal{I}x)(t) = (x(t_{i+1}) - x(t_i)) \frac{t - t_i}{t_{i+1} - t_i} + x(t_i), \quad \text{for } t \in [t_i, t_{i+1}).$$

Suppose that realisations $x \sim P$ and $y \sim Q$ are almost surely in $C^2([0, 1])$ and in particular satisfy $\mathbb{E}_{x \sim P}[\|x''\|_{\mathcal{X}}^2] < \infty$ and $\mathbb{E}_{y \sim Q}[\|y''\|_{\mathcal{X}}^2] < \infty$. Then

$$\mathbb{E}_{x \sim P}[\|x - \mathcal{R}\mathcal{I}x\|_{\mathcal{X}}] \leq \frac{1}{N(n)^2} \mathbb{E}_{x \sim P}[\|x''\|_{\mathcal{X}}],$$

and analogously for $y \sim Q$. Therefore if $N(n) \sim n^\alpha$ with $\alpha > 1/4$ then the conditions of Theorem 16 are satisfied. Conditions for P, Q to have samples lie almost surely in $C^2([0, 1])$ in the case where P, Q are Gaussian measures has been studied widely (Lukić and Beder, 2001; Steinwart, 2018; Karvonen et al., 2020).

6.1.2 KERNEL INTERPOLANT OF QUASI-UNIFORMLY SAMPLED DATA

Let $\mathcal{X} = L^2(\mathcal{D})$ with $\mathcal{D} \subset \mathbb{R}^d$ compact. As in the previous example, \mathcal{I} will be the evaluation operator over a set of points $\Xi_{N(n)} = \{t_i\}_{i=1}^{N(n)}$ but now we assume the points are placed quasi-uniformly in the scattered data approximation sense, for example regularly placed grid points, see Wynne et al. (2021) and Wendland (2005, Chapter 4) for other method to obtain quasi-uniform points.

We set \mathcal{R} as the kernel interpolant using a kernel k_0 with RKHS norm equivalent to $W_2^\nu(\mathcal{D})$ with $\nu > d/2$, this is achieved by the common Matérn and Wendland kernels (Wendland, 2005; Kanagawa et al., 2018). For this choice of recovery operator, $(\mathcal{R}\mathcal{I}x)(t) = k_0(t, \Xi)K_0^{-1}\mathcal{I}(x)$ where $k_0(t, \Xi) = (k_0(t, t_1), \dots, k_0(t, t_{N(n)}))$ and K_0 is the kernel matrix of k_0 over $\Xi_{N(n)}$.

Suppose the realisations of P and Q lie almost surely in $W_2^\tau(\mathcal{D})$ for some $\tau > d/2$, this assumption is discussed when P, Q are Gaussian processes in Kanagawa et al. (2018, Section 4), then

$$\mathbb{E}_{x \sim P}[\|x - \mathcal{R}\mathcal{I}x\|_{\mathcal{X}}] \leq CN(n)^{-(\tau \wedge \nu)/d},$$

for some constant $C > 0$, with an identical result holding for realisations of Q (Wynne et al., 2021; Narcowich et al., 2006). It follows that choosing $N(n) \sim n$ guarantees that the conditions of Theorem 16 hold in this setting. Here we see that to maintain a scaling of $N(n)$ independent of dimension d we need the signal smoothness ν, τ to increase with d .

Note that the case where \mathcal{I} is pointwise evaluation corrupted by noise may be treated in a similar way by using results from Bayesian non-parametrics, for example van der Vaart and van Zanten (2011, Theorem 5). In this case \mathcal{R} would be the posterior mean of a Gaussian process that is conditioned on $\mathcal{I}(x)$.

6.1.3 PROJECTION ONTO AN ORTHONORMAL BASIS

Let \mathcal{X} be an arbitrary real, separable Hilbert space and $\{e_n\}_{n=1}^\infty$ be an orthonormal basis. Suppose that \mathcal{I} is a projection operator onto the first $N(n)$ elements of the basis $\mathcal{I}x = (\langle x, e_1 \rangle_{\mathcal{X}}, \dots, \langle x, e_{N(n)} \rangle_{\mathcal{X}})^\top$ and \mathcal{R} constructs a function from basis coefficients $\mathcal{R}(\beta_1, \dots, \beta_{N(n)}) = \sum_{i=1}^{N(n)} \beta_i e_i$ meaning $\mathcal{R}\mathcal{I}x = \sum_{i=1}^{N(n)} \langle x, e_i \rangle_{\mathcal{X}} e_i$. A typical example on $L^2([0, 1])$ would be a Fourier series representation of the samples $\{x_i\}_{i=1}^n$ and $\{y_i\}_{i=1}^n$ from which the functions can be recovered efficiently via an inverse Fast Fourier Transform. By Parseval's theorem $\mathbb{E}_{x \sim P}[\|x - \mathcal{R}\mathcal{I}x\|_{\mathcal{X}}^2] = \sum_{i=N(n)+1}^\infty |\langle x, e_i \rangle_{\mathcal{X}}|^2 \rightarrow 0$ as $N(n) \rightarrow \infty$. For the conditions of Theorem 16 to hold, we require that $n^{1/2} \mathbb{E}_{x \sim P}[\sum_{i=N(n)+1}^\infty |\langle x, e_i \rangle_{\mathcal{X}}|^2] \rightarrow 0$ as $n \rightarrow \infty$ which means $N(n)$ will need to grow in a way to compensate for the auto-correlation of realisations of P and Q .

In this setting the use of the integral kernels described in Section 5.2 are particularly convenient. Indeed, let $C = \sum_{i=1}^\infty \lambda_i e_i \otimes e_i \in L_+^1(\mathcal{X})$ and consider the integral kernel k_{C, k_0} .

An evaluation of the kernel $k(x, y)$ can then be approximated using a random Fourier feature approach (Rahimi and Recht, 2008) by

$$k_{C, k_0}(x, y) \approx \frac{1}{n_S} \sum_{l=1}^{n_S} k_0 \left(\sum_{i=1}^{N(n)} \lambda_i^{\frac{1}{2}} x_i \eta_i^l, \sum_{i=1}^{N(n)} \lambda_i^{\frac{1}{2}} y_i \eta_i^l \right),$$

where $x_i = \langle x, e_i \rangle_{\mathcal{X}}$, $y_i = \langle y, e_i \rangle_{\mathcal{X}}$ and $\eta_i^l \sim \mathcal{N}(0, 1)$ i.i.d. for $i = 1, \dots, N(n)$ and $l = 1, \dots, n_S$ for some $n_S \in \mathbb{N}$. This permits opportunities to reduce the computational cost of MMD tests as judicious choices of Φ will permit accurate approximations of $k(x, y)$ using n_S small. Similarly, the weights, λ_i can be chosen to reduce the dimensionality of the functions x_i and y_i .

6.2 Explicit Calculations for Gaussian Processes

A key property of the SE- T kernel is that the mean-embedding $\Phi_{k_T} P$ and $\text{MMD}_{k_T}(P, Q)^2$ have closed form solutions when P, Q are Gaussian measures. Using the natural correspondence between Gaussian measures and Gaussian processes from Section 2 we may get closed form expressions for Gaussian processes. This addresses the open question regarding the link between Bayesian non-parametrics methods and kernel mean-embeddings that was discussed in Muandet et al. (2017, Section 6.2).

Before stating the next result we need to introduce the concept of determinant for an operator, for $S \in L_1(\mathcal{X})$ define $\det(I+S) = \prod_{n=1}^{\infty} (1+\lambda_n)$ where $\{\lambda_n\}_{n=1}^{\infty}$ are the eigenvalues of S . The equality $\det((I+S)(I+R)) = \det(I+S)\det(I+R)$ holds and is frequently used.

Theorem 17 *Let k_T be the SE- T kernel for some $T \in L^+(\mathcal{X})$ and $P = N_{a,S}$ be a non-degenerate Gaussian measure on \mathcal{X} then*

$$\Phi_{k_T}(N_{a,S})(x) = \det(I+TST)^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+TST)^{-1} T(x-a), T(x-a) \rangle_{\mathcal{X}}}.$$

Theorem 18 *Let k_T be the SE- T kernel for some $T \in L^+(\mathcal{X})$ and $P = N_{a,S}, Q = N_{b,R}$ be non-degenerate Gaussian measures on \mathcal{X} then*

$$\begin{aligned} \text{MMD}_{k_T}(P, Q)^2 &= \det(I+2TST)^{-\frac{1}{2}} + \det(I+2TRT)^{-\frac{1}{2}} \\ &\quad - 2 \det \left((I+TST)(I+(TRT)^{\frac{1}{2}}(I+TST)^{-1}(TRT)^{\frac{1}{2}}) \right)^{-\frac{1}{2}} \\ &\quad \times e^{-\frac{1}{2} \langle (I+T(S+R)T)^{-1} T(a-b), T(a-b) \rangle_{\mathcal{X}}}. \end{aligned}$$

These results outline the geometry of Gaussian measures with respect to the distance induced by the SE- T kernel. We see that the means only occur in the formula through their difference and if both mean elements are zero then the distance is measured purely in terms of the spectrum of the covariance operators.

Corollary 19 *Under the Assumptions of Theorem 18 and that T, S, R commute then*

$$\begin{aligned} \text{MMD}_{k_T}(P, Q)^2 &= \det(I+2TST)^{-\frac{1}{2}} + \det(I+2TRT)^{-\frac{1}{2}} \\ &\quad - 2 \det(I+T(S+R)T)^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+T(S+R)T)^{-1} T(a-b), T(a-b) \rangle_{\mathcal{X}}}. \end{aligned}$$

Since the variance terms ξ_1, ξ_2 from Section 3 are simply multiple integrals of the SE-T kernel against Gaussian measures we may obtain closed forms for them too. Theorem 20 is a particular instance of the more general Theorem 27 in the Appendix.

Theorem 20 *Let $P = N_S, Q = N_{m,S}$ be non-degenerate Gaussian measures on \mathcal{X} , $T \in L_1^+(\mathcal{X})$ and assume T and S commute then when using the SE-T kernel*

$$\begin{aligned} \xi_1 &= 2 \det(\Sigma_S)^{-\frac{1}{2}} \left(1 + e^{-\langle (I+3TST)^{-1}Tm, Tm \rangle_{\mathcal{X}}} - 2e^{-\frac{1}{2}\langle (I+2TST)\Sigma_S^{-1}Tm, Tm \rangle_{\mathcal{X}}} \right. \\ &\quad \left. - 2 \det(I + 2TST)^{-1} \left(1 + e^{-\langle (I+2TST)^{-1}Tm, Tm \rangle_{\mathcal{X}}} - 2e^{-\frac{1}{2}\langle (I+2TST)^{-1}Tm, Tm \rangle_{\mathcal{X}}} \right), \right. \\ \xi_2 &= 2 \det(I + 4TST)^{-\frac{1}{2}} \left(1 + e^{-\langle (I+4TST)^{-1}Tm, Tm \rangle_{\mathcal{X}}} \right. \\ &\quad \left. - 2 \det(I + 2TST)^{-1} \left(1 + e^{-\langle (I+2TST)^{-1}Tm, Tm \rangle_{\mathcal{X}}} - 4e^{-\frac{1}{2}\langle (I+2TST)^{-1}Tm, Tm \rangle_{\mathcal{X}}} \right) \right. \\ &\quad \left. - 8 \det(\Sigma_S)^{-\frac{1}{2}} e^{-\frac{1}{2}\langle (I+2TST)\Sigma_S^{-1}Tm, Tm \rangle_{\mathcal{X}}} \right), \end{aligned}$$

where $\Sigma_S = (I + TST)(I + 3TST)$.

6.3 Weak Convergence and MMD

We know that MMD is a metric on \mathcal{P} when k is characteristic, so it is natural to identify the topology it generates and in particular how it relates to the standard topology for elements of \mathcal{P} , the weak topology.

Theorem 21 *Let \mathcal{X} be a Polish space, k a bounded, continuous, characteristic kernel on $\mathcal{X} \times \mathcal{X}$ and $P \in \mathcal{P}$ then $P_n \xrightarrow{w} P$ implies $\text{MMD}_k(P_n, P) \rightarrow 0$ and if $\{P_n\}_{n=1}^\infty \subset \mathcal{P}$ is tight then $\text{MMD}_k(P_n, P) \rightarrow 0$ implies $P_n \xrightarrow{w} P$ where \xrightarrow{w} denotes weak convergence.*

For a discussion on weak convergence and tightness see Billingsley (1971). The tightness is used to compensate for the lack of compactness of \mathcal{X} which is often required in analogous finite dimensional results. In particular, in Chevyrev and Oberhauser (2018) an example where $\text{MMD}_k(P_n, P) \rightarrow 0$ but P_n does not converge to P was given without the assumption of tightness. A precise characterisation of the relationship between MMD and weak convergence over a Polish space is an open problem.

7. Practical Considerations for Kernel Selection

We now present examples and techniques to choose kernels and construct maps T that are admissible. Two main categories will be discussed, integral operators induced by kernels and situation specific kernels.

For the first category assume $\mathcal{X} = \mathcal{Y} = L^2(\mathcal{D})$ for some compact $\mathcal{D} \subset \mathbb{R}^d$ and let k_0 be a measurable kernel over $\mathcal{D} \times \mathcal{D}$ and set $T = C_{k_0}$ where C_{k_0} is the covariance operator associated with k_0 , see Section 2. We call k_0 an *admissible kernel* if C_{k_0} is admissible. If k_0 is continuous then by Mercer's theorem $C_{k_0}x = \sum_{n=1}^\infty \lambda_n \langle x, e_n \rangle e_n$ for some positive sequence $\{\lambda_n\}_{n=1}^\infty$ and orthonormal set $\{e_n\}_{n=1}^\infty$ (Steinwart and Christmann, 2008, Chapter 4.5). To be admissible C_{k_0} needs to be injective which is equivalent to $\{e_n\}_{n=1}^\infty$ forming a basis (Steinwart and Scovel, 2012, Proof of Theorem 3.1). Call k_0 *integrally strictly positive definite* (ISPD) if $\int_{\mathcal{D}} \int_{\mathcal{D}} x(s)k_0(s, t)x(t)dsdt > 0$ for all non-zero $x \in \mathcal{X}$. Recall

that if k_0 is translation invariant then by Theorem 23 there exists a measure μ_{k_0} such that $\hat{\mu}_{k_0}(s - t) = k(s, t)$.

Proposition 22 *Let $\mathcal{D} \subset \mathbb{R}^d$ be compact and k_0 a continuous kernel on \mathcal{D} , if k_0 is ISPD then k_0 is admissible. In particular, if k_0 is continuous and translation invariant and μ_{k_0} has full support on \mathcal{D} then k_0 is admissible.*

For multiple examples of ISPD kernels see Sriperumbudur et al. (2011) and of μ_{k_0} see Sriperumbudur et al. (2010). Using the product to convolution property of the Fourier transform one can construct k_0 such that μ_{k_0} has full support relatively easily or modify standard integral operators which aren't admissible. For example, for some $F \in \mathbb{N}$ consider the kernel $k_{\cos(F)}(s, t) = \sum_{n=0}^{F-1} \cos(2\pi n(s - t))$ on $[0, 1]^2$ whose spectral measure is a sum of Dirac measures so does not have full support. If the Dirac measures are convolved with a Gaussian then they would be smoothed out and would result in a measure with full support. Since convolution in the frequency domain corresponds to a product in space domain the new kernel $k_{\text{c-exp}(F,l)}(s, t) = e^{-\frac{1}{2l^2}(s-t)^2} k_{\cos(F)}(s, t)$ satisfies the conditions of Proposition 22. This technique of frequency modification has found success in modelling audio signals (Wilkinson, 2019, Section 3.4). In general, any operator of the form $Tx = \sum_{n=1}^{\infty} \lambda_n \langle x, e_n \rangle_{\mathcal{X}} e_n$ for positive, bounded $\{\lambda_n\}_{n=1}^{\infty}$ and an orthonormal basis $\{e_n\}_{n=1}^{\infty}$ is admissible even if it is not induced by a kernel, for example the functional Mahalanobis distance (Berrendero et al., 2020).

The second category is scenario specific choices. By this we mean kernels whose structure is specified to the testing problem at hand. For example, while the kernel two-sample test may be applied for distributions with arbitrary difference one may tailor it for a specific testing problem, such a difference of covariance operator. If one does only wish to test for difference in covariance operator of the two probability measures then an appropriate kernel would be $k_{\text{cov}}(x, y) = \langle x, y \rangle_{\mathcal{X}}^2$ which is not characteristic but $\text{MMD}_{k_{\text{cov}}}(P, Q) = 0$ if and only if P, Q have the same covariance operator. However, a practitioner may want a kernel which emphasises difference in covariance operator, due to prior knowledge regarding the data, while still being able to detect arbitrary difference, in case the difference is more complicated than initially thought. We now present examples of T which do this. To emphasise higher order moments, let $\mathcal{X} \subset L^4(\mathcal{D})$ and \mathcal{Y} the direct sum of $L^2(\mathcal{D})$ with itself equipped with the norm $\|(x, x')\|_{\mathcal{Y}}^2 = \|x\|_{L^2(\mathcal{D})}^2 + \|x'\|_{L^2(\mathcal{D})}^2$ and $T(x) = (x, x^2)$. This map captures second order differences and first order differences individually, as opposed to the polynomial map which combines them. Alternatively, one might be in a scenario where the difference in the distributions is presumed to be in the lower frequencies. In this case a map of the form $T(x) = \sum_{n=1}^F \lambda_n \langle x, e_n \rangle_{\mathcal{X}} e_n$ could be used for decreasing, positive λ_n and some orthogonal e_n . This will not be characteristic, since T only acts on F frequencies, however if F is picked large enough then good performance could still be obtained in practice. For example, λ_n, e_n could be calculated empirically from the function samples using functional principal component analysis and F could be picked so that the components explain a certain percentage of total variance. See Horváth and Kokoszka (2012) for a deeper discussion on functional principal components and its central role in functional data analysis.

All of the choices outlined above have associated hyperparameters, for example if $T = C_{k_0}$ then hyperparameters of k_0 are hyperparameters of T such as the bandwidth. It is

outside the scope of this paper to investigate new methods to choose these parameters but we do believe it is important future work. Multiple methods for finite dimensional data have been proposed using the surrogates for test power outlined in Section 4 (Sutherland et al., 2016; Gretton et al., 2012b; Liu et al., 2020) which could have potential for use in the infinite dimensional scenario.

8. Numerical Simulations

In this section we perform numerical simulations on real and synthetic data to reinforce the theoretical results. Code is available at <https://github.com/georgewynne/Kernel-Functional-Data>.

8.1 Power Scaling of Functional Data

Verification of the power scaling when performing the mean shift two-sample test using functional data, discussed in Section 4, is performed. Specifically we perform the two-sample test using the SE- I kernel with $x \sim \mathcal{GP}(0, k_l)$ and $y \sim \mathcal{GP}(m, k_l)$ where $m(t) = 0.05$ for $t \in [0, 1]$ and $k_l(s, t) = e^{-\frac{1}{2l^2}(s-t)^2}$ with 50 samples from each distribution. This is repeated 500 times to calculate power with 1000 permutations used in the bootstrap to simulate the null. The observation points are a uniform grid on $[0, 1]$ with N points, meaning N will be the dimension of the observed discretised function vectors. The parameter l dictates the dependency of the fluctuations. Small l means less dependency between the random function values so the covariance matrix is closer to the identity. When the random functions are m with $\mathcal{N}(0, 1)$ i.i.d. corruption the corresponding value of l is zero which essentially means $k_l(x, y) = \delta_{xy}$. In this case the scaling of power is expected to follow (6) and grow asymptotically as \sqrt{N} . On the other hand if $l > 0$ the fluctuations within each random function are dependent and we expect scaling as (7) which does not grow asymptotically with N .

Figure 1 confirms this theory showing that power increases with a finer observation mesh only when there is no dependence in the random functions values. We see some increase of power as the mesh gets finer for the case of small dependency however the rate of increase is much smaller than the i.i.d. setting.




Figure 1: Test power as mesh size decreases given different point dependency strengths

Abbreviation	Description	Reference
ID	SE- T kernel, $T = I$	Section 5
FPCA	SE- T kernel, T based on functional principle components	Section 7
SQR	SE- T kernel, T squaring feature expansion	Section 7
CEXP	SE- T kernel, T based on the cosine-exponential kernel	Section 7
COV	Covariance kernel $k(x, y) = \langle x, y \rangle_{\mathcal{X}}^2$	Section 7
FAD	Functional Anderson-Darling	(Pomann et al., 2016)
CVM	Functional Cramer-von Mises	(Hall and Keilegom, 2007)
BOOT-HS	Bootstrap Hilbert-Schmidt	(Papadoditis and Sapatinas, 2016)
FPCA- χ	Functional Principal Component χ^2	(Fremdt et al., 2012)

Table 1: Summary of two-sample tests and kernels used in numerical experiments

8.2 Synthetic Data

The tests are all performed using the k_{cov} kernel from Section 7 and the SE- T kernel for four different choices of T and unless stated otherwise $\mathcal{Y} = L^2([0, 1])$ and we use the short hand L^2 for $L^2([0, 1])$ and n_x, n_y will denote the sample sizes of the two samples. To calculate power each test is repeated 500 times and 1000 permutations are used in the bootstrap to simulate the null distribution.

COV will denote the k_{cov} kernel, which can only detect difference of covariance operator. ID will denote $T = I$. CEXP will denote $T = C_{k_0}$ with $k_0 = k_{\text{c-exp}(20, \sqrt{10})}$ the cosine exponential kernel. SQR will denote $T(x) = (x, x^2)$ with \mathcal{Y} the direct sum of $L^2([0, 1])$ with itself as detailed in Section 7. FPCA will denote $Tx = \sum_{n=1}^F \lambda_n^{1/2} \langle x, e_n \rangle e_n$ where λ_n, e_n are empirical functional principal components and principal values computed from the union of the two collections of samples with F chosen such that 95% of variance is explained. The abbreviations are summarised in Table 1 along with references to the other tests being compared against.

For all four instances of the SE- T kernel $\exp(-\frac{1}{2\gamma^2} \|T(x) - T(y)\|_{\mathcal{Y}}^2)$ we use, for all but SQR scenario, the median heuristic $\gamma^2 = \text{Median}\{\|T(a) - T(b)\|_{\mathcal{Y}}^2 : a, b \in \{x_i\}_{i=1}^{n_x} \cup \{y_i\}_{i=1}^{n_y}, a \neq b\}$. As the SQR scenario involves two norms in the exponent two calculations of median heuristic are needed so that the kernel used is $\exp(-\frac{1}{2\gamma_1^2} \|x - y\|_{L^2}^2 - \frac{1}{2\gamma_2^2} \|x^2 - y^2\|_{L^2}^2)$ with $\gamma_j^2 = \text{Median}\{\|a^j - b^j\|_{L^2}^2 : a, b \in \{x_i\}_{i=1}^{n_x} \cup \{y_i\}_{i=1}^{n_y}, a \neq b\}$ for $j = 1, 2$.

DIFFERENCE OF MEAN

We compare to the Functional Anderson-Darling (FAD) test in Pomann et al. (2016) which involves computing functional principal components and then doing multiple Anderson-Darling tests. Independent realisations $\{x_i\}_{i=1}^{n_x}$ and $\{y_j\}_{j=1}^{n_y}$ of the random functions x, y over $[0, 1]$ are observed on a grid of 100 uniform points with $n_x = n_y = 100$ and observation noise $\mathcal{N}(0, 0.25)$. The two distributions are

$$\begin{aligned} x(t) &\sim t + \xi_{10}\sqrt{2}\sin(2\pi t) + \xi_5\sqrt{2}\cos(2\pi t), \\ y(t) &\sim t + \delta t^3 + \eta_{10}\sqrt{2}\sin(2\pi t) + \eta_5\sqrt{2}\cos(2\pi t), \end{aligned}$$

with $\xi_5, \eta_5 \stackrel{i.i.d}{\sim} \mathcal{N}(0, 5)$ and $\xi_{10}, \eta_{10} \stackrel{i.i.d}{\sim} \mathcal{N}(0, 10)$. The δ parameter measures the deviation from the null hypothesis that x, y have the same distribution. The range of the parameter is $\delta \in \{0, 0.5, 1, 1.5, 2\}$.

Figure 2 shows CEXP performing best among all the choices which makes sense since this choice explicitly smooths the signal to make the mean more identifiable compared to the noise. We see that FPCA performs poorly because the principal components are deduced entirely from the covariance structure and do not represent the mean difference well. Likewise COV performs poorly since it can only detect difference in covariance, not mean. Except from FPCA and COV all choices of T out perform the FAD method. This is most likely because the FAD method involves computing multiple principle components, an estimation which is inherently random and results in components which do not identify the mean shift well, and computes multiple FAD tests with a Bonferroni correction which can cause too harsh a requirement for significance. The empirical sizes of each method are all very close to the nominal 5% level, see Table 4 in the Appendix for the values used to plot Figure 2.

Figure 3 shows an ROC curve. On the x -axis is α the false positive rate parameter in the test, see Section 3, and on the y -axis is the power of the test, meaning the true positive rate. The plot was obtained using $\delta = 1.25, n_x = 50, n_y = 50$ with the same observation locations and noise as described above. We see that COV and FPCA performs trivially weakly implying the calculated principal components are uninformative for identifying the difference in mean. CEXP has the best curve ID, SQR perform equally well.




Figure 2: Test power under mean difference for different kernels.

DIFFERENCE OF VARIANCE

We investigate two synthetic data sets, the first from Pomann et al. (2016) and the second from Paparoditis and Sapatinas (2016). The first represents a difference in covariance in a specific frequency and the second a difference across all frequencies.




Figure 3: ROC curve for different kernels.

In the first data set $n_x = n_y = 100$, observations are made on a uniform grid of 100 points and the observation noise is $\mathcal{N}(0, 0.25)$. The two distributions are

$$\begin{aligned} x(t) &\sim \xi_{10}\sqrt{2}\sin(2\pi t) + \xi_5\sqrt{2}\cos(2\pi t), \\ y(t) &\sim \eta_{10+\delta}\sqrt{2}\sin(2\pi t) + \eta_5\sqrt{2}\cos(2\pi t), \end{aligned}$$

with $\xi_5, \eta_5 \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 5)$ and $\xi_{10} \sim \mathcal{N}(0, 10)$ and $\eta_{10+\delta} \sim \mathcal{N}(0, 10+\delta)$. Therefore the difference in covariance structure is manifested in the first frequency. The range of the parameter is $\delta \in \{0, 5, 10, 15, 20\}$ and we again compare against the FAD test.

Figure 4 shows that COV performs the best which is to be expected since it is specifically designed to only detect change in covariance. SQR and FPCA perform well since they are designed to capture covariance information too. CEXP performs almost identically to ID since it is designed to improve performance on mean shift tests, not covariance shift. The values of the plot are in Table 5 in the Appendix.

The second dataset is from Paparoditis and Sapatinas (2016) and we compare against the data reported there of a bootstrap Hilbert-Schmidt norm (BOOT-HS) test (Paparoditis and Sapatinas, 2016, Section 2.2) and a functional principal component chi-squared (FPCA- χ) test (Fremdt et al., 2012), which is similar to the test in Panaretos et al. (2010). The number of function samples is $n_x = n_y = 25$ and each sample is observed on a uniform grid over $[0, 1]$ consisting of 500 points. The first distribution is defined as

$$x(t) \sim \sum_{n=1}^{10} \xi_n n^{-\frac{1}{2}} \sqrt{2} \sin(\pi n t) + \eta_n n^{-\frac{1}{2}} \sqrt{2} \cos(\pi n t),$$

where ξ_n, η_n are i.i.d. Student's t -distribution random variables with 5 degrees of freedom. For $\delta \in \mathbb{R}$ the other function distribution is $y \sim \delta x'$ where x' is an i.i.d. copy of x . When $\delta = 1$ the two distributions are the same. The entire covariance structure of Y is different from that of X when $\delta \neq 1$ which is in contrast the previous numerical example where the




Figure 4: Test power under variance difference in one frequency for different kernels.

covariance structure differed at only one frequency. The range of the deviation parameter is $\delta \in \{1, 1.2, 1.4, 1.6, 1.8, 2\}$.

Figure 5 shows again that COV and SQR performs the best. The BOOT-HS and FPCA- χ tests are both conservative, providing rejection rates below 5% when the null is true as opposed to the kernel based tests which all lie at or very close to the 5% level. The values of the plot are in Table 6 in the Appendix.




Figure 5: Test power under variance difference across all frequencies for different kernels.

DIFFERENCE OF HIGHER ORDERS

Data from Hall and Keilegom (2007) is used when performing the test. The random functions x, y are distributed as

$$x(t) \sim \sum_{n=1}^{15} e^{-\frac{n}{2}} \xi_n^x \psi_n(t),$$

$$y(t) \sim \sum_{n=1}^{15} e^{-\frac{n}{2}} \xi_{n,1}^y \psi_n(t) + \delta \sum_{n=1}^{15} n^{-2} \xi_{n,2}^y \psi_n^*(t),$$

with $\xi_n^x, \xi_{n,1}^y, \xi_{n,2}^y \stackrel{i.i.d}{\sim} \mathcal{N}(0, 1)$, $\psi_1(t) = 1$, $\psi_n(t) = \sqrt{2} \sin((k-1)\pi t)$ for $n > 1$ and $\psi_1^*(t) = 1$, $\psi_n^*(t) = \sqrt{2} \cos((k-1)\pi(2t-1))$ if $n > 1$ is even, $\psi_n^*(t) = \sqrt{2} \sin((k-1)\pi(2t-1))$ if $n > 1$ is odd. The observation noise for x is $\mathcal{N}(0, 0.01)$ and for y is $\mathcal{N}(0, 0.09)$. The range of the parameter is $\delta \in \{0, 1, 2, 3, 4\}$ and we compare against the FAD test and the Cramer-von Mises test in Hall and Keilegom (2007). The number of samples is $n_x = n_y = 15$ and for each random function 20 observation locations are sampled randomly according to p_x or p_y with p_x being the uniform distribution on $[0, 1]$ and p_y the distribution with density function $0.8 + 0.4t$ on $[0, 1]$.

Since the data is noisy and irregularly sampled, curves were fit to the data before the test was performed. The posterior mean of a Gaussian process with noise parameter $\sigma^2 = 0.01$ was fit to each data sample using a Matérn-1.5 kernel $k_{\text{Mat}}(s, t) = (1 + \sqrt{3}(s-t))e^{-\sqrt{3}(s-t)}$.

Figure 6 shows that the COV, SQR perform the best with other choices of T performing equally. Good power is still obtained against the existing methods despite the function reconstructions, validating the theoretical results of Section 6. The values of the plot are in Table 7 in the Appendix.




Figure 6: Test power under difference of higher orders for different kernels.

8.3 Real Data

BERKELEY GROWTH DATA

We now perform tests on the Berkeley growth dataset which contains the height of 39 male and 54 female children from age 1 to 18. The data can be found in the R package `fda`. We perform the two sample test on this data for the five different choices of T with γ chosen via the median heuristic outlined in the previous subsection. To identify the effect on test performance of sample size we perform random subsampling of the datasets and repeat the test to calculate test power. For each sample size $M \in \{5, 10, 15, 20\}$ we sample M functions from each data set and perform the test, this is repeated 500 times to calculate test power. The results are plotted in Figure 7. Similarly, to investigate the size of the test we sample two disjoint subsets of size $M \in \{5, 10, 15, 20\}$ from the female data set and perform the test and record whether the null was incorrectly rejected, this is repeated 500 times to obtain a rate of incorrect rejection of the null, the results are reported in Table 2.

Figure 7 shows ID, SQR performing the best, COV performs weaker than CEXP suggesting that it is not just a difference of covariance operator that distinguishes the two samples. Table 2 shows nearly all the tests have the correct empirical size.




Figure 7: Test power under subsamples of size M using Berkeley growth data.

NEU STEEL DATA

We perform the two-sample test on two classes from the North Eastern University steel defect dataset (Song and Yan, 2013; He et al., 2020; Dong et al., 2019). The dataset consists of 200×200 pixel grey scale images of defects of steel surfaces with 6 different classes of defects and 300 images in each class. We perform the test on the two classes which are most visually similar, called rolled-in scale and crazing. See the URL (Song and Yan, 2020) for further description of the dataset. For each sample size $M \in \{5, 10, 15, 20\}$ we sample M images from each class and perform the test, this is repeated 500 times to calculate test power. Again we assess the empirical size by sampling two distinct subsets from one class, the rolled-in class, for sample sizes $M \in \{5, 10, 15, 20\}$ and repeat this 500

M	ID	CEXP	SQR	FPCA	COV
5	4.6	4.4	4.4	3.6	4.8
10	5.2	5.4	5.0	5.6	5.4
15	4.6	4.8	5.6	4.8	7.2
20	5.2	5.4	6.0	6.4	5.6

Table 2: Empirical size, meaning the rate of rejection of the null when the null is true, of the two-sample test performed on the female Berkeley growth data for different sample sizes across different choices of T . The values are written as percentages, a value above 5 shows too frequent rejection and below shows too conservative a test.

times and report the rate of incorrect null rejection. For CEXP we use the two dimensional tensor product kernel induced by the CEXP kernel with 20 frequencies and $l = \sqrt{10}/200$ to normalise the size of the images.

Figure 8 shows SQR having the best performance, CEXP performs well and so does ID. Table 3 shows that the empirical size is inflated under some choices of T including SQR which should be taken into consideration when considering overall performance.




Figure 8: Test power under subsamples of size M using NEU steel data.

9. Conclusion

In this paper we studied properties of kernels on real, separable Hilbert spaces, and the associated Maximum Mean Discrepancy distance. Based on this, we formulated a novel kernel-based two-sample testing procedure for functional data. The development of kernels on Hilbert spaces was motivated by the observation that certain scaling of kernel parameters in the finite dimensional regime can result in a kernel over a Hilbert space. Indeed, multiple theoretical properties emerged as natural infinite dimensional generalisations of the finite

M	ID	CEXP	SQR	FPCA	COV
5	4.4	4.8	5.6	4.2	3.8
10	5.8	5.8	6.0	5.8	6.6
15	6.4	5.8	5.8	6.6	6.6
20	5.0	6.0	5.2	5.8	6.0

Table 3: Empirical size, meaning the rate of rejection of the null when the null is true, of the two-sample test performed on the rolled-in scale class from the NEU steel defect data, for different sample sizes across different choices of T . The values are written as percentages, a value above 5 shows too frequent rejection and below shows too conservative a test.

dimensional case. The development of kernels defined directly over Hilbert spaces facilitates the use of hyperparameters adapted for functional data, such as the choice of T in the SE- T kernel, which can result in greater test power.

While other nonparametric two-sample tests for functional data have been proposed recently, we believe that kernel-based approaches offer unique advantages. In particular, the ability to choose the kernel to reflect *a priori* knowledge about the data, such as any underlying dependencies, or to emphasise at which spatial scales the comparison should be made between the samples can be of significant benefit to practitioners. The construction of kernels which are tailor-made for two-sample testing of specific forms of functional data, for example time series and spatial data, is an interesting and open question, which we shall defer to future work.

The theory of kernels on function spaces is of independent interest and our work highlights how existing results on probability measures on infinite dimensional spaces can be applied to kernel methods, for example the use of the Minlos-Sazanov theorem when proving characteristicness. Recent theoretical developments relating to kernels and MMD on general topological spaces in the absence of compactness or local compactness have revealed the challenges in establishing important properties of such metrics, for example determining weak convergence of sequences of probability measures (Simon-Gabriel and Schölkopf, 2018; Simon-Gabriel et al., 2020; Chevyrev and Oberhauser, 2018), which have important implications for the development of effective MMD-based tests for functional data.

A further application of kernels on function spaces is statistical learning of maps between function spaces, in particular, the challenge of learning surrogate models for large-scale PDE systems which can be viewed as a nonlinear deterministic maps from an input function space, initial or boundary data, to an output function space, a system response. Here there are fundamental challenges to be addressed relating to the universality properties of such kernels. Preliminary work in (Nelsen and Stuart, 2020) indicates that this is a promising direction of research.

Acknowledgments

GW was supported by an EPSRC Industrial CASE award [EP/S513635/1] in partnership with Shell UK Ltd. AD was supported by the Lloyds Register Foundation Programme on Data Centric Engineering and by The Alan Turing Institute under the EPSRC grant [EP/N510129/1]. We thank Sebastian Vollmer for helpful comments.

Appendix A. Proofs of Theoretical Results

A.1 Bochner and Minlos-Sazanov Theorem

Bochner's theorem provides an exact relationship between continuous, translation invariant kernels on \mathbb{R}^d , meaning $k(x, y) = \phi(x - y)$ for some continuous ϕ , and the Fourier transforms of finite Borel measures on \mathbb{R}^d . For a proof see Wendland (2005, Theorem 6.6).

Theorem 23 (Bochner) *A continuous function $\phi: \mathbb{R}^d \rightarrow \mathbb{C}$ is positive definite if and only if it is the Fourier transform of a finite Borel measure μ_ϕ on \mathbb{R}^d*

$$\hat{\mu}_\phi(x) := \int_{\mathbb{R}^d} e^{ix^T y} d\mu_\phi(y) = \phi(x).$$

Bochner's theorem does not continue to hold in infinite dimensions, for example the kernel $k(x, y) = e^{-\frac{1}{2}\|x-y\|_{\mathcal{X}}^2}$ when \mathcal{X} is an infinite dimensional, real, separable Hilbert space is not the Fourier transform of a finite Borel measure on \mathcal{X} (Maniglia and Rhandi, 2004, Proposition 1.2.11). Instead, a stronger continuity property is required, this is the content of the Minlos-Sazanov theorem. For a proof see (Maniglia and Rhandi, 2004, Theorem 1.1.5) or Vakhania et al. (1987, Theorem VI.1.1).

Theorem 24 (Minlos-Sazanov) *Let \mathcal{X} be a real, separable Hilbert space and $\phi: \mathcal{X} \rightarrow \mathbb{C}$ a positive definite function on \mathcal{X} then the following are equivalent*

1. ϕ is the Fourier transform of a finite Borel measure on \mathcal{X}
2. There exists $C \in L_1^+(\mathcal{X})$ such that ϕ is continuous with respect to the norm induced by C given by $\|x\|_C^2 = \langle Cx, x \rangle_{\mathcal{X}}$.

The existence of such an operator is a much stronger continuity property than standard continuity on \mathcal{X} and will be crucial in the proof of Theorem 26, one of our main results. To see that continuity with respect to such a C is stronger than usual continuity consider the following example. Fix any $\varepsilon > 0$ and assume we only know that $\phi: \mathcal{X} \rightarrow \mathbb{R}$ is continuous and for simplicity assume that $\phi(0) = 1$, then we know there exists some $\delta > 0$ such that $\|x\|_{\mathcal{X}} < \delta$ implies $|\phi(x) - 1| < \varepsilon$ meaning we have control of $\phi(x)$ over the bounded set $\|x\|_{\mathcal{X}} < \delta$. On the other hand, if ϕ is continuous with respect to $\|\cdot\|_C$ for some $C \in L_1^+(\mathcal{X})$ then we know there exists some $\delta' > 0$ such that $\|x\|_C < \delta'$ implies $|\phi(x) - 1| < \varepsilon$ so we have control of $\phi(x)$ over the *unbounded* set $\|x\|_C < \delta'$. To see this set is unbounded let $\{\lambda_n, e_n\}_{n=1}^\infty$ be the orthonormal eigensystem of C and for $n \in \mathbb{N}$ let $y_n = \frac{\delta' e_n}{2\lambda_n}$ if $\lambda_n > 0$ otherwise $y_n = ne_n$, then since $C \in L_1^+(\mathcal{X})$ we know $\lambda_n \rightarrow 0$ so $\|y_n\|_{\mathcal{X}} \rightarrow \infty$. Since we used elements from the eigensystem it is clear that $\|y_n\|_C \leq \delta'/2$. Therefore we have constructed a subset of the ball with respect to $\|\cdot\|_C$ of radius δ' that has unbounded norm.

A.2 Proofs for Section 4

Proof [Proposition 2]

We begin by outlining the scaling that occurs for each of the two cases. For economy of notation we set $\gamma = \gamma_N$. If $k_0(s, t) = \delta_{st}$ then $\text{Tr}(\Sigma^i) = N$ for all $i \in \mathbb{N}$ and $\langle m_N, \Sigma m_N \rangle = \|m_N\|_2^2 = \Theta(N)$ by the Riemann scaling. If k_0 is continuous and bounded then $\text{Tr}(\Sigma^i) = \Theta(N^i)$ since $N^{-i} \text{Tr}(\Sigma^i) \rightarrow \text{Tr}(C_{k_0}^i)$ and $\langle m_N, \Sigma^i m_N \rangle = \Theta(N^{i+1})$ since $N^{-(i+1)} \langle m_N, \Sigma^i m_N \rangle \rightarrow \int_{\mathcal{D}} \int_{\mathcal{D}} m(s) C_{k_0}^i m(t) ds dt$.

We shall outline the proof of the result in the second case, the proof for the first case is entirely analogous and is completed by substituting the above scaling results where appropriate. Throughout this proof we shall repeatedly use the following Taylor series approximation, for a positive definite matrix $\Sigma \in \mathbb{R}^{N \times N}$, based on the classical Girard-Waring formulas.

$$\det(I + \varepsilon \Sigma) = 1 + \varepsilon \text{Tr}(\Sigma) + \frac{\varepsilon^2}{2} R_1(s) = 1 + \varepsilon \text{Tr}(\Sigma) + \frac{\varepsilon^2}{2} (\text{Tr}(\Sigma)^2 - \text{Tr}(\Sigma^2)) + \frac{\varepsilon^3}{6} R_2(s'),$$

for some $0 < s, s' \leq \varepsilon$ where

$$R_1(s) = \left(\sum \frac{\lambda_i}{(1 + s\lambda_i)} \right)^2 - \det(I + s\Sigma) \sum \frac{\lambda_i^2}{(1 + s\lambda_i)^2},$$

and

$$\begin{aligned} R_2(s) &= -2(1 + \det(I + s\Sigma)) \sum \frac{\lambda_i}{(1 + s\lambda_i)} \sum \frac{\lambda_i^2}{(1 + s\lambda_i)^2} \\ &\quad + 2 \det(I + s\Sigma) \sum \frac{\lambda_i^3}{(1 + s\lambda_i)^3}, \end{aligned}$$

where $\{\lambda_1, \dots, \lambda_N\}$ are the eigenvalues of Σ . We may bound these remainders

$$R_1(s) \leq \left(\sum \frac{\lambda_i}{(1 + s\lambda_i)} \right)^2 \leq \text{Tr}(\Sigma)^2 = O(N^2),$$

and

$$R_2(s) \leq 2 \det(I + s\Sigma) \sum \frac{\lambda_i^3}{(1 + s\lambda_i)^3} \leq 2e^{\varepsilon \text{Tr}(\Sigma)} \text{Tr}(\Sigma^3).$$

In our scenario $\varepsilon = \gamma^{-2} = O(N^{-2\alpha}) = o(N^{-1})$ therefore $e^{\varepsilon \text{Tr}(\Sigma)}$ is bounded since $\text{Tr}(\Sigma) = O(N)$. Finally, $\text{Tr}(\Sigma^3) = O(N^3)$ therefore $R_2(s) = O(N^3)$. Setting $T = \frac{1}{\gamma} I$ and $R = S = \Sigma$ in Corollary 19 and applying the above Taylor expansion along with the expansions for $x^{-1/2}$ and $(I + \varepsilon \Sigma)^{-1}$ with Lagrange remainder term

$$\begin{aligned} \text{MMD}_k^2(P, Q) &= 2 \det(I + 2\Sigma/\gamma^2)^{-\frac{1}{2}} \left(1 - e^{-\frac{1}{2\gamma^2} \langle (I + 2\Sigma/\gamma^2)^{-1} \mu, \mu \rangle} \right) \\ &= \left(1 + \frac{1}{\gamma^2} r_1 \right) \left(\frac{1}{\gamma^2} \langle (I + 2\Sigma/\gamma^2)^{-1} \mu, \mu \rangle + \frac{1}{\gamma^4} r_2 \right) \\ &= \left(1 + \frac{1}{\gamma^2} r_1 \right) \left(\frac{1}{\gamma^2} \langle \mu, \mu \rangle + \frac{1}{\gamma^4} r_2 + \frac{1}{\gamma^4} r_3 \right), \end{aligned}$$

where r_1, r_2, r_3 are remainder terms.

To leading order $r_1/\gamma^2 = O(N^{1-2\alpha})$, $r_2/\gamma^4 = O(N^{2-4\alpha})$, $r_3/\gamma^4 = O(N^{2-4\alpha})$ so it follows that $\text{MMD}_k^2(P, Q) = \gamma^{-2}\|m_N\|_2 + r(N)$ where, to leading order, $r(N) = O(N^{2-4\alpha})$. We now consider the denominator. To this end, applying Theorem 20 with $T = \gamma^{-1}I$ and $S = \Sigma$ we obtain

$$\begin{aligned}\xi_2 &= 2 \det(I + 4\Sigma/\gamma^2)^{-\frac{1}{2}} \left(1 + e^{-\frac{1}{\gamma^2}\langle \mu, (I+4\Sigma/\gamma^2)^{-1}\mu \rangle} \right) \\ &\quad - 2 \det(I + 2\Sigma/\gamma^2)^{-1} \left(1 + e^{-\frac{1}{\gamma^2}\langle \mu, (I+2\Sigma/\gamma^2)^{-1}\mu \rangle} - 4e^{-\frac{1}{2\gamma^2}\langle \mu, (I+2\Sigma/\gamma^2)^{-1}\mu \rangle} \right) \\ &\quad - 8 \det((I + \Sigma/\gamma^2)(I + 3\Sigma/\gamma^2))^{-1/2} e^{-\frac{1}{2\gamma^2}\langle \mu, (I+2\Sigma/\gamma^2)(I+3\Sigma/\gamma^2)^{-1}(I+\Sigma/\gamma^2)^{-1}\mu \rangle}.\end{aligned}$$

We split this into two terms, $\xi_2 = A_1 + A_2$, where

$$A_1 = 4 \det(I + 4\Sigma/\gamma^2)^{-\frac{1}{2}} + 4 \det(I + 2\Sigma/\gamma^2)^{-1} - 8 \det((I + \Sigma/\gamma^2)(I + 3\Sigma/\gamma^2))^{-1/2},$$

and

$$\begin{aligned}A_2 &= 2 \det(I + 4\Sigma/\gamma^2)^{-\frac{1}{2}} \left(-1 + e^{-\frac{1}{\gamma^2}\langle \mu, (I+4\Sigma/\gamma^2)^{-1}\mu \rangle} \right) \\ &\quad - 2 \det(I + 2\Sigma/\gamma^2)^{-1} \left(3 + e^{-\frac{1}{\gamma^2}\langle \mu, (I+2\Sigma/\gamma^2)^{-1}\mu \rangle} - 4e^{-\frac{1}{2\gamma^2}\langle \mu, (I+2\Sigma/\gamma^2)^{-1}\mu \rangle} \right) \\ &\quad - 8 \det((I + \Sigma/\gamma^2)(I + 3\Sigma/\gamma^2))^{-1/2} \left(-1 + e^{-\frac{1}{2\gamma^2}\langle \mu, (I+2\Sigma/\gamma^2)(I+3\Sigma/\gamma^2)^{-1}(I+\Sigma/\gamma^2)^{-1}\mu \rangle} \right).\end{aligned}$$

Applying Taylor's theorem for the determinant and exponential terms we obtain:

$$\begin{aligned}A_1 &= 4 \left[1 - \frac{2}{\gamma^2} \text{Tr}(\Sigma) + \frac{2}{\gamma^4} (\text{Tr}(\Sigma))^2 + \frac{4}{\gamma^4} \text{Tr}(\Sigma^2) + r_1 \right] \\ &\quad + 4 \left[1 - \frac{2}{\gamma^2} \text{Tr}(\Sigma) + \frac{2}{\gamma^4} (\text{Tr}(\Sigma))^2 + \frac{2}{\gamma^4} \text{Tr}(\Sigma^2) + r_2 \right] \\ &\quad - 8 \left[1 - \frac{1}{2\gamma^2} \text{Tr}(\Sigma) + \frac{1}{8\gamma^4} \text{Tr}(\Sigma)^2 + \frac{1}{4\gamma^4} \text{Tr}(\Sigma^2) + r_3 \right] \\ &\quad \times \left[1 - \frac{3}{2\gamma^2} \text{Tr}(\Sigma) + \frac{9}{8\gamma^4} (\text{Tr}(\Sigma))^2 + \frac{9}{4\gamma^4} \text{Tr}(\Sigma^2) + r_4 \right],\end{aligned}$$

where to leading order $r_1, r_2, r_3, r_4 = O(N^{3-6\alpha} + N^{4-8\alpha})$. After simplification we obtain $A_1 = \frac{4}{\gamma^4} \text{Tr}(\Sigma^2) + r(N)$ where $r(N) = O(N^{3-6\alpha})$. Similarly writing $A_2 = B_1 + B_2 + B_3$, we obtain:

$$\begin{aligned}B_1 &= \frac{2}{\gamma^2} (1 - 2\text{Tr}(\Sigma)/\gamma^2 + r_5) \left(-\langle \mu, \mu \rangle + \frac{4}{\gamma^2} \langle \mu, \Sigma\mu \rangle + \frac{1}{2\gamma^2} \langle \mu, \mu \rangle^2 + r_6 \right), \\ B_2 &= -\frac{2}{\gamma^2} (1 - 2\text{Tr}(\Sigma)/\gamma^2 + r_7) \left(\langle \mu, \mu \rangle - \frac{2}{\gamma^2} \langle \mu, \Sigma\mu \rangle + r_8 \right), \\ B_3 &= -\frac{8}{\gamma^2} (1 - 2\text{Tr}(\Sigma)/\gamma^2 + r_9) \left(-\frac{1}{2} \langle \mu, \mu \rangle + \frac{1}{\gamma^2} \langle \mu, \Sigma\mu \rangle + \frac{1}{8\gamma^2} \langle \mu, \mu \rangle^2 + r_{10} \right),\end{aligned}$$

where r_5, \dots, r_{10} are remainder terms which satisfy $r_i = O(N^{3-6\alpha} + N^{4-8\alpha})$ for $i = 5, \dots, 10$. Expanding $B_1 + B_2 + B_3$ and collecting into powers of γ , we see that the constant

terms and the terms with denominator γ^2 cancel out. Collecting terms with γ^4 denominator gives $A_2 = \frac{4}{\gamma^4} \langle \mu, \Sigma \mu \rangle + r(N)$ where $r(d) = O(d^{3-6\alpha})$ is a remainder term containing the higher order terms. Combining the leading order expressions for A_1 and A_2 , and collecting together the remainder terms we obtain

$$\frac{\text{MMD}_k^2(P, Q)}{\sqrt{\xi_2}} = \frac{\|\mu\|_2^2/\gamma^2 + r(N)}{\sqrt{4\text{Tr}(\Sigma^2)/\gamma^4 + 4\langle \mu, \Sigma \mu \rangle/\gamma^4 + r'(N)}},$$

where $r(N) = O(N^{2-4\alpha})$ and $r'(N) = O(N^{3-6\alpha})$. It follows that

$$\frac{\text{MMD}_k^2(P, Q)}{\sqrt{\xi_2}} = \frac{\|\mu\|_2^2}{\sqrt{4\text{Tr}(\Sigma^2) + 4\langle \mu, \Sigma \mu \rangle}} \left(\frac{1 + \frac{r(N)}{\|\mu\|_2^2/\gamma^2}}{\sqrt{1 + \frac{r'(N)}{4\text{Tr}(\Sigma^2)/\gamma^4 + 4\langle \mu, \Sigma \mu \rangle/\gamma^4}}} \right). \quad (11)$$

As discussed at the start of the proof $\|m_N\|_2^2 = \Theta(N)$ and $\text{Tr}(\Sigma^2), \langle m_N, \Sigma m_N \rangle = \Theta(N^2)$ meaning that $\gamma^2 r(N)/\|m_N\|_2^2 = O(N^{1-2\alpha})$ and so converges to zero. Likewise the fraction in the square root in the denominator converges to zero, meaning the term in the brackets of (11) converge to 1 as $N \rightarrow \infty$, yielding the advertised result once numerator and denominator is divided by N since

$$\begin{aligned} N^{-2} \text{Tr}(\Sigma^2) &\rightarrow \int_{\mathcal{D}} \int_{\mathcal{D}} k_0(s, t) ds dt = \|C_k\|_{HS}^2, \\ N^{-2} \langle m_N, \Sigma m_N \rangle &\rightarrow \int_{\mathcal{D}} \int_{\mathcal{D}} m(s) k_0(s, t) m(t) ds dt = \|C_k^{1/2} m\|_{L^2(\mathcal{D})}^2. \end{aligned}$$

For the case of identity matrix substituting the corresponding scaling of $\text{Tr}(\Sigma^2), \langle m_N, \Sigma m_N \rangle$ in the relevant places will produce the result. ■

Proof [Lemma 3]

We use the standard fact that for a real valued random variables η the following inequality holds $|\mathbb{E}[\eta] - \text{Median}[\eta]| \leq \text{Var}[\eta]^{\frac{1}{2}}$ and we will be using this inequality on $\|x - y\|_2^2$ where $x \sim P, y \sim Q$. For economy of notation let $\mu = \mu_1 - \mu_2$ and $\Sigma = \Sigma_1 + \Sigma_2$. Using standard Gaussian integral identities we obtain

$$\begin{aligned} \mathbb{E}[\|x - y\|_2^2]^2 &= (\text{Tr}(\Sigma) + \|\mu\|_2^2)^2, \\ \mathbb{E}[\|x - y\|_2^4] &= 2\text{Tr}(\Sigma^2) + 4\langle \mu, \Sigma \mu \rangle + (\text{Tr}(\Sigma) + \|\mu\|_2^2)^2, \end{aligned}$$

therefore

$$\text{Var}[\|x - y\|_2^2] = \mathbb{E}[\|x - y\|_2^4] - \mathbb{E}[\|x - y\|_2^2]^2 = 2\text{Tr}(\Sigma^2) + 4\|\mu\|_2^2 \text{Tr}(\Sigma).$$

Substituting into the inequality at the start of the proof

$$\left| \frac{\text{Median}[\|x - y\|_2^2]}{\mathbb{E}[\|x - y\|_2^2]} - 1 \right|^2 \leq \frac{2\text{Tr}(\Sigma^2) + 4\|\mu\|_2^2 \text{Tr}(\Sigma)}{(\text{Tr}(\Sigma) + \|\mu\|_2^2)^2} = 2 \left(1 - \frac{\|\mu\|_2^4}{(\text{Tr}(\Sigma) + \|\mu\|_2^2)^2} \right),$$

which completes the first part of the proof. The special case follows from dividing the numerator and denominator of the fraction in the right hand side by N^2 and using the assumption that the kernels are continuous and the mesh satisfies the Riemann scaling property. Writing m_N for the discretised version of $m_1 - m_2$ gives $N^{-2}\|m_N\|_2^4 \rightarrow \|m_1 - m_2\|_{L^2(\mathcal{D})}^2$, $N^{-1}\text{Tr}(\Sigma) \rightarrow \text{Tr}(C_{k_1} + C_{k_2})$, $N^{-1}\|m_N\|_2^2 \rightarrow \|m_1 - m_2\|_{L^2(\mathcal{D})}^2$. ■

A.3 Proofs for Section 5

Lemma 25 *The SE-T function is a kernel.*

Proof [Lemma 25] Consider first the case $T = I$, it is shown in Schoenberg (1938, Theorem 3) that if $k(x, y) = \phi(\|x - y\|_{\mathcal{Y}}^2)$ for ϕ a completely monotone function then k is positive definite on \mathcal{Y} and it is well known that e^{-ax} is such a function for $a > 0$ therefore k_I is a kernel. Now take k_T to be the SE-T kernel then for any $N \in \mathbb{N}$, $\{a_n\}_{n=1}^N \subset \mathbb{R}$, $\{x_n\}_{n=1}^N \subset \mathcal{X}$ we have $\sum_{n,m=1}^N a_n a_m k_T(x_n, x_m) = \sum_{n,m=1}^N a_n a_m k_I(T(x_n), T(x_m)) \geq 0$ ■

Proof [Theorem 5] This proof uses the argument of Minh (2009, Theorem 1). The plan is to first show the function space stated in the theorem is an RKHS and that the kernel is the SE-T kernel so by uniqueness of kernel for RKHS we are done. This is done using the Aronszajn theorem (Minh, 2009, Theorem 9) which identifies the kernel as an infinite sum of basis functions, see also Paulsen and Raghupathi (2016, Theorem 2.4).

First we prove that $\mathcal{H}_{k_T}(\mathcal{X})$ is a separable Hilbert space. That it is an inner product space is clear from the definition of the inner product and the assumption that $\lambda_n > 0$ for all $n \in \mathbb{N}$. The definition of the inner product means completeness of $\mathcal{H}_{k_T}(\mathcal{X})$ equivalent to completeness of the weighted l^2 space given by

$$l_{\lambda}^2(\Gamma) = \left\{ (w_{\gamma})_{\gamma \in \Gamma} : \|(w_{\gamma})_{\gamma \in \Gamma}\|_{l_{\lambda}^2(\Gamma)}^2 := \sum_{\gamma \in \Gamma} \frac{\gamma!}{\lambda^{\gamma}} w_{\gamma}^2 < \infty \right\}$$

which can be easily seen to be complete since Γ is countable and $\gamma!/\lambda^{\gamma}$ is positive for all $\gamma \in \Gamma$. To see that $\mathcal{H}_{k_T}(\mathcal{X})$ is separable observe that from the definition of the inner product, the countable set of functions $\phi_{\gamma}(x) = \sqrt{\lambda^{\gamma}/\gamma!} e^{-\frac{1}{2}\|Tx\|_{\mathcal{X}}^2} x^{\gamma}$ is orthonormal and spans $\mathcal{H}_{k_T}(\mathcal{X})$ hence is an orthonormal basis.

Next we prove that $\mathcal{H}_{k_T}(\mathcal{X})$ is an RKHS. Expanding the kernel through the exponential function gives

$$k_T(x, y) = e^{-\frac{1}{2}\|Tx\|_{\mathcal{X}}^2} e^{-\frac{1}{2}\|Ty\|_{\mathcal{X}}^2} e^{\langle Tx, Ty \rangle_{\mathcal{X}}} = e^{-\frac{1}{2}\|Tx\|_{\mathcal{X}}^2} e^{-\frac{1}{2}\|Ty\|_{\mathcal{X}}^2} \sum_{n=0}^{\infty} \frac{\langle Tx, Ty \rangle_{\mathcal{X}}^n}{n!},$$

and by the assumption on T we know $\langle Tx, Ty \rangle_{\mathcal{X}}^n = (\sum_{m=1}^{\infty} \lambda_m x_m y_m)^n = \sum_{|\gamma|=n} \frac{n!}{\gamma!} \lambda^{\gamma} x^{\gamma} y^{\gamma}$ where $x_m = \langle x, e_m \rangle_{\mathcal{X}}$, similarly for y_m , therefore

$$k_T(x, y) = e^{-\frac{1}{2}\|Tx\|_{\mathcal{X}}^2} e^{-\frac{1}{2}\|Ty\|_{\mathcal{X}}^2} \sum_{|\gamma| \geq 0} \frac{\lambda^{\gamma}}{\gamma!} x^{\gamma} y^{\gamma} = \sum_{|\gamma| \geq 0} \phi_{\gamma}(x) \phi_{\gamma}(y).$$

So for any $F \in \mathcal{H}_{k_T}(\mathcal{X})$ we have $\langle F, k_T(\cdot, x) \rangle_{\mathcal{H}_{k_T}(\mathcal{X})} = \sum_{\gamma \in \Gamma} \langle F, \phi_\gamma \rangle_{\mathcal{H}_{k_T}(\mathcal{X})} \phi_\gamma(x) = F(x)$ so k_T is a reproducing kernel of $\mathcal{H}_{k_T}(\mathcal{X})$ so by uniqueness of reproducing kernels we have that $\mathcal{H}_{k_T}(\mathcal{X})$ is the RKHS of k_T . \blacksquare

Proof [Theorem 7] If $k_T(x, y) = \hat{\mu}(x - y)$ for some Borel measure on \mathcal{X} then (2) lets us write

$$\begin{aligned} \text{MMD}_{k_T}(P, Q)^2 &= \int_{\mathcal{X}} \int_{\mathcal{X}} k_T(x, y) d(P - Q)(x) d(P - Q)(y) \\ &= \int_{\mathcal{X}} \int_{\mathcal{X}} \int_{\mathcal{X}} e^{i\langle h, x-y \rangle_{\mathcal{X}}} d\mu(h) d(P - Q)(x) d(P - Q)(y) \\ &= \int_{\mathcal{X}} \left(\int_{\mathcal{X}} e^{i\langle h, x \rangle_{\mathcal{X}}} d(P - Q)(x) \right) \left(\int_{\mathcal{X}} e^{-i\langle h, y \rangle_{\mathcal{X}}} d(P - Q)(y) \right) d\mu(h) \quad (12) \\ &= \int_{\mathcal{X}} \left| \hat{P}(h) - \hat{Q}(h) \right|^2 d\mu(h), \end{aligned}$$

where (12) is obtained by using Fubini's theorem to swap the integrals which is permitted since $|e^{i\langle h, x-y \rangle_{\mathcal{X}}}| = 1$ and is therefore integrable with respect to P and Q .

Fourier transforms of finite Borel measures on \mathcal{X} are uniformly continuous (Albeverio and Mazzucchi, 2015, Proposition 2.21) therefore if μ has full support, meaning that $\mu(U) > 0$ for every open $U \subset \mathcal{X}$, then we may immediately conclude that $\hat{P} = \hat{Q}$ and since the Fourier transform of finite Borel measures on \mathcal{X} is injective (Da Prato, 2006, Proposition 1.7) we may conclude that $P = Q$ meaning that Φ_{k_T} is injective. Assume T is injective. If k_T is the SE- T kernel then Da Prato (2006, Proposition 1.25) shows that $\mu = N_{T^2}$ has full support since T^2 is also injective. Therefore k_T is characteristic. For the converse direction we use the contrapositive and assume that T is not injective meaning there exists $x^* \in \mathcal{X}$ with $x^* \neq 0$ and $Tx^* = 0$. Set $P = \delta_0$ and $Q = \delta_{x^*}$ the Dirac measures on $0, x^*$ then $\Phi_{k_T}(P) = k_T(0, \cdot) = k_T(x^*, \cdot) = \Phi_{k_T}(Q)$ therefore Φ_{k_T} is not injective so k_T is not characteristic. \blacksquare

Proof [Theorem 9] The result is a corollary of the next theorem which is where Theorem 24 is employed.

Theorem 26 *Let \mathcal{X} be a real, separable Hilbert space and $T = I$ then the SE- T kernel is characteristic.*

Proof [Theorem 26] The idea of this proof is to use the contrapositive and assume $P \neq Q$ and conclude that $\text{MMD}_{k_I}(P, Q) > 0$. The main tool shall be Theorem 24 since $P \neq Q$ implies $\hat{P} \neq \hat{Q}$ and Theorem 24 implies that these Fourier transforms vary slowly in some sense so there will be a set of big enough measure such that $\hat{P}(x) \neq \hat{Q}(x)$ for x in this set, which will allow us to deduce $\text{MMD}_{k_I}(P, Q) > 0$.

Suppose P, Q are Borel probability measures on \mathcal{X} with $P \neq Q$ then $\hat{P} \neq \hat{Q}$ (Da Prato, 2006, Proposition 1.7) so there exists $x^* \in \mathcal{X}, \varepsilon > 0$ such that $|\hat{P}(x^*) - \hat{Q}(x^*)| > \varepsilon$. By Theorem 24 there exists $S, R \in L_1^+(\mathcal{X})$ such that \hat{P} (respectively \hat{Q}) is continuous with respect to the norm induced by S (respectively R). For $r > 0$ let $B_S(x^*, r) = \{x \in$

$\mathcal{X}: \langle S(x - x^*), x - x^* \rangle_{\mathcal{X}} < r^2$ be the ball based at x^* of radius r with respect to the norm induced by S , and $B_R(x^*, r)$ will denote the analogous ball with respect to the norm induced by R . By the continuity properties of \widehat{P}, \widehat{Q} there exists $r > 0$ such that

$$\begin{aligned} x \in B_S(x^*, r) &\implies |\widehat{P}(x) - \widehat{P}(x^*)| < \frac{\varepsilon}{4} \\ x \in B_R(x^*, r) &\implies |\widehat{Q}(x) - \widehat{Q}(x^*)| < \frac{\varepsilon}{4}. \end{aligned}$$

The set $B_S(x^*, r) \cap B_R(x^*, r)$ is non-empty since it contains x^* and if $x \in B_S(x^*, r) \cap B_R(x^*, r)$ then by reverse triangle inequality

$$\begin{aligned} |\widehat{P}(x) - \widehat{Q}(x)| &= |\widehat{P}(x) - \widehat{P}(x^*) + \widehat{P}(x^*) - \widehat{Q}(x^*) + \widehat{Q}(x^*) - \widehat{Q}(x)| \\ &\geq |\widehat{P}(x^*) - \widehat{Q}(x^*)| - |\widehat{P}(x) - \widehat{P}(x^*)| - |\widehat{Q}(x) - \widehat{Q}(x^*)| \\ &> \varepsilon - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} = \frac{\varepsilon}{2}. \end{aligned} \tag{13}$$

Now define the operator $U = S + R$ which is positive, symmetric and trace class since both S and R have these properties. Note that $B_U(x^*, r) \subset B_S(x^*, r) \cap B_R(x^*, r)$ because

$$\begin{aligned} \|x - x^*\|_U^2 &= \langle U(x - x^*), x - x^* \rangle_{\mathcal{X}} \\ &= \langle S(x - x^*), x - x^* \rangle_{\mathcal{X}} + \langle R(x - x^*), x - x^* \rangle_{\mathcal{X}} \\ &= \|x - x^*\|_S^2 + \|x - x^*\|_R^2. \end{aligned}$$

Since U is a positive, compact, symmetric operator there exists a decomposition into its eigenvalues $\{\lambda_n\}_{n=1}^{\infty}$, a non-negative sequence converging to zero, and eigenfunctions $\{e_n\}_{n=1}^{\infty}$ which form an orthonormal basis of \mathcal{X} . We will later need to associate a non-degenerate Gaussian measure with U . To this end define V to be the positive, symmetric, trace class operator with eigenvalues $\{\rho_n\}_{n=1}^{\infty}$ where $\rho_n = \lambda_n$ if $\lambda_n > 0$ otherwise $\rho_n = n^{-2}$, and eigenfunctions $\{e_n\}_{n=1}^{\infty}$ inherited from U . Then by construction V is injective, positive, symmetric and trace class. The V induced norm dominates the U induced norm therefore $B_V(x^*, r) \subset B_U(x^*, r)$ so for $x \in B_V(x^*, r)$ we have $|\widehat{P}(x) - \widehat{Q}(x)| > \varepsilon/2$.

Now we construct an operator which will approximate I , define the operator $I_m x = \sum_{n=1}^{\infty} \omega_n^{(m)} \langle x, e_n \rangle_{\mathcal{X}} e_n$ where $\omega_n^{(m)} = 1$ for $n \leq m$ and $\omega_n^{(m)} = n^{-2}$ for $n > m$ and $\{e_n\}_{n=1}^{\infty}$ is the eigenbasis of V , then $I_m \in L_1^+(\mathcal{X})$ for every $m \in \mathbb{N}$. It is easy to see $k_{I_m^{1/2}}$ converges pointwise to k_I as $m \rightarrow \infty$ since $e^{-\frac{1}{2}x}$ is a continuous function on \mathbb{R} and $\|x\|_{I_m^{1/2}}^2 \rightarrow \|x\|_I^2 = \|x\|_{\mathcal{X}}^2$ however clearly $I_m^{1/2}$ does not converge in operator norm to I since I is not a compact operator. Since $k_{I_m} \leq 1$ for all m we may use the bounded convergence theorem to obtain

$$\begin{aligned} \text{MMD}_{k_I}(P, Q)^2 &= \int_{\mathcal{X}} \int_{\mathcal{X}} k_I(x, y) d(P - Q)(x) d(P - Q)(y) \\ &= \lim_{m \rightarrow \infty} \int_{\mathcal{X}} \int_{\mathcal{X}} k_{I_m^{1/2}}(x, y) d(P - Q)(x) d(P - Q)(y) \\ &= \lim_{m \rightarrow \infty} \int_{\mathcal{X}} |\widehat{P}(x) - \widehat{Q}(x)|^2 dN_{I_m}(x), \end{aligned} \tag{14}$$

where (14) is by the same reasoning as in the proof of Theorem 7. In light of the lower bound we derived earlier over $B_V(x^*, r)$ of the integrand in (13)

$$\text{MMD}_{k_I}(P, Q)^2 = \lim_{m \rightarrow \infty} \int_{\mathcal{X}} |\widehat{P}(x) - \widehat{Q}(x)|^2 dN_{I_m}(x) \geq \lim_{m \rightarrow \infty} \int_{B_V(x^*, r)} \frac{\varepsilon^2}{4} dN_{I_m}(x),$$

so if we can lower bound $N_{I_m}(B_V(x^*, r))$ by a positive constant independent of m then we are done. This set is the ball with respect to $V \in L_1^+(\mathcal{X})$ which is a somehow large set, see the discussion after Theorem 24, and we will use a push-forward of measure argument.

Define $T(x) = x - x^*$ then $N_{I_m}(B_V(x^*, r)) = T_{\#}N_{I_m}(B_V(0, r))$ and (Da Prato, 2006, Proposition 1.17) tells us $T_{\#}N_{I_m}(B_V(0, r)) = N_{-x^*, I_m}(B_V(0, r))$. Next we note that $N_{-x^*, I_m}(B_V(0, r)) = V_{\#}^{\frac{1}{2}}N_{-x^*, I_m}(B(0, r))$ and (Da Prato, 2006, Proposition 1.18) tells us that

$$V_{\#}^{\frac{1}{2}}N_{-x^*, I_m}(B(0, r)) = N_{-V^{\frac{1}{2}}x^*, V^{\frac{1}{2}}I_mV^{\frac{1}{2}}}(B(0, r)).$$

For ease of notation let $y^* = -V^{\frac{1}{2}}x^*$ and since we choose to construct I_m from the eigenbasis of V we have $V_m x := V^{\frac{1}{2}}I_mV^{\frac{1}{2}}x = \sum_{n \in \mathbb{N}} \rho_n^{(m)} \langle x, e_n \rangle_{\mathcal{X}} e_n$ where $\rho_n^{(m)} = \rho_n$ for $n \leq m$ and $\rho_n^{(m)} = \rho_n n^{-2}$ for $n > m$ so $V_m \in L_1^+(\mathcal{X})$ and is injective for every $m \in \mathbb{N}$. We follow the proof of (Da Prato, 2006, Proposition 1.25) and define the sets

$$A_l = \left\{ x \in \mathcal{X} : \sum_{n=1}^l \langle x, e_n \rangle_{\mathcal{X}}^2 \leq \frac{r^2}{2} \right\} \quad B_l = \left\{ x \in \mathcal{X} : \sum_{n=l+1}^{\infty} \langle x, e_n \rangle_{\mathcal{X}}^2 < \frac{r^2}{2} \right\}.$$

Since V_m is non-degenerate for every $m \in \mathbb{N}$ we have that for every $l \in \mathbb{N}$ the events A_l, B_l are independent under N_{y^*, V_m} (Da Prato, 2006, Example 1.22) meaning $\forall m, l \in \mathbb{N}$ we have

$$N_{y^*, V_m}(B(0, r)) \geq N_{y^*, V_m}(A_l \cap B_l) = N_{y^*, V_m}(A_l)N_{y^*, V_m}(B_l),$$

and by the measure theoretic Chebyshev inequality, for every $l \in \mathbb{N}$

$$\begin{aligned} N_{y^*, V_m}(B_l) &\geq 1 - N_{y^*, V_m}(B_l^c) \geq 1 - \frac{2}{r^2} \sum_{n=l+1}^{\infty} \int_{\mathcal{X}} \langle x, e_n \rangle_{\mathcal{X}}^2 dN_{y^*, V_m} \\ &= 1 - \frac{2}{r^2} \left(\sum_{n=l+1}^{\infty} \rho_n^{(m)} + \langle y^*, e_n \rangle_{\mathcal{X}}^2 \right) \\ &\geq 1 - \frac{2}{r^2} \left(\sum_{n=l+1}^{\infty} \rho_n + \langle y^*, e_n \rangle_{\mathcal{X}}^2 \right). \end{aligned}$$

As the final line involves the tail of a finite sum with no dependency on m there exists an $L \in \mathbb{N}$ such that $N_{y^*, V_m}(B_L) > \frac{1}{2}$ for every $m \in \mathbb{N}$ and $l \geq L$. Note that for $m > L$ we have $N_{y^*, V_m}(A_L) = N_{y^*, V}(A_L)$ since A_L depends only on the first L coordinates and $\rho_n^{(m)} = \rho_n$ for $n \leq L$ if $m > L$. So for $m > L$

$$N_{y^*, V_m}(A_L) = N_{y^*, V}(A_L) \geq N_{y^*, V} \left(B \left(0, \frac{r}{\sqrt{2}} \right) \right) > c > 0,$$

for some c since V is non-degenerate (Da Prato, 2006, Proposition 1.25).

Overall we have shown that there exists an $L \in \mathbb{N}$ such that for $m > L$ we have $N_{I_m}(B_V(x^*, r)) = N_{y^*, V_m}(B(0, r)) > \frac{c}{2}$. Therefore, by substituting back into the lower bound for $\text{MMD}_{k_I}(P, Q)^2$

$$\begin{aligned} \text{MMD}_{k_I}(P, Q)^2 &= \lim_{m \rightarrow \infty} \int_{\mathcal{X}} |\widehat{P}(x) - \widehat{Q}(x)|^2 dN_{I_m}(x) \\ &\geq \lim_{m \rightarrow \infty} \int_{B_V(x^*, r)} \frac{\varepsilon^2}{4} dN_{I_m}(x) > \frac{\varepsilon^2 c}{8} > 0, \end{aligned}$$

which implies by contrapositive that k_I is characteristic. ■

With Theorem 26 proved we proceed to prove Theorem 9. By (2)

$$\begin{aligned} \text{MMD}_{k_T}(P, Q)^2 &= \int \int k_T(x, x') dP(x) dP(x') + \int \int k_T(y, y') dQ(y) dQ(y') \\ &\quad - 2 \int \int k_T(x, y) dQ(x) dP(y) \\ &= \int \int k_I(x, x') dT_{\#}P(x) dT_{\#}P(x') + \int \int k_I(y, y') dT_{\#}Q(y) dT_{\#}Q(y') \\ &\quad - 2 \int \int k_I(x, y) dT_{\#}P(x) dT_{\#}Q(y) \\ &= \text{MMD}_{k_I}(T_{\#}P, T_{\#}Q)^2. \end{aligned} \tag{15}$$

Using Theorem 26 we know that if $\text{MMD}_{k_T}(P, Q) = 0$ then $T_{\#}P = T_{\#}Q$ and all that is left to show is that the assumption on T implies $P = Q$. By the definition of push-forward measure we know that for every $B \in \mathcal{B}(\mathcal{Y})$ we have the equality $P(T^{-1}B) = Q(T^{-1}B)$. By the assumptions on $\mathcal{X}, \mathcal{Y}, T$ we know that $T(A) \in \mathcal{B}(\mathcal{Y})$ for every $A \in \mathcal{B}(\mathcal{X})$ (Kechris, 1995, Theorem 15.1). Hence for any $A \in \mathcal{B}(\mathcal{X})$ take $B = T(A)$ then $P(A) = P(T^{-1}B) = Q(T^{-1}B) = Q(A)$, which shows $P = Q$. ■

Proof [Proposition 11] For any $N \in \mathbb{N}$ take any $\{a_n\}_{n=1}^N \subset \mathbb{R}, \{x_n\}_{n=1}^N \subset \mathcal{X}$ then

$$\sum_{n,m=1}^N k_{C,k_0}(x_n, x_m) = \int_{\mathcal{X}} \sum_{n,m=1}^N k_0(\langle x_n, h \rangle_{\mathcal{X}}, \langle x_m, h \rangle_{\mathcal{X}}) dN_C(h) \geq 0,$$

since this is the integral of a non-negative quantity as k_0 is a kernel. Symmetry of k_{C,k_0} follows since k_0 is symmetric meaning k_{C,k_0} is a kernel. Expanding k_0 using its spectral measure we have

$$k_{C,k_0}(x, y) = \int_{\mathcal{X}} \int_{\mathbb{R}} e^{i\langle x-y, rh \rangle_{\mathcal{X}}} d\mu(r) dN_C(h) = \hat{\nu}(x - y),$$

where $\nu(A) = \int_{\mathcal{X}} \int_{\mathbb{R}} \mathbb{1}_A(rh) d\mu(r) dN_C(h)$ for all $A \in \mathcal{B}(\mathcal{X})$. This is the law of the \mathcal{X} valued random variable ξX where $\xi \sim \mu$ and $X \sim N_C$ are independent. We will show that ν has

full support on \mathcal{X} from which it follows that k_{C,k_0} is characteristic by following the proof of Theorem 7.

Fix any open ball $B = B(h, r) \subset \mathcal{X}$ then given the assumption on μ by intersecting with $(0, \infty)$ or $(-\infty, 0)$ we may assume that a, b have the same sign. Assume that $a, b > 0$, the proof for when $a, b < 0$ is analogous. We first treat the case $h \neq 0$. Set $\delta = \min(\frac{1}{2}(\frac{b}{a} - 1), \frac{r}{2\|h\|_{\mathcal{X}}})$ so that $(a, a(1 + \delta)) \subset (a, b)$. Now consider the ball $B' = B(\frac{h}{a(1+\delta)}, \frac{r}{4a(1+\delta)})$, take any $c \in (a, a(1 + \delta))$ and any $x \in B'$ then

$$\begin{aligned} \|cx - h\|_{\mathcal{X}} &\leq \left\| \xi x - \frac{\xi h}{a(1+\delta)} \right\|_{\mathcal{X}} + \left\| \frac{\xi h}{a(1+\delta)} - h \right\|_{\mathcal{X}} \\ &\leq \frac{cr}{4a(1+\delta)} + \|h\|_{\mathcal{X}} \left(1 - \frac{c}{a(1+\delta)} \right) \\ &< \frac{r}{4} + \|h\|_{\mathcal{X}} \left(1 - \frac{1}{1+\delta} \right) < \frac{r}{4} + \frac{r}{2} < r. \end{aligned}$$

Therefore for any $c \in (a, a(1 + \delta))$ we have $cB' \subset B$. Hence $\mathbb{P}(\xi X \in B) = \nu(B) \geq \mu((a, a(1 + \delta)))N_C(B') > 0$ by the assumptions on μ and the way N_C is non-degenerate.

The case $h = 0$ is analogous, take $B' = B(0, \frac{r}{2b})$ then for every $c \in (a, b)$ we have $cB' \subset B$ and we again conclude $\nu(B) > 0$. \blacksquare

Proof [Corollary 13] The idea of the proof is to represent the IMQ- T kernel as an integral of the SE- T kernel then use the same limit argument as in the proof of Theorem 26 and push-forward argument of Theorem 9.

Throughout this proof k_T^{IMQ} and k_T^{SE} will denote the IMQ- T and SE- T kernels respectively. By the same proof technique as Theorem 9 it suffices to prove that k_I^{IMQ} is characteristic. Let $\{e_n\}_{n=1}^{\infty}$ be any orthonormal basis of \mathcal{X} and let $I_m x = \sum_{n=1}^m \langle x, e_n \rangle_{\mathcal{X}} e_n$ so that I_m converges to I pointwise. Then by the same limiting argument in the proof of Theorem 26 using bounded convergence theorem, letting N_1 be the $\mathcal{N}(0, 1)$ measure on \mathbb{R} we have

$$\begin{aligned} \text{MMD}_{k_I^{\text{IMQ}}}(P, Q)^2 &= \lim_{m \rightarrow \infty} \int_{\mathcal{X}} \int_{\mathcal{X}} k_{I_m}^{\text{IMQ}}(x, y) d(P - Q)(x) d(P - Q)(y) \\ &= \lim_{m \rightarrow \infty} \int_{\mathcal{X}} \int_{\mathcal{X}} \int_{\mathbb{R}} k_{I_m}^{\text{SE}}(zx, zy) dN_1(z) d(P - Q)(x) d(P - Q)(y) \quad (16) \end{aligned}$$

$$= \int_{\mathbb{R}} \text{MMD}_{k_I^{\text{SE}}}(z_{\#}P, z_{\#}Q)^2 dN_1(z), \quad (17)$$

where (16) is from the integral representation of $k_{I_m}^{\text{IMQ}}$ in (10) which can be used since $I_m \in L_1^+(\mathcal{X})$, (17) is obtained by using Fubini's theorem and bounded convergence theorem and $z_{\#}P$ denotes the push-forward of the measure under the linear map from \mathcal{X} to \mathcal{X} defined

as multiplication by the scalar z . The integrand in (17) can be rewritten as

$$\begin{aligned}
 \text{MMD}_{k_I^{\text{SE}}}(z_{\#}P, z_{\#}Q)^2 &= \int_{\mathcal{X}} \int_{\mathcal{X}} k_I^{\text{SE}}(x, y) d(z_{\#}P - z_{\#}Q)(x) d(z_{\#}P - z_{\#}Q)(y) \\
 &= \int_{\mathcal{X}} \int_{\mathcal{X}} k_I^{\text{SE}}(zx, zy) d(P - Q)(x) d(P - Q)(y) \\
 &= \int_{\mathcal{X}} \int_{\mathcal{X}} e^{-\frac{z^2}{2} \|x-y\|_{\mathcal{X}}^2} d(P - Q)(x) d(P - Q)(y) \\
 &= \text{MMD}_{k_{z^2 I}^{\text{SE}}}(P, Q)^2,
 \end{aligned}$$

which makes it clear that $\text{MMD}_{k_I^{\text{SE}}}(z_{\#}P, z_{\#}Q)^2$ is a continuous, non-negative function of z and equals 0 if and only if $z = 0$. Using this we deduce

$$\begin{aligned}
 \text{MMD}_{k_I^{\text{IMQ}}}(P, Q)^2 &= \int_{\mathbb{R}} \text{MMD}_{k_I^{\text{SE}}}(z_{\#}P, z_{\#}Q)^2 dN_{\sigma^2}(z) \\
 &= \int_{\mathbb{R}} \text{MMD}_{k_{z^2 I}^{\text{SE}}}(P, Q)^2 dN_{\sigma^2}(z) > 0,
 \end{aligned}$$

since N_{σ^2} has strictly positive density. This proves that the IMQ- I kernel is characteristic and by the same push-forward argument as in the proof of Theorem 9 we may conclude that the IMQ- T kernel is characteristic too. \blacksquare

A.4 Proofs for Section 6

Proof [Proposition 14] For any i, j using the assumption on k we have

$$|k(x_i, y_j) - k(\mathcal{R}I x_i, \mathcal{R}I y_j)| \leq L \|x_i - y_j - \mathcal{R}I x_i + \mathcal{R}I y_j\|_{\mathcal{X}} \quad (18)$$

$$\leq L (\|\mathcal{R}I x_i - x_i\|_{\mathcal{X}} + \|\mathcal{R}I y_j - y_j\|_{\mathcal{X}}), \quad (19)$$

where (18) is by assumption and (19) uses the triangle inequality. Using (3) gives

$$\begin{aligned}
 &\left| \widehat{\text{MMD}}_k(X_n, Y_n)^2 - \widehat{\text{MMD}}_k(\mathcal{R}I X_n, \mathcal{R}I Y_n)^2 \right| \\
 &\leq \frac{1}{n(n-1)} \sum_{i \neq j}^n |h(z_i, z_j) - h(\mathcal{R}I z_i, \mathcal{R}I z_j)| \\
 &\leq \frac{2L}{n(n-1)} \sum_{i \neq j}^n \|\mathcal{R}I x_i - x_i\|_{\mathcal{X}} + \|\mathcal{R}I x_j - x_j\|_{\mathcal{X}} + \|\mathcal{R}I y_i - y_i\|_{\mathcal{X}} + \|\mathcal{R}I y_j - y_j\|_{\mathcal{X}} \quad (20)
 \end{aligned}$$

$$= \frac{4L}{n} \sum_{i=1}^n \|\mathcal{R}I x_i - x_i\|_{\mathcal{X}} + \|\mathcal{R}I y_i - y_i\|_{\mathcal{X}}, \quad (21)$$

where (20) follows from expanding using the definition of h in Section 3 and using the triangle inequality and (21) follows from counting the number of pairs of indices in the sum. \blacksquare

Proof [Corollary 15] This can be deduced by the Lipschitz constants of $e^{-x^2/2}$ and $(x^2 + 1)^{-1/2}$. Then the proof of Proposition 14 may be continued in the same manner. ■

Proof [Theorem 16] Define the two random variables

$$\begin{aligned} A_n &= n^{\frac{1}{2}} \left(\widehat{\text{MMD}}_k(\mathcal{R}IX_n, \mathcal{R}IY_n)^2 - \widehat{\text{MMD}}_k(X_n, Y_n)^2 \right) \\ B_n &= n^{\frac{1}{2}} \left(\widehat{\text{MMD}}_k(X_n, Y_n)^2 - \text{MMD}_k(P, Q)^2 \right). \end{aligned}$$

It is known $B_n \xrightarrow{d} \mathcal{N}(0, \xi)$ (Gretton et al., 2012a, Corollary 16) so the proof is complete by Slutsky's theorem if $A_n \xrightarrow{\mathbb{P}} 0$. Fix any $\varepsilon > 0$ then by Proposition 2

$$\begin{aligned} \mathbb{P}(|A_n| > \varepsilon) &\leq \mathbb{P}\left(\frac{4L}{n^{\frac{1}{2}}} \sum_{i=1}^n \|\mathcal{R}Ix_i - x_i\|_{\mathcal{X}} + \|\mathcal{R}Iy_i - y_i\|_{\mathcal{X}} > \varepsilon\right) \\ &\leq \frac{4L}{\varepsilon n^{\frac{1}{2}}} \mathbb{E}\left[\sum_{i=1}^n \|\mathcal{R}Ix_i - x_i\|_{\mathcal{X}} + \|\mathcal{R}Iy_i - y_i\|_{\mathcal{X}}\right] \end{aligned} \quad (22)$$

$$= \frac{4Ln^{\frac{1}{2}}}{\varepsilon} \mathbb{E}[\|\mathcal{R}Ix - x\|_{\mathcal{X}} + \|\mathcal{R}Iy - y\|_{\mathcal{X}}] \quad (23)$$

$$\rightarrow 0, \quad (24)$$

where (22) is by Markov's inequality, (23) is by the assumption that the samples from P, Q and discretisation U, V are i.i.d. across samples and (24) is by assumption. ■

Proof [Theorem 17] Note that $\Phi_{k_T}(N_{a,S})(x) = \Phi_{k_I}(N_{Ta,STS})(Tx)$ (Da Prato, 2006, Proposition 1.18). The proof simply uses (Da Prato and Zabczyk, 2002, Proposition 1.2.8) to calculate the Gaussian integrals.

$$\begin{aligned} \Phi_{k_I}(N_{a,S})(x) &= \int_{\mathcal{X}} e^{-\frac{1}{2}\langle x-y, x-y \rangle_{\mathcal{X}}} dN_{a,S}(y) \\ &= e^{-\frac{1}{2}\langle a-x, a-x \rangle_{\mathcal{X}}} \int_{\mathcal{X}} e^{-\frac{1}{2}\langle y, y \rangle_{\mathcal{X}}} e^{-\langle y, a-x \rangle_{\mathcal{X}}} dN_S(y) \\ &= \det(I + S)^{-\frac{1}{2}} e^{-\frac{1}{2}\langle a-x, a-x \rangle_{\mathcal{X}}} e^{\frac{1}{2}\langle (I+S)^{-\frac{1}{2}} S^{\frac{1}{2}}(a-x) \rangle_{\mathcal{X}}} \end{aligned} \quad (25)$$

$$\begin{aligned} &= \det(I + S)^{-\frac{1}{2}} e^{-\frac{1}{2}\langle a-x, a-x \rangle_{\mathcal{X}}} e^{\frac{1}{2}\langle S^{\frac{1}{2}}(I+S)^{-1} S^{\frac{1}{2}}(a-x), a-x \rangle_{\mathcal{X}}} \\ &= \det(I + S)^{-\frac{1}{2}} e^{-\frac{1}{2}\langle (I-S^{\frac{1}{2}}(I+S)^{-1} S^{\frac{1}{2}})(a-x), (a-x) \rangle_{\mathcal{X}}} \\ &= \det(I + S)^{-\frac{1}{2}} e^{-\frac{1}{2}\langle (I+S)^{-1}(x-a), x-a \rangle_{\mathcal{X}}}, \end{aligned} \quad (26)$$

where (25) is due to Da Prato and Zabczyk (2002, Proposition 1.2.8). The last equality is due to the Sherman-Morrison-Woodbury identity for operators (Hsing and Eubank, 2015, Theorem 3.5.6). Substituting in Ta for a , STS for S and Tx for x gives the desired expression, as discussed at the start of the proof. ■

Proof [Theorem 18] The idea of the proof is that $\text{MMD}_{k_T}(P, Q)^2$ is simply double integrals of the kernel with respect to Gaussian measures. One integral was completed in Theorem 17 and we apply Da Prato and Zabczyk (2002, Proposition 1.2.8) again. Note $\text{MMD}_{k_T}(N_{a,S}, N_{b,R})^2 = \text{MMD}_{k_I}(N_{Ta, TST}, N_{Tb, TRT})^2$ so it suffices to do the calculations for k_I and substitute the other values in. Also since k_T is translation invariant we may without loss of generality assume $a = 0$ and replace b with $a - b$ at the end.

$$\begin{aligned} & \int_{\mathcal{X}} \int_{\mathcal{X}} k_I(x, y) dN_S(x) dN_{b,R}(y) \\ &= \det(I + S)^{-\frac{1}{2}} \int_{\mathcal{X}} e^{-\frac{1}{2} \langle (I+S)^{-1}(y-b), y-b \rangle_{\mathcal{X}}} dN_R(y) \end{aligned} \quad (27)$$

$$\begin{aligned} &= \det(I + S)^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+S)^{-1}b, b \rangle_{\mathcal{X}}} \int_{\mathcal{X}} e^{-\frac{1}{2} \langle (I+S)^{-1}y, y \rangle_{\mathcal{X}}} e^{\langle y, (I+S)^{-1}b \rangle_{\mathcal{X}}} dN_R(y) \\ &= \det(I + S)^{-\frac{1}{2}} \det \left(I + R^{\frac{1}{2}}(I + S)^{-1}R^{\frac{1}{2}} \right)^{-\frac{1}{2}} \end{aligned} \quad (28)$$

$$\begin{aligned} & \times e^{-\frac{1}{2} \langle (I+S)^{-1}b, b \rangle_{\mathcal{X}}} e^{-\frac{1}{2} \langle (I+S)^{-1}R^{\frac{1}{2}}(I+R^{\frac{1}{2}}(I+S)^{-1}R^{\frac{1}{2}})^{-1}R^{\frac{1}{2}}(I+S)^{-1}b, b \rangle_{\mathcal{X}}} \\ &= \det(I + S)^{-\frac{1}{2}} \det \left(I + R^{\frac{1}{2}}(I + S)^{-1}R^{\frac{1}{2}} \right)^{-\frac{1}{2}} \end{aligned} \quad (29)$$

$$\begin{aligned} & \times e^{-\frac{1}{2} \langle (I-(I+S)^{-1}R^{\frac{1}{2}}(I+R^{\frac{1}{2}}(I+S)^{-1}R^{\frac{1}{2}})^{-1}R^{\frac{1}{2}}(I+S)^{-1}b, b \rangle_{\mathcal{X}}} \\ &= \det(I + S)^{-\frac{1}{2}} \det \left(I + R^{\frac{1}{2}}(I + S)^{-1}R^{\frac{1}{2}} \right)^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+S+R)^{-1}b, b \rangle_{\mathcal{X}}}, \end{aligned} \quad (30)$$

where (27) is obtained by substituting the result of Theorem 17, (28) is applying Da Prato and Zabczyk (2002, Proposition 1.2.8), (29) is just rearranging terms and (30) is using the Sherman-Morrison-Woodbury identity for operators (Hsing and Eubank, 2015, Theorem 3.5.6). The proof is completed by using the expression of MMD in terms of three double integrals and substituting in the appropriate values of S, R, b inline with the description at the start of the proof. In particular when $b = 0$ and $S = R$

$$\begin{aligned} \det(I + S) \det \left(I + S^{\frac{1}{2}}(I + S)^{-1}S^{\frac{1}{2}} \right) &= \det \left((I + S)(I + (I + S)^{-1}S) \right) \\ &= \det(I + 2S), \end{aligned}$$

by the Sherman-Morrison-Woodbury identity for operators. ■

Proof [Theorem 20]

A more general result for which Theorem 20 is a specific case shall be proved.

Theorem 27 *Let $P = N_{a,S}, Q = N_{b,R}$ be two non-degenerate Gaussian measures on \mathcal{X} , $C \in L^+(\mathcal{X})$ and assume C, S, R all commute then when using the SE-T kernel*

$$\begin{aligned} \xi_1 &= \alpha(T, S, R, a, b) + \alpha(T, R, S, a, b) \\ \xi_2 &= \beta(T, S, R, a, b) + \beta(T, R, S, a, b), \end{aligned}$$

where

$$\begin{aligned}
 \alpha(T, S, R, a, b) &= \det((I + TST)(I + 3TST))^{-\frac{1}{2}} - \det(I + 2TST)^{-1} \\
 &\quad + \det((I + TRT)(I + T(2S + R)T))^{-\frac{1}{2}} e^{-\langle (I+T(2S+R)T)^{-1}T(a-b), T(a-b) \rangle_{\mathcal{X}}} \\
 &\quad - \det(I + T(S + R)T)^{-1} e^{-\langle (I+T(S+R)T)^{-1}T(a-b), T(a-b) \rangle_{\mathcal{X}}} \\
 &\quad - 2 \det(\Sigma_S)^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+2TST)\Sigma_S^{-1}T(a-b), T(a-b) \rangle_{\mathcal{X}}} \\
 &\quad + 2 \det((I + 2TST)(I + T(S + R)T))^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+T(S+R)T)^{-1}T(a-b), T(a-b) \rangle_{\mathcal{X}}},
 \end{aligned}$$

$$\begin{aligned}
 \beta(T, S, R, a, b) &= \det(I + 4TST)^{-\frac{1}{2}} - \det(I + 2TST)^{-1} \\
 &\quad + \det(I + 2T(S + R)T)^{-\frac{1}{2}} e^{-\langle (I+2T(S+R)T)^{-1}(a-b), a-b \rangle_{\mathcal{X}}} \\
 &\quad - \det(I + T(S + R)T)^{-1} e^{-\langle (I+T(S+R)T)^{-1}(a-b), a-b \rangle_{\mathcal{X}}} \\
 &\quad + 4 \det((I + T(S + R)T)(I + 2TST))^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+T(S+R)T)^{-1}(a-b), a-b \rangle_{\mathcal{X}}} \\
 &\quad - 4 \det(\Sigma_S)^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+2TST)\Sigma_S^{-1}(a-b), a-b \rangle_{\mathcal{X}}},
 \end{aligned}$$

and $\Sigma_X = (I + TST)(I + TRT) + TXT(2I + T(S + R)T)$ for $X \in \{S, R\}$.

Proof [Theorem 27] As in the proof of Theorem 18 it suffices to consider $T = I$ and $a = 0$. Set $k = k_I$ and $\langle \cdot, \cdot \rangle = \langle \cdot, \cdot \rangle_{\mathcal{X}}$ for ease of notation. The expression for ξ_1 is derived first. The simplifications in Sutherland (2019) reveal

$$\begin{aligned}
 \xi_1 &= \mathbb{E}_x[\mathbb{E}_{x'}[k(x, x')]^2] - \mathbb{E}_{x,x'}[k(x, x')]^2 \\
 &\quad + \mathbb{E}_y[\mathbb{E}_{y'}[k(y, y')]^2] - \mathbb{E}_{y,y'}[k(y, y')]^2 \\
 &\quad + \mathbb{E}_x[\mathbb{E}_y[k(x, y)]^2] - \mathbb{E}_{x,y}[k(x, y)]^2 \\
 &\quad + \mathbb{E}_y[\mathbb{E}_x[k(y, x)]^2] - \mathbb{E}_{y,x}[k(y, x)]^2 \\
 &\quad - 2\mathbb{E}_x[\mathbb{E}_{x'}[k(x, x')]\mathbb{E}_y[k(x, y)]] + 2\mathbb{E}_{x,x'}[k(x, x')]\mathbb{E}_{x,y}[k(x, y)] \\
 &\quad - 2\mathbb{E}_y[\mathbb{E}_{y'}[k(y, y')]\mathbb{E}_x[k(y, x)]] + 2\mathbb{E}_{y,y'}[k(y, y')]\mathbb{E}_{x,y}[k(x, y)].
 \end{aligned}$$

To calculate this only three of the terms need to be calculated then the rest are deduced by substituting in different values. For example $\mathbb{E}_{x,x'}[k(x, x')]$ can be deduced from the formula for $\mathbb{E}_{x,y}[k(x, y)]$ by setting $b = 0$ and $S = R$ since this would make y acts as an independent copy of x in the expectation. The three terms needed are

$$\mathbb{E}_{x,y}[k(x, y)] \tag{31}$$

$$\mathbb{E}_x[\mathbb{E}_y[k(x, y)]^2] \tag{32}$$

$$\mathbb{E}_x[\mathbb{E}_{x'}[k(x, x')]\mathbb{E}_y[k(x, y)]] \tag{33}$$

Expression (31) was derived in the proof of Theorem 18 as

$$\mathbb{E}_{x,y}[k(x, y)] = \det(I + S + R)^{-\frac{1}{2}} e^{-\frac{1}{2} \langle (I+S+R)^{-1}b, b \rangle}.$$

Next a formula for (32) is derived. First note $\mathbb{E}_y[k(x, y)]$ is the content of Theorem 17. The rest follows by using (Da Prato and Zabczyk, 2002, Proposition 1.2.8) and rearranging terms.

$$\begin{aligned} \mathbb{E}_x[\mathbb{E}_y[k(x, y)]^2] &= \det(I + R)^{-1} \int_{\mathcal{X}} e^{-\langle (I+R)^{-1}(x-b), x-b \rangle} dN_S(x) \\ &= \det((I + R)(I + R + 2S))^{-\frac{1}{2}} e^{-\langle (I+R)^{-1}b, b \rangle} \\ &\quad \times e^{\langle 2S(I+R)^{-1}(I+2S(I+R)^{-1})^{-1}(I+R)^{-1}b, b \rangle} \\ &= \det((I + R)(I + R + 2S))^{-\frac{1}{2}} e^{-\langle (I+R+2S)^{-1}b, b \rangle}. \end{aligned}$$

Finally (33) is derived which involves the longest calculations. The terms in the first expectation are the content of Theorem 17.

$$\begin{aligned} &\mathbb{E}_x[\mathbb{E}_{x'}[k(x, x')]\mathbb{E}_y[k(x, y)]] \\ &= \det((I + S)(I + R))^{-\frac{1}{2}} \int_{\mathcal{X}} e^{-\frac{1}{2}\langle (I+S)^{-1}x, x \rangle} e^{-\frac{1}{2}\langle (I+R)^{-1}(x-b), x-b \rangle} dN_S(x) \\ &= \det((I + S)(I + R))^{-\frac{1}{2}} e^{-\frac{1}{2}\langle (I+R)^{-1}b, b \rangle} \\ &\quad \times \int_{\mathcal{X}} e^{-\frac{1}{2}\langle ((I+S)^{-1} + (I+R)^{-1})x, x \rangle} e^{-\frac{1}{2}\langle (I+R)^{-1}b, x \rangle} dN_S(x) \\ &= \det((I + S)(I + R))^{-\frac{1}{2}} \det(I + S((I + S)^{-1} + (I + R)^{-1}))^{-\frac{1}{2}} \\ &\quad \times e^{\frac{1}{2}\langle [(I+S((I+S)^{-1} + (I+R)^{-1}))^{-1}S(I+R)^{-1} - I]b, (I+R)^{-1}b \rangle} \\ &= \det(\Sigma_S)^{-\frac{1}{2}} e^{-\frac{1}{2}\langle (I+2S)\Sigma_S^{-1}b, b \rangle}, \end{aligned}$$

where $\Sigma_S = (I + S)(I + R) + S(2I + S + R)$. The last equality is obtained by rearranging the terms in the exponent and determinant. Substituting into the formula for ξ_1 the derivations for (31), (32) and (33) completes the derivation for ξ_1 . The simplification of ξ_2 in (Sutherland, 2019) is

$$\begin{aligned} \xi_2 &= \mathbb{E}_{x, x'}[k(x, x')^2] - \mathbb{E}_{x, x'}[k(x, x')]^2 + \mathbb{E}_{y, y'}[k(y, y')^2] - \mathbb{E}_{y, y'}[k(y, y')]^2 \\ &\quad + 2\mathbb{E}_{x, y}[k(x, y)^2] - 2\mathbb{E}_{x, y}[k(x, y)]^2 \\ &\quad - 4\mathbb{E}_x[\mathbb{E}_{x'}[k(x, x')]\mathbb{E}_y[k(x, y)]] + 4\mathbb{E}_{x, x'}[k(x, x')]\mathbb{E}_{x, y}[k(x, y)] \\ &\quad - 4\mathbb{E}_y[\mathbb{E}_{y'}[k(y, y')]\mathbb{E}_x[k(y, x)]] + 4\mathbb{E}_{y, y'}[k(y, y')]\mathbb{E}_{y, x}[k(y, x)], \end{aligned}$$

only the terms involving k^2 need to be calculated. Note that $k_I^2 = k_{\sqrt{2}I}$ meaning if S, R, b are replaced by $2S, 2R, \sqrt{2}b$ then the formula for (31) immediately gives a formula for the terms involving k^2 . Combining these derived formulas gives the desired expression for ξ_2 . ■

Theorem 20 is recovered by substituting $S = R, a = 0, b = m$ into Theorem 27. ■

Proof [Theorem 21] Suppose $P_n \xrightarrow{w} P$ then by Simon-Gabriel and Schölkopf (2018, Lemma 10), which holds in our case since the key intermediate result Berg et al. (1984, Theorem 3.3) only requires \mathcal{X} to be a Hausdorff space, we have $\text{MMD}(P_n, P) \rightarrow 0$.

Suppose $\text{MMD}(P_n, P) \rightarrow 0$, by Prokhorov’s theorem (Billingsley, 1971, Section 5) we know that $\{P_n\}_{n=1}^\infty$ is relatively compact. Since k is characteristic we know that \mathcal{H}_k is a separating set in the sense of Ethier and Kurtz (1986, Chapter 4) and $\text{MMD}(P_n, P) \rightarrow 0$ implies that for every $F \in \mathcal{H}_k$ that $\lim_{n \rightarrow \infty} \int F dP_n = \int F dP$ therefore Ethier and Kurtz (1986, Lemma 3.4.3) applies and we may conclude that $P_n \xrightarrow{w} P$. ■

A.5 Proof for Section 7

Proof [Proposition 22] Suppose k_0 is ISPD and C_{k_0} isn’t injective. Then there exists non-zero $x \in L^2(\mathcal{D})$ such that $C_{k_0}x = 0$ so $\int_{\mathcal{D}} \int_{\mathcal{D}} x(s)k_0(s, t)x(t)dsdt = \langle x, C_{k_0}x \rangle_{L^2(\mathcal{D})} = \langle x, 0 \rangle_{L^2(\mathcal{D})} = 0$ contradicting k_0 being ISPD. Combining Sriperumbudur et al. (2011, Proposition 5) and Sriperumbudur et al. (2010, Theorem 9) shows that if μ_{k_0} has full support then k_0 is ISPD. ■

Appendix B. Numerical Values Used in Plots

δ	ID	CEXP	SQR	FPCA	COV	FAD
0.0	0.05	0.05	0.048	0.048	0.048	0.05
0.5	0.05	0.068	0.058	0.054	0.044	0.057
1.0	0.113	0.264	0.106	0.07	0.042	0.062
1.5	0.234	0.958	0.288	0.108	0.044	0.096
2.0	0.552	1.0	0.68	0.144	0.064	0.13

Table 4: Test size, the first row, and test power, other rows, under varying values of δ in the mean shift experiment described by Pomann et al. (2016) plotted in Figure 2.

References

- Sergio Albeverio and Sonia Mazzucchi. An introduction to infinite-dimensional oscillatory and probabilistic integrals. In *Stochastic Analysis: A Series of Lectures*, pages 1–54. Springer Basel, 2015.
- Miguel A. Arcones and Evarist Giné. On the bootstrap of U and V statistics. *The Annals of Statistics*, 20(2):655–674, 1992.
- Alexander Aue, Gregory Rice, and Ozan Sönmez. Detecting and dating structural breaks in functional data without dimension reduction. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3):509–529, 2018.

δ	ID	CEXP	SQR	FPCA	COV	FAD
0	0.056	0.044	0.053	0.052	0.05	0.05
5	0.168	0.164	0.244	0.212	0.484	0.08
10	0.528	0.526	0.662	0.62	0.93	0.29
15	0.842	0.84	0.898	0.878	0.992	0.55
20	0.964	0.966	0.974	0.986	1.0	0.812

Table 5: Test size, the first row, and test power, other rows, under varying values of δ in the variance shift in one frequency experiment described by Pomann et al. (2016) plotted in Figure 4.

δ	ID	CEXP	SQR	FPCA	COV	BOOT-HS	FPCA- χ
1.0	0.064	0.064	0.062	0.062	0.05	0.025	0.006
1.2	0.108	0.11	0.152	0.098	0.11	0.05	0.016
1.4	0.194	0.184	0.336	0.178	0.264	0.168	0.011
1.6	0.354	0.326	0.552	0.314	0.47	0.338	0.01
1.8	0.544	0.522	0.754	0.506	0.632	0.556	0.036
2.0	0.716	0.69	0.856	0.642	0.764	0.666	0.07

Table 6: Test size, the first row, and test power, other rows, under varying values of δ in the scalar multiplication experiment described by Paparoditis and Sapatinas (2016) plotted in Figure 5.

Francis Bach. On the equivalence between kernel quadrature rules and random feature expansions. *The Journal of Machine Learning Research*, 18(1):714–751, 2017.

Michal Benko, Wolfgang Härdle, and Alois Kneip. Common functional principal components. *The Annals of Statistics*, 37(1):1–34, 2009.

Christian Berg, Jens Peter Reus Christensen, and Paul Ressel. *Harmonic Analysis on Semigroups*. Springer New York, 1984.

Alain Berlinet and Christine Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer US, 2004.

José R. Berrendero, Beatriz Bueno-Larraz, and Antonio Cuevas. On Mahalanobis distance in functional settings. *Journal of Machine Learning Research*, 21(9):1–33, 2020.

Patrick Billingsley. *Weak Convergence of Measures*. Society for Industrial and Applied Mathematics, 1971.

K. M. Borgwardt, A. Gretton, M. J. Rasch, H.-P. Kriegel, B. Scholkopf, and A. J. Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):49–57, 2006.

δ	ID	CEXP	SQR	FPCA	COV	FAD	CVM
0	0.046	0.052	0.064	0.052	0.072	0.05	0.05
1	0.246	0.23	0.544	0.22	0.474	0.066	0.088
2	0.912	0.862	0.994	0.828	0.968	0.28	0.2
3	0.996	0.992	1.0	0.988	0.996	0.54	0.26
4	1.0	0.998	1.0	0.994	0.1.0	0.73	0.32

Table 7: Test size, the first row, and test power, other rows, under varying values of δ in the higher order difference experiment described by Hall and Keilegom (2007) plotted in Figure 6.

Béatrice Bucchia and Martin Wendler. Change-point detection and bootstrap for Hilbert space valued random fields. *Journal of Multivariate Analysis*, 155:344–368, 2017.

Alejandra Cabana, Ana M Estrada, Jairo Pena, and Adolfo J Quiroz. Permutation tests in the two-sample problem for functional data. In *Functional Statistics and Related Fields*, pages 77–85. Springer, 2017.

C. Carmeli, E. de Vito, A. Toigo, and V. Umanità. Vector valued reproducing kernel hilbert spaces and universality. *Analysis and Applications*, 08(01):19–61, 2010.

Shubhadeep Chakraborty and Xianyang Zhang. A new framework for distance and kernel-based metrics in high dimensions. *arXiv:1909.13469*, 2019.

Huaihou Chen, Philip T Reiss, and Thaddeus Tarpey. Optimally weighted L2 distance for functional data. *Biometrics*, 70(3):516–525, 2014.

Ilya Chevyrev and Harald Oberhauser. Signature moments to characterize laws of stochastic processes. *arXiv:1810.10971*, 2018.

Andreas Christmann and Ingo Steinwart. Universal kernels on non-standard input spaces. *Advances in Neural Information Processing Systems 23*, pages 406–414, 2010.

Antonio Cuevas. A partial overview of the theory of statistics with functional data. *Journal of Statistical Planning and Inference*, 147:1–23, 2014.

G. Da Prato and J. Zabczyk. *Second Order Partial Differential Equations in Hilbert Spaces*. Cambridge University Press, 2002.

Giuseppe Da Prato. *An Introduction to Infinite-Dimensional Analysis*. Springer Berlin Heidelberg, 2006.

Hongwen Dong, Kechen Song, Yu He, Jing Xu, Yunhui Yan, and Qinggang Meng. PGA-Net: pyramid feature fusion and global context attention network for automated surface defect detection. *IEEE Transactions on Industrial Informatics*, 2019.

Stewart N. Ethier and Thomas G. Kurtz, editors. *Markov Processes*. John Wiley & Sons, Inc., 1986.

- Gregory Fasshauer and Michael McCourt. *Kernel-based Approximation Methods using MATLAB*. World Scientific, June 2014.
- Frédéric Ferraty and Philippe Vieu. Curves discrimination: a nonparametric functional approach. *Computational Statistics & Data Analysis*, 44(1-2):161–173, 2003.
- Stefan Fremdt, Josef G. Steinbach, Lajos Horváth, and Piotr Kokoszka. Testing the equality of covariance operators in functional samples. *Scandinavian Journal of Statistics*, 40(1):138–152, 2012.
- Jerome H Friedman and Lawrence C Rafsky. Multivariate generalizations of the Wald-Wolfowitz and Smirnov two-sample tests. *The Annals of Statistics*, pages 697–717, 1979.
- Thomas Gärtner. A survey of kernels for structured data. *ACM SIGKDD Explorations Newsletter*, 5(1):49, 2003.
- Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J. Smola. A kernel method for the two-sample-problem. *Advances in Neural Information Processing Systems 19*, pages 513–520, 2007.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(1):723–773, 2012a.
- Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K. Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. *Advances in Neural Information Processing Systems 25*, pages 1205–1213, 2012b.
- Peter Hall and Ingrid Van Keilegom. Two-sample tests in functional data analysis starting from discrete data. *Statistica Sinica*, 17(4):1511–1531, 2007.
- Yu He, Kechen Song, Qinggang Meng, and Yunhui Yan. An end-to-end steel surface defect detection approach via fusing multiple hierarchical features. *IEEE Transactions on Instrumentation and Measurement*, 69(4):1493–1504, 2020.
- Lajos Horváth and Piotr Kokoszka. *Inference For Functional Data With Applications*. Springer Science & Business Media, 2012.
- Lajos Horváth, Piotr Kokoszka, and Ron Reeder. Estimation of the mean of functional time series and a two-sample problem. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(1):103–122, 2012.
- Lajos Horváth, Piotr Kokoszka, and Gregory Rice. Testing stationarity of functional time series. *Journal of Econometrics*, 179(1):66–82, 2014.
- Tailen Hsing and Randall Eubank. *Theoretical Foundations of Functional Data Analysis, with an Introduction to Linear Operators*. John Wiley & Sons, Ltd, 2015.

- Hachem Kadri, Emmanuel Duflos, Philippe Preux, Stéphane Canu, Alain Rakotomamonjy, and Julien Audiffren. Operator-valued kernels for learning from functional response data. *Journal of Machine Learning Research*, 17(20):1–54, 2016.
- M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv:1807.02582*, 2018.
- Toni Karvonen, George Wynne, Filip Tronarp, Chris Oates, and Simo Särkkä. Maximum likelihood estimation and uncertainty quantification for gaussian process approximation of deterministic functions. *SIAM/ASA Journal on Uncertainty Quantification*, 8(3):926–958, 2020.
- Alexander S. Kechris. *Classical Descriptive Set Theory*. Springer New York, 1995.
- Patrick Kidger, Patric Bonnier, Imanol Perez Arribas, Cristopher Salvi, and Terry Lyons. Deep signature transforms. In *Advances in Neural Information Processing Systems 32*, pages 3105–3115. 2019.
- AN Kolmogorov-Smirnov, A Kolmogorov, and M Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *Giornale dell’Istituto Italiano degli Attuari*, 1933.
- Feng Liu, Wenkai Xu, Jie Lu, Guangquan Zhang, Arthur Gretton, and D. J. Sutherland. Learning deep kernels for non-parametric two-sample tests. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Miles Lopes, Laurent Jacob, and Martin J Wainwright. A more powerful two-sample test in high dimensions using random projection. *Advances in Neural Information Processing Systems 24*, pages 1206–1214, 2011.
- Milan N. Lukić and Jay H. Beder. Stochastic processes with sample paths in reproducing kernel Hilbert spaces. *Transactions of the American Mathematical Society*, 353(10):3945–3969, 2001.
- Stefania Maniglia and Abdelaziz Rhandi. Gaussian measures on separable Hilbert spaces and applications, 2004.
- KV Mardia and IL Dryden. The statistical analysis of shape data. *Biometrika*, 76(2): 271–281, 1989.
- H. Q. Minh. Some properties of Gaussian reproducing kernel Hilbert spaces and their implications for function approximation and learning theory. *Constructive Approximation*, 32(2):307–338, 2009.
- Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2):1–141, 2017.
- Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.

- F. J. Narcowich, J. D. Ward, and H. Wendland. Sobolev error estimates and a Bernstein inequality for scattered data interpolation via radial basis functions. *Constructive Approximation*, 24(2):175–186, 2006.
- Nicholas H Nelsen and Andrew M Stuart. The random feature model for input-output maps between Banach spaces. *arXiv:2005.10224*, 2020.
- Ivan Nourdin and Giovanni Peccati. *Normal Approximations with Malliavin Calculus*. Cambridge University Press, 2009.
- Victor M. Panaretos, David Kraus, and John H. Maddocks. Second-order comparison of Gaussian random functions and the geometry of DNA minicircles. *Journal of the American Statistical Association*, 105(490):670–682, 2010.
- E. Paparoditis and T. Sapatinas. Bootstrap-based testing of equality of mean functions or equality of covariance operators for functional data. *Biometrika*, 103(3):727–733, 2016.
- Vern I. Paulsen and Mrinal Raghupathi. *An introduction to the theory of reproducing kernel Hilbert spaces*, volume 152 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2016.
- Bruno Pelletier. Kernel density estimation on riemannian manifolds. *Statistics & Probability Letters*, 73(3):297–304, 2005.
- Gina-Maria Pomann, Ana-Maria Staicu, and Sujit Ghosh. A two-sample distribution-free test for functional data with application to a diffusion tensor imaging study of multiple sclerosis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 65(3):395–414, January 2016.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in Neural Information Processing Systems 20*, pages 1177–1184, 2008.
- Aaditya Ramdas, Sashank J Reddi, Barnabas Poczos, Aarti Singh, and Larry Wasserman. Adaptivity and computation-statistics tradeoffs for kernel and distance based high dimensional two sample testing. *arXiv preprint arXiv:1508.00655*, 2015a.
- Aaditya Ramdas, Sashank Jakkam Reddi, Barnabás Póczos, Aarti Singh, and Larry Wasserman. On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 3571–3577, 2015b.
- C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Sashank Reddi, Aaditya Ramdas, Barnabas Poczos, Aarti Singh, and Larry Wasserman. On the High Dimensional Power of a Linear-Time Two Sample Test under Mean-shift Alternatives. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pages 772–780, 2015.

- Saburoou Saitoh and Yoshihiro Sawano. *Theory of reproducing kernels and applications*. Springer, 2016.
- Paul Schmid. On the Kolmogorov and Smirnov limit theorems for discontinuous distribution functions. *The Annals of Mathematical Statistics*, 29(4):1011–1027, 1958.
- I. J. Schoenberg. Metric spaces and completely monotone functions. *The Annals of Mathematics*, 39(4):811, 1938.
- Dino Sejdinovic, Bharath Sriperumbudur, Arthur Gretton, and Kenji Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The Annals of Statistics*, 41(5):2263–2291, 2013.
- Robert J. Serfling. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, Inc., November 1980.
- John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- Carl-Johann Simon-Gabriel and Bernhard Schölkopf. Kernel distribution embeddings: Universal kernels, characteristic kernels and kernel metrics on distributions. *Journal of Machine Learning Research*, 19(44):1–29, 2018.
- Carl-Johann Simon-Gabriel, Alessandro Barp, and Lester Mackey. Metrizing weak convergence with Maximum Mean Discrepancies. *arXiv:2006.09268*, 2020.
- Nickolay Smirnov. Table for estimating the goodness of fit of empirical distributions. *The Annals of Mathematical Statistics*, 19(2):279–281, 1948.
- Kechen Song and Yunhui Yan. A noise robust method based on completed local binary patterns for hot-rolled steel strip surface defects. *Applied Surface Science*, 285:858–864, 2013.
- Kechen Song and Yunhui Yan. NEU Steel Dataset Description. http://faculty.neu.edu.cn/yunhyan/NEU_surface_defect_database.html, 2020. Last accessed 29/07/2020.
- Bharath K. Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert R.G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010. ISSN 1532-4435.
- Bharath K. Sriperumbudur, Kenji Fukumizu, and Gert R. G. Lanckriet. Universality, characteristic kernels and RKHS embedding of measures. *Journal of Machine Learning Research*, 12:2389–2410, 2011. ISSN 1532-4435.
- I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, 2008.
- Ingo Steinwart. Convergence types and rates in generic karhunen-loève expansions with applications to sample path properties. *Potential Analysis*, 51(3):361–395, 2018.

- Ingo Steinwart and Clint Scovel. Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constructive Approximation*, 35(3):363–417, 2012.
- T.J. Sullivan. *Introduction to Uncertainty Quantification*. Springer International Publishing, 2015.
- Dougal J. Sutherland. Unbiased estimators for the variance of MMD estimators. *arXiv:1906.02104*, 2019.
- Dougal J. Sutherland, Hsiao-Yu Tung, Heiko Strathmann, Soumyajit De, Aaditya Ramdas, Alex Smola, and Arthur Gretton. Generative models and model criticism via optimized maximum mean discrepancy. In *Proceedings of the 8th International Conference on Learning Representations*, 2016.
- Gábor J Székely. E-statistics: The energy of statistical samples. *Bowling Green State University, Department of Mathematics and Statistics Technical Report*, 3(05):1–18, 2003.
- Gábor J Székely and Maria L Rizzo. Testing for equal distributions in high dimension. *InterStat*, 5(16.10):1249–1272, 2004.
- N. N. Vakhania, V. I. Tarieladze, and S. A. Chobanyan. *Probability Distributions on Banach Spaces*. Springer Netherlands, 1987.
- A. W. van der Vaart and J. H. van Zanten. Information rates of nonparametric Gaussian process methods. *Journal of Machine Learning Research*, 12:2095–2119, 2011.
- A Wald and J Wolfowitz. On a test whether two samples are from the same distribution. *The Annals of Mathematical Statistics*, 11:147–162, 1940.
- H. Wendland. *Scattered Data Approximation*. Cambridge University Press, 2005.
- William J. Wilkinson. *Gaussian Process Modelling for Audio Signals*. PhD thesis, Queen Mary Univeristy London, 2019.
- George Wynne, François-Xavier Briol, and Mark Girolami. Convergence guarantees for gaussian process means with misspecified likelihoods and smoothness. *Journal of Machine Learning Research*, 22(123):1–40, 2021.
- Haizhang Zhang, Yuesheng Xu, and Qinghui Zhang. Refinement of operator-valued reproducing kernels. *Journal of Machine Learning Research*, 13(4):91–136, 2012.
- Jin-Ting Zhang, Xuehua Liang, and Shengning Xiao. On the two-sample Behrens-Fisher problem for functional data. *Journal of Statistical Theory and Practice*, 4(4):571–587, 2010.
- Changbo Zhu, Shun Yao, Xianyang Zhang, and Xiaofeng Shao. Distance-based and RKHS-based dependence metrics in high dimension. *Annals of Statistics*, 2019. To appear.