

Learning sets with separating kernels

Ernesto De Vito^{a,*}, Lorenzo Rosasco^{b,c,d}, Alessandro Toigo^{e,f}

^a DIMA, Università di Genova, Genova, Italy

^b DIBRIS, Università di Genova & Istituto Italiano di Tecnologia, Italy

^c Massachusetts Institute of Technology, USA

^d Istituto Italiano di Tecnologia, Italy

^e Dipartimento di Matematica, Politecnico di Milano, Milano, Italy

^f I.N.F.N., Sezione di Milano, Milano, Italy

Received 11 April 2012

Received in revised form 31 October 2013

Accepted 7 November 2013

Available online 15 November 2013

Communicated by Ding-Xuan Zhou

1. Introduction

In this paper we study the problem of learning from data the set where the data probability distribution is concentrated. Our study is more broadly motivated by questions in unsupervised learning, such as the problem of inferring geometric properties of probability distributions from random samples.

In recent years, there has been great progress in the theory and algorithms for supervised learning, i.e. function approximation problems from random noisy data [10,22,29,55,74]. On the other hand, while there are a number of methods and studies in unsupervised learning, e.g. algorithms for clustering, dimensionality reduction, dictionary learning (see Chapter 14 of [38]), many interesting problems remain largely unexplored.

Our analysis starts with the observation that many studies in unsupervised learning hinge on at least one of the following two assumptions. The first is that the data are distributed according to a probability distribution which is absolutely continuous with respect to a reference measure, such as the Lebesgue measure. In this case it is possible to define a density and the corresponding density level sets. Studies in this scenario include [8,30,44,69] to name a few. Such an assumption prevents considering the case where the data are represented in a high-dimensional Euclidean space but are concentrated on a Lebesgue negligible subset, as a lower-dimensional submanifold. This motivates the second assumption – sometimes called *manifold assumption* – postulating that the data lie on a low-dimensional Riemannian manifold embedded in a Euclidean space. This latter idea has triggered a large number of different algorithmic and theoretical studies (see for example [4,6,20,21,27,59]). Though the manifold assumption has proved useful in some applications, there are many practical scenarios where it might not be satisfied. This observation has motivated considering more general situations such as *manifold plus noise* models [18,52], and models where the data are described by combinations of more than one manifold [46,76].

* Corresponding author.

E-mail addresses: devito@dim.unige.it (E. De Vito), lrosasco@mit.edu (L. Rosasco), alessandro.toigo@polimi.it (A. Toigo).

Here we consider a different point of view and work in a setting where the data are described by an abstract probability space and a *similarity function* induced by a reproducing kernel [65]. In this framework, we consider the basic problem of estimating the set where the data distribution is concentrated (see Section 1.2 for a detailed discussion of related works). A special class of reproducing kernels, that we call separating kernels, plays a special role in our study. First, it allows to define a suitable metric on the probability space and makes the support of the distribution well defined; second, it leads to a new analytical characterization of the support in terms of the null space of the integral operator associated to the reproducing kernel.

This last result is the key towards a new computational approach to learn the support from data, since the integral operator can be approximated with high probability from random samples [58,65]. Estimation of the null space of the integral operator can be unstable, and regularization techniques can be used to obtain stable estimators. In this paper we study a class of regularization techniques proposed to solve ill-posed problems [34] and already studied in the context of supervised learning [3,48]. Regularization is achieved by *filtering* out the small eigenvalues of the sample empirical matrix defined by the kernel. Different algorithms are defined by different filter functions and have different computational properties. Consistency and stability properties for a large class of spectral filters and of the corresponding algorithms are established in a unified framework. Numerical experiments show that the proposed algorithms are competitive, and often better, than other state of the art techniques.

The paper is divided into two parts. The first part includes Section 2, where we establish several mathematical results relating reproducing kernel Hilbert spaces of functions on a set X and the geometry of the set X itself. In particular, in this section we introduce the concept of separating kernel, which we further explore in Section 3. These results are of interest in their own right, and are at the heart of our approach. In the second part of the paper we discuss the problem of learning the support from data. More precisely, in Section 4 we illustrate some algorithms for learning the support of a distribution from random samples. In Section 5 we establish universal consistency for the proposed methods and discuss stability to random sampling. We conclude in Sections 6 and 7 with some further discussions and some numerical experiments, respectively. A conference version of this paper appeared in [28]. We now start by describing in some more detail our results and discussing some related works.

1.1. Summary of main results

In this section we briefly describe the main ideas and results in the paper.

The setting we consider is described by a probability space (X, ρ) and a measurable reproducing kernel K on the set X [2]. The data are independent and identically distributed (i.i.d.) samples x_1, \dots, x_n , each one drawn from X with probability ρ . The reproducing kernel K reflects some prior information on the problem and, as we discuss in the following, will also define the geometry of X . The goal is to use the sample points x_1, \dots, x_n to estimate the region where the probability measure ρ is concentrated.

To fix some ideas, the space X can be thought of as a high-dimensional Euclidean space and the distribution ρ as being concentrated on a region X_ρ , which is a smaller – and potentially lower dimensional – subset of X (e.g. a linear subspace or a manifold). In this example, the goal is to build from data an estimator X_n which is, with high probability, close to X_ρ with respect to a suitable metric.

We first note that a precise definition of X_ρ requires some care. If ρ is assumed to have a continuous density with respect to some fixed reference measure (for example, the Lebesgue measure in the Euclidean space), then the region X_ρ can be easily defined to be the closure of the set of points where the density function is non-zero. Nevertheless, this assumption would prevent considering the situation where the data are concentrated on a “small”, possibly lower dimensional, subset of X . Note that, if the set X were endowed with a topological structure and ρ were defined on the corresponding Borel σ -algebra, it would be natural to define X_ρ as the support of the measure ρ , i.e. the smallest *closed* subset of X having measure one. However, since the set X is only assumed to be a measurable space, no a priori given topology is available. Here we also remark that the definition of X_ρ is not the only point where some further structure on X would be useful. Indeed, when defining a learning error, a notion of distance between the set X_ρ and its estimator X_n is also needed and hence some metric structure on X is required.

The idea is to use the properties of the reproducing kernel K to induce a metric structure – and consequently a topology – on X . Indeed, under some mild technical assumptions on K , the function

$$d_K(x, y) = \sqrt{K(x, x) + K(y, y) - 2K(x, y)} \quad \forall x, y \in X$$

defines a metric on X , thus making X a topological space. Then, it is natural to define X_ρ to be the support of ρ with respect to such metric topology. Moreover, the Hausdorff distance d_H induced by the metric d_K provides a notion of distance between closed sets.

The problem we consider can now be restated as follows: we want to learn from data an estimator X_n of X_ρ , such that $\lim_{n \rightarrow \infty} d_H(X_n, X_\rho) = 0$ almost surely. While X_ρ is now well defined, it is not clear how to build an estimator from data. A main result in the paper, given in Theorem 3, provides a new analytic characterization of X_ρ , which immediately suggests a new computational solution for the corresponding learning problem. To derive and state this result, we introduce a new notion of reproducing kernels, called separating kernels, that, roughly speaking, captures the sense in which the reproducing




Fig. 1. The separating property is illustrated in a simple situation where $X = \mathbb{R}^2$. In the top pictures, the support X_ρ is a line passing through the origin and is separated by the linear kernel $K(x, y) = x^T y$: for all $x \notin X_\rho$, there exists a function $f \in \mathcal{H}$ (a linear function on X) which is zero on X_ρ and such that $f(x) \neq 0$. The pictures on the right are a plot of the plane $y = f(x_1, x_2)$. In the bottom pictures, the support is a segment passing through the origin. The linear kernel is too simple to separate this set: all planes are going to be zero also outside of the support (the dotted line in the picture).

kernel and the probability distribution need to be related. We say that a reproducing kernel Hilbert space \mathcal{H} (or equivalently its kernel) *separates* a subset $C \subset X$, if, for any $x \notin C$, there exists $f \in \mathcal{H}$ such that

$$f(x) \neq 0 \quad \text{and} \quad f(y) = 0 \quad \forall y \in C.$$

If K separates all possible closed subsets in X , we say that it is *completely separating*. Fig. 1 illustrates the notion of separating kernel in the simple example of the linear kernel in a Euclidean space.

Now, [Theorem 3](#) states that, if either K is completely separating, or at least separates X_ρ , then X_ρ is the level set of a suitable distribution dependent continuous function F_ρ . More precisely, let \mathcal{H} be the reproducing kernel Hilbert space associated to K [2], $T : \mathcal{H} \rightarrow \mathcal{H}$ the integral operator with kernel K , and denote by T^\dagger its pseudo-inverse. If we consider the function F_ρ on X , defined by

$$F_\rho(x) = \langle T^\dagger T K_x, K_x \rangle \quad \forall x \in X,$$

and K separates X_ρ , then we prove that

$$X_\rho = \{x \in X \mid F_\rho(x) = 1\},$$

(where for simplicity we are assuming $K(x, x) = 1$ for all $x \in X$).

The above result is crucial since the integral operator T can be approximated with high probability from data (see [58] and references therein). However, since the definition of F_ρ involves the pseudo-inverse of T , the support estimation problem can be unstable [71] and regularization techniques are needed to ensure stability. With this in mind, we propose and study a family of spectral regularization techniques which are classical in inverse problems [34] and have been considered in supervised learning in [3,48]. We define an estimator by

$$X_n = \{x \in X \mid F_n(x) \geq 1 - \tau_n\},$$

where $F_n(x) = (1/n) \mathbf{K}_x^\top g_{\lambda_n}(\mathbf{K}_n/n) \mathbf{K}_x$, with $(\mathbf{K}_n)_{i,j} = K(x_i, x_j)$, \mathbf{K}_x is the column vector whose i -th entry is $K(x_i, x)$, and \mathbf{K}_x^\top is its transpose. Here $g_{\lambda_n}(\mathbf{K}_n/n)$ is a matrix defined via spectral calculus by a spectral filter function g_{λ_n} that suppresses the contribution of the eigenvalues smaller than λ_n . Examples of spectral filters include Tikhonov regularization and truncated singular values decomposition [48], to name a few.

This class of methods can be studied within a unified framework, and the error analysis in the paper establishes strong universal consistency if X_ρ is separated by K . More precisely, under the latter assumption, we show in [Theorem 6](#) that,

$$\lim_{n \rightarrow \infty} d_H(X_n, X_\rho) = 0 \quad \text{almost surely,}$$

provided that X is compact and the sequences $(\tau_n)_{n \geq 1}, (\lambda_n)_{n \geq 1}$ are chosen so that,

$$\tau_n = 1 - \min_{1 \leq i \leq n} F_n(x_i), \quad \lim_{n \rightarrow \infty} \lambda_n = 0, \quad \sup_{n \geq 1} (L_{\lambda_n} \log n) / \sqrt{n} < +\infty,$$

where L_{λ_n} is the Lipschitz constant of the function $r_{\lambda_n}(\sigma) = \sigma g_{\lambda_n}(\sigma)$. The above result is universal in the sense that consistency can be shown without assuming regularity condition on ρ or X_ρ .

The proof of the above result crucially depends on estimating the deviation between F_n and F_ρ . Indeed, for the above choice of the sequence $(\lambda_n)_{n \geq 1}$ we show that

$$\lim_{n \rightarrow \infty} \sup_{x \in X} |F_\rho(x) - F_n(x)| = 0 \quad \text{almost surely.}$$

Under suitable distribution dependent assumptions, the above result can be further developed to obtain finite sample bounds quantifying stability to random sampling. Indeed, if the couple (ρ, K) is such that $\sup_{x \in X} \|T^{-s/2} T^\dagger T K_x\| < +\infty$, with $0 < s \leq 1$, and the eigenvalues of the (compact and positive) operator T satisfy $\sigma_j \sim j^{-1/b}$ for some $0 < b \leq 1$, then we prove in [Theorem 7](#) that, for $n \geq 1$ and $\delta > 0$, we have

$$\sup_{x \in X} |F_n(x) - F_\rho(x)| \leq C_{s,b,\delta} \left(\frac{1}{n} \right)^{\frac{s}{2s+b+1}}$$

with probability at least $1 - 2e^{-\delta}$, for $\lambda_n = n^{-1/(2s+b+1)}$ and a suitable constant $C_{s,b,\delta}$ which does not depend on n .

Finally, we remark that our construction relies on the assumption that the kernel K separates the support X_ρ . The question then arises whether there exist kernels that can separate a large number of, and perhaps *all*, closed subsets, namely kernels that are *completely separating*. Indeed, a positive answer can be given and, for translation invariant kernels on \mathbb{R}^d , [Theorem 4](#) actually gives a sufficient condition for a kernel to be completely separating in terms of its Fourier transform. As a consequence, the Abel kernel $K(x, y) = e^{-\|x-y\|/\sigma}$ on the Euclidean space $X = \mathbb{R}^d$ is completely separating. Interestingly, the Gaussian kernel $K(x, y) = e^{-\|x-y\|^2/\sigma^2}$, which is very popular in machine learning, is not.

1.2. State of the art

The problem of building an estimator X_n of a subset $X_\rho \subset X$ which is consistent with respect to some kind of metric among sets has been considered in seemingly diverse fields for different application purposes, from anomaly detection – see [\[17\]](#) for a review – to surface estimation [\[60\]](#). We give a summary of the main approaches, with basic references for further details.

Support and level set estimation. Support estimation (also called set estimation) is a part of the theory of non-parametric statistics. We refer to [\[24,25\]](#) for a detailed review on this topic. Usually, the space X is \mathbb{R}^d with the Euclidean metric d , and X_ρ is the corresponding support of ρ . If X_ρ is convex, a natural estimator is the convex hull of the data $X_n = \text{conv}\{x_1, \dots, x_n\}$, for which convergence rates can be derived with respect to the Hausdorff distance [\[33,56\]](#). If X_ρ is not convex, Devroye and Wise [\[30\]](#) propose the estimator

$$X_n = \bigcup_{i=1}^n B(x_i, \epsilon_n),$$

where $B(x, \epsilon)$ is the ball of center x and radius ϵ , and ϵ_n slowly goes to zero when n tends to infinity. Consistency and minimax converges rates are studied in [\[30,44\]](#) with respect to the distance

$$d_\mu(C_1, C_2) = \mu(C_1 \Delta C_2),$$

where $C_1 \Delta C_2 = (C_1 \setminus C_2) \cup (C_2 \setminus C_1)$ and μ is a suitable known measure.

If ρ has a density f with respect to some known measure μ , a traditional approach is based on a non-parametric estimator f_n of f , a so-called *plug-in* estimator. A kernel based class of plug-in estimators is proposed in [\[23\]](#), namely

$$X_n = \left\{ x \in X \mid f_n(x) \geq c_n \right\} \quad \text{with} \quad f_n(x) = \frac{1}{nh_n^d} \sum_{i=1}^n K\left(\frac{x-x_i}{h_n}\right),$$

where h_n is a regularization parameter and c_n is a suitable threshold. Convergence rates with respect to d_μ are provided in [\[23\]](#).

A related problem is level set estimation, where the goal is to detect the high density regions $\{x \in X \mid f(x) \geq c\}$. Consistency and optimal convergence rates for different plug-in estimators

$$X_n = \left\{ x \in X \mid f_n(x) \geq c \right\}$$

have been studied with respect to both d_H and d_μ , see for example [\[8,63,72\]](#) for a slightly different approach.

One class learning algorithm. In machine learning, set estimation has been viewed as a classification problem where we have at our disposal only positive examples. An interesting discussion on the relation between density level set estimation, binary classification and anomaly detection is given in [\[69\]](#). In this context, some algorithms inspired by Support Vector

Machine (SVM) have been studied in [61,69,75]. A kernel method based on kernel principal component analysis is presented in [39] and is essentially a special case of our framework.

Manifold learning. As we mentioned before, a setting which is of special interest is the one in which X is \mathbb{R}^d and X_ρ is a low-dimensional Riemannian submanifold. In this case, the error of an estimator is studied in terms of the error functional

$$d_\rho(X_\rho, X_n) = \int_{X_\rho} d(x, X_n) d\rho(x),$$

where d is the Euclidean metric. Some results in this framework are given in [1,49,51].

Computational geometry. A classic situation, considered for example in image reconstruction problems, is when the set X_ρ is a hyper-surface of \mathbb{R}^d and the data x_1, \dots, x_n are either chosen deterministically or sampled uniformly. The goal in this case is to find a smooth function f that gives the Cartesian equation of the hyper-surface, see for example [40,42,47].

2. Kernels, integral operators and geometry in probability spaces

In this section we establish the results that provide the foundations of our approach. The basic framework in this paper is described by a triple (X, ρ, K) , where

- X is a set (endowed with a σ -algebra \mathcal{A}_X);
- ρ is a probability measure defined on X ;
- K is a (real) reproducing kernel on X , i.e. a real function on $X \times X$ of positive type.

We interpret X as the data space and ρ as the probability distribution generating the data. Roughly speaking, the kernel K provides a natural *similarity measure* on X and it defines its geometry.

We denote by \mathcal{H} the reproducing kernel Hilbert space associated with the reproducing kernel K (we refer to [2,68] for an exhaustive review on the theory of reproducing kernel Hilbert spaces). The scalar product and norm in \mathcal{H} are denoted by $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$, respectively. We recall that the elements of \mathcal{H} are real functions on X , and the reproducing property $f(x) = \langle f, K_x \rangle$ holds true for all $x \in X$ and $f \in \mathcal{H}$, where $K_x \in \mathcal{H}$ is defined by $K_x(y) = K(y, x)$.

In order to prove our results, we need some technical conditions on K .

Assumption 1. The kernel K has the following properties:

- a) for all $x, y \in X$ with $x \neq y$ we have $K_x \neq K_y$;
- b) the associated reproducing kernel Hilbert space \mathcal{H} is separable;
- c) the real function K is measurable with respect to the product σ -algebra $\mathcal{A}_X \otimes \mathcal{A}_X$;
- d) for all $x \in X$, $K(x, x) = 1$.

Assumptions 1.a), 1.b) and 1.c) are minimal requirements. In particular, Assumptions 1.a) and 1.b) are needed in order to define a separable metric structure on X , while Assumption 1.c) ensures that such metric topology is compatible with the σ -algebra \mathcal{A}_X (see Proposition 1 below). In Proposition 2, the combination of 1.a), 1.b) and 1.c) will allow us to define the support X_ρ of the probability measure ρ , as anticipated in Section 1.1. Assumption 1.d), instead, is a normalization requirement, and could be replaced by a suitable boundedness condition (in fact, even weaker integrability conditions could also be considered). We choose the normalization $K(x, x) = 1 \forall x \in X$ since it makes equations more readable, and it is not restrictive in view of Proposition 13 in Appendix A.1.

We now show how the above assumptions allow us to define a metric on X and to characterize the corresponding support of ρ in terms of the integral operator with kernel K .

2.1. Metric induced by a kernel

Our first result makes X a separable metric space isometrically embedded in \mathcal{H} . This point of view is developed in [65]. The relation between metric spaces isometrically embedded in Hilbert spaces and kernels of positive type was studied by Schoenberg around 1940. A recent discussion on this topic can be found in Chapter 2 §3 of [7].

Proposition 1. Under Assumption 1.a), the map $d_K : X \times X \rightarrow [0, +\infty[$ defined by

$$d_K(x, y) = \|K_x - K_y\| = \sqrt{K(x, x) + K(y, y) - 2K(x, y)} \quad (1)$$

is a metric on X . Furthermore

- i) the map $x \mapsto K_x$ is an isometry from X into \mathcal{H} ;
- ii) the kernel K is a continuous function on $X \times X$, and each $f \in \mathcal{H}$ is a continuous function.

If also [Assumption 1.b](#)) is satisfied, then

iii) the metric space (X, d_K) is separable.

Finally, if also [Assumption 1.c](#)) holds true, then

iv) the closed subsets of X are measurable (with respect to \mathcal{A}_X);

v) if Y is a topological space endowed with its Borel σ -algebra and $f : X \rightarrow Y$ is continuous, then f is measurable; in particular, the functions in \mathcal{H} are measurable.

Proof. Many of these properties are known in the literature, see for example [\[15,68\]](#) and references therein. For the reader's convenience, we give a self-contained short proof.

[Assumption 1.a](#)) states that the map $x \mapsto K_x$ is injective. Since $d_K(x, y) = \|K_x - K_y\|$ by definition, d_K is the metric on X making $x \mapsto K_x$ an isometry, as claimed in item i). About ii), the kernel K is continuous since $K(x, y) = \langle K_y, K_x \rangle$ and the map $x \mapsto K_x$ is continuous by item i); furthermore, the elements of \mathcal{H} are continuous by the reproducing property $f(x) = \langle f, K_x \rangle$.

If also [Assumption 1.b](#)) holds true, then the set $\{K_x \mid x \in X\}$ is separable in \mathcal{H} , and so is X as the map $x \mapsto K_x$ is isometric from X onto $\{K_x \mid x \in X\}$. Item iii) then follows.

Suppose now that also [Assumption 1.c](#)) holds true. Then the map d_K is a measurable map, so that the open balls of X are measurable. Since X is separable, any open set is a countable union of open balls, hence it is measurable. It follows that the closed subsets are measurable, too, hence item iv).

Let Y and f be as in item v). If $A \subset Y$ is closed, then $f^{-1}(A)$ is closed in X , hence measurable by item iv). It follows that $f^{-1}(A)$ is measurable for all Borel sets $A \subset Y$, i.e. f is measurable. Since the elements of \mathcal{H} are continuous by ii), they are measurable, and item v) is proved. \square

In the rest of the paper we will always consider X as a topological metric space with metric d_K . Note that d_K is the metric induced on X by the norm of \mathcal{H} through the embedding $x \mapsto K_x$. The next result shows that under our assumptions we can define the set X_ρ as the smallest closed subset of X having measure one.

Proposition 2. Under [Assumptions 1.a](#)), [1.b](#)) and [1.c](#)), there exists a unique closed subset $X_\rho \subset X$ with $\rho(X_\rho) = 1$ satisfying the following property: if C is a closed subset of X and $\rho(C) = 1$, then $C \supset X_\rho$.

Proof. Define the measurable set X_ρ as

$$X_\rho = \bigcap_{\substack{C \text{ closed} \\ \rho(C)=1}} C.$$

Clearly, X_ρ is closed and measurable by [Proposition 1](#). Since X is separable, there exists a sequence of closed subsets $(C_j)_{j \geq 1}$ such that every closed subset $C = \bigcap C_{j_k}$, for some suitable subsequence. Hence, $X_\rho = \bigcap_{j \mid \rho(C_j)=1} C_j$ and, as a consequence, $\rho(X_\rho) = 1$. \square

We add one remark. The set X_ρ is called *the support* of the measure ρ and clearly depends both on the probability distribution and on the topology induced by the kernel K through the metric d_K on X .

2.2. Separating kernels

The following definition of separating kernel plays a central role in our approach.

Definition 1. We say that the reproducing kernel Hilbert space \mathcal{H} separates a subset $C \subset X$, if, for all $x \notin C$, there exists $f \in \mathcal{H}$ such that

$$f(x) \neq 0 \quad \text{and} \quad f(y) = 0 \quad \forall y \in C. \tag{2}$$

In this case we also say that the corresponding reproducing kernel separates C .

We add some comments. First, in [\(2\)](#) the function f depends on x and C . Second, the reproducing property and [\(2\)](#) imply that $K_x \neq 0$ and $K_x \neq K_y$ for all $x \notin C$ and $y \in C$ (compare with [Assumption 1.a](#))). Finally, we stress that a different notion of *separating* property is given in [\[68\]](#).

Remark 1. Given an arbitrary reproducing kernel Hilbert space \mathcal{H} , there exist sets that are not separated by \mathcal{H} . For example, if $X = \mathbb{R}^d$ and \mathcal{H} is the reproducing kernel Hilbert space with linear kernel $K(x, y) = x^T y$, the only sets separated by \mathcal{H} are the linear manifolds, that is, the set of points defined by homogeneous linear equations (see Fig. 1). A natural question is then whether there exist kernels capable of separating large classes of subsets and in particular all the closed subsets. Section 3 answers positively to this question, introducing the notion of completely separating kernels.

Next, we provide an equivalent characterization of the separating property, which will be the key to a computational approach to support estimation. For any set C , let $P_C : \mathcal{H} \rightarrow \mathcal{H}$ be the orthogonal projection onto the closed subspace

$$\mathcal{H}_C = \overline{\text{span}\{K_x \mid x \in C\}},$$

i.e. the closure of the linear space generated by the family $\{K_x \mid x \in C\}$. Note that $P_C^2 = P_C$, $P_C^\top = P_C$ and

$$\ker P_C = \{K_x \mid x \in C\}^\perp = \{f \in \mathcal{H} \mid f(x) = 0 \forall x \in C\}.$$

Moreover, define the function

$$F_C : X \rightarrow \mathbb{R}, \quad F_C(x) = \langle P_C K_x, K_x \rangle. \quad (3)$$

Remark 2. The Hilbert space \mathcal{H}_C is a closed subspace of the reproducing kernel Hilbert space \mathcal{H} , and it is itself a reproducing kernel Hilbert space of functions on X with reproducing kernel $K_C(x, y) = \langle P_C K_y, P_C K_x \rangle = \langle P_C K_y, K_x \rangle$. Note that $K_C(x, y) = K(x, y)$ for all $x, y \in C$ by definition of P_C . Clearly, the function F_C corresponds to the value of K_C on the diagonal.

Then, we have the following theorem.

Theorem 1. For any subset $C \subset X$, the following facts are equivalent:

- i) \mathcal{H} separates the set C ;
- ii) for all $x \notin C$, $K_x \notin \text{ran } P_C$;
- iii) $C = \{x \in X \mid F_C(x) = K(x, x)\}$;
- iv) $\{K_x \mid x \in C\} = \{K_x \mid x \in X\} \cap \text{ran } P_C$.

Under Assumption 1.a), if C is separated by \mathcal{H} , then C is closed with respect to the metric d_K .

Proof. We first prove that i) \Rightarrow ii). Given $x \notin C$, by assumption there is $f \in \mathcal{H}$ such that $\langle f, K_x \rangle = f(x) \neq 0$, i.e. $K_x \notin \{f\}^\perp$, and $\langle f, K_y \rangle = f(y) = 0$ for all $y \in C$, i.e. $f \in \ker P_C = \text{ran } P_C^\perp$. It follows that $\text{ran } P_C \subset \{f\}^\perp$, and then $K_x \notin \text{ran } P_C$.

We prove ii) \Rightarrow iii). If $x \in C$, then $K_x \in \text{ran } P_C$ by definition of P_C , so that $F_C(x) = K(x, x)$. Hence $C \subset \{x \in X \mid F_C(x) = K(x, x)\}$. If $x \notin C$, then by assumption $P_C K_x \neq K_x$, i.e. $(I - P_C)K_x \neq 0$. By the equality

$$\|(I - P_C)K_x\|^2 = \langle K_x, K_x \rangle - \langle P_C K_x, K_x \rangle - \langle K_x, P_C K_x \rangle + \langle P_C K_x, P_C K_x \rangle = K(x, x) - F_C(x),$$

this implies $F_C(x) \neq K(x, x)$. Hence $C \supset \{x \in X \mid F_C(x) = K(x, x)\}$.

We prove iii) \Rightarrow i). If $x \notin C$, define $f = (I - P_C)K_x \in \ker P_C$, so that $f(y) = 0$ for all $y \in C$. Furthermore, $f(x) = K(x, x) - F_C(x) \neq 0$. Thus, f separates the set C .

Finally, iv) is a restatement of ii) taking into account that $K_x \in \text{ran } P_C$ for all $x \in C$ by construction.

Under Assumption 1.a), the map $x \mapsto F_C(x) - K(x, x) = \langle P_C K_x, K_x \rangle - K(x, x)$ is continuous by Proposition 1. By item iii), C is the 0-level set of this function, hence C is closed. \square

Proposition 13 in Appendix A.1 shows that the reproducing kernel K can be normalized under the mild assumption that $K(x, x) \neq 0$ for all $x \in X$, so that Assumption 1.d) can be satisfied up to a rescaling of K .

2.2.1. A special case: metric spaces

It may be the case that the set X has its own metric d_X , and the σ -algebra \mathcal{A}_X is the Borel σ -algebra associated with the topology induced by d_X . The following proposition shows that the metrics d_K and d_X induce the same topology on X , provided that \mathcal{H} separates all the d_X -closed subsets and the corresponding kernel is continuous.

Proposition 3. Let X be a separable metric space with respect to a metric d_X , and \mathcal{A}_X the corresponding Borel σ -algebra. Let \mathcal{H} be a reproducing kernel Hilbert space on X with kernel K . Assume that the kernel K is a continuous function with respect to d_X and that the space \mathcal{H} separates every subset of X which is closed with respect to d_X . Then

- i) Assumptions 1.a), 1.b) and 1.c) hold true, and $K(x, x) > 0$ for all $x \in X$;
- ii) a set is closed with respect to d_K if and only if it is closed with respect to d_X .

Proof. The kernel is measurable and the space \mathcal{H} is separable by Proposition 5.1 and Corollary 5.2 in [15]. Since the points are closed sets for d_X and the d_X -closed sets are separated by \mathcal{H} , then $K_x \neq 0$ (i.e. $K(x, x) > 0$) for all $x \in X$ and $K_x \neq K_y$ if $x \neq y$ by the discussion following Definition 1.

We show that d_X and d_K are equivalent metrics. Take a sequence $(x_j)_{j \geq 1}$ such that for some $x \in X$ it holds that $\lim_{j \rightarrow \infty} d_X(x_j, x) = 0$. Since K is continuous with respect to d_X , we have $\lim_{j \rightarrow \infty} d_K(x_j, x) = 0$. Hence, the d_K -closed sets are d_X -closed, too. Conversely, if the set C is d_X -closed, since \mathcal{H} separates C , Theorem 1 implies that $C = \{x \in X \mid K(x, x) - F_C(x) = 0\}$, which is a d_K -closed set by d_K -continuity of the map $x \mapsto K(x, x) - F_C(x)$. \square

Item ii) of the above proposition states that the metrics d_K and d_X are equivalent and implies that the set X_ρ defined in Proposition 2 coincides with the support of ρ with respect to the topology induced by d_X .

2.3. The integral operator defined by the kernel

We denote by \mathcal{S}_1 the Banach space of the trace class operators on \mathcal{H} , with trace class norm

$$\|A\|_{\mathcal{S}_1} = \text{tr}[(A^\top A)^{\frac{1}{2}}] = \sum_{i \in I} \langle (A^\top A)^{\frac{1}{2}} e_i, e_i \rangle,$$

where $\{e_i\}_{i \in I}$ is any orthonormal basis of \mathcal{H} . Furthermore, we let \mathcal{S}_2 be the separable Hilbert space of the Hilbert–Schmidt operators on \mathcal{H} , with Hilbert–Schmidt norm

$$\|A\|_{\mathcal{S}_2}^2 = \text{tr}[A^\top A] = \sum_{i \in I} \|Ae_i\|^2.$$

Finally, if A is any bounded operator on \mathcal{H} , we denote by $\|A\|_\infty$ its uniform operator norm. It is standard that $\|A\|_\infty \leq \|A\|_{\mathcal{S}_2} \leq \|A\|_{\mathcal{S}_1}$. Moreover, for all functions $f_1, f_2 \in \mathcal{H}$, the rank-one operator $f_1 \otimes f_2$ on \mathcal{H} defined by

$$(f_1 \otimes f_2)(f) = \langle f, f_2 \rangle f_1 \quad \forall f \in \mathcal{H}$$

is trace class, and $\|f_1 \otimes f_2\|_{\mathcal{S}_1} = \|f_1 \otimes f_2\|_{\mathcal{S}_2} = \|f_1\| \|f_2\|$.

We recall a few facts on integral operators with kernel K (see [15] for proofs and further discussions). Under Assumption 1, the \mathcal{S}_1 -valued map $x \mapsto K_x \otimes K_x$ is Bochner-integrable with respect to ρ , and its integral

$$T = \int_X K_x \otimes K_x d\rho(x) \tag{4}$$

defines a positive trace class operator T with $\|T\|_{\mathcal{S}_1} = \text{tr}[T] = 1$ (a short proof is given in Proposition 14 of Appendix A.2). Using the reproducing property of \mathcal{H} , it is straightforward to see that T is simply the integral operator with kernel K acting on \mathcal{H} , i.e.

$$(Tf)(x) = \int_X K(x, y) f(y) d\rho(y) \quad \forall f \in \mathcal{H}.$$

The following is a key result in our approach.

Theorem 2. Under Assumption 1, the null space of T is

$$\ker T = \{K_x \mid x \in X_\rho\}^\perp = \ker P_{X_\rho}, \tag{5}$$

where X_ρ is the support of ρ as defined in Proposition 2.

Proof. Note that, for all $f \in \mathcal{H}$, the set

$$C_f = \{x \in X \mid f(x) = 0\} = \{x \in X \mid \langle f, K_x \rangle = 0\}$$

is closed since f is continuous. We now prove Eq. (5). Since T is a positive operator, spectral theorem implies that $Tf = 0$ if and only if $\langle Tf, f \rangle = 0$. The definition of T and the reproducing property give that

$$\langle Tf, f \rangle = \int_X \langle (K_x \otimes K_x) f, f \rangle d\rho(x) = \int_X |\langle K_x, f \rangle|^2 d\rho(x) = \int_X |f(x)|^2 d\rho(x),$$

hence the condition $\langle Tf, f \rangle = 0$ is equivalent to the fact that $f(x) = 0$ for ρ -almost every $x \in X$. Hence $f \in \ker T$ if and only if $\rho(C_f) = 1$, i.e. $C_f \supset X_\rho$, or equivalently $\langle f, K_x \rangle = 0 \forall x \in X_\rho$. Eq. (5) then follows. \square

In the following, we will use the abbreviated notation $P_\rho = P_{X_\rho}$. Note that the space \mathcal{H} splits into the direct sum $\mathcal{H} = \mathcal{H}_\rho \oplus \mathcal{H}_\rho^\perp$, where

$$\begin{aligned}\mathcal{H}_\rho &= \text{ran } P_\rho = \overline{\text{ran } T} = \overline{\text{span}\{K_x \mid x \in X_\rho\}}, \\ \mathcal{H}_\rho^\perp &= \ker P_\rho = \ker T = \{f \in \mathcal{H} \mid f(x) = 0 \ \forall x \in X_\rho\}.\end{aligned}$$

Remark 3. The reproducing kernel Hilbert space \mathcal{H}_ρ (see Remark 2) has been considered before [70], and in particular in the context of semi-supervised manifold regularization [5], where X_ρ is assumed to be an embedded manifold. The corresponding reproducing kernel is $K_\rho(x, y) = \langle P_\rho K_y, K_x \rangle$ and $F_{X_\rho}(x) = K_\rho(x, x)$. See also the discussion in Section 6.

Under Assumption 1, we also introduce the integral operator $L_K : L^2(X, \rho) \rightarrow L^2(X, \rho)$,

$$(L_K \phi)(x) = \int_X K(x, y) \phi(y) d\rho(y) \quad \forall \phi \in L^2(X, \rho),$$

which is a positive trace class operator, too. Note the difference between the operators T and L_K : although their definitions are formally the same, the respective domains and images change.

Since T and L_K are positive trace class operators, by the Hilbert–Schmidt theorem each of them admits an orthonormal family of eigenvectors in \mathcal{H} and $L^2(X, \rho)$, respectively, with a corresponding family of positive eigenvalues. The two spectral decompositions are strongly related, as we now briefly recall (see also Proposition 8 of [58] and Theorem 2.11 of [70]).

Denote by $(\sigma_j)_{j \in J}$ the (finite or countable) family of strictly positive eigenvalues of L_K , where each eigenvalue is repeated according to its (finite) multiplicity. For each $j \in J$ select a corresponding eigenvector $\phi_j \in L^2(X, \rho)$ in such a way that the sequence $(\phi_j)_{j \in J}$ is orthonormal in $L^2(X, \rho)$. Hilbert–Schmidt theorem provides that

$$L_K = \sum_{j \in J} \sigma_j \phi_j \otimes \phi_j, \tag{6}$$

where the series converges in trace norm. In general, each element ϕ_j is an equivalence class of functions defined ρ -almost everywhere. In particular, the value of ϕ_j is not defined outside X_ρ . However, in each equivalence class we can choose a unique continuous function, denoted again by ϕ_j , which is defined at every point of X by means of the *extension equation* [19,58]

$$\phi_j(x) = \sigma_j^{-1} \int_X K(x, y) \phi_j(y) d\rho(y) \quad \forall x \in X. \tag{7}$$

With this choice, which will be implicitly assumed in the following, the family $(\sigma_j)_{j \in J}$ coincides with the family of strictly positive eigenvalues of T (with the same multiplicities), $(\sqrt{\sigma_j} \phi_j)_{j \in J}$ is an orthonormal family in \mathcal{H} of eigenfunctions of T , and

$$T = \sum_{j \in J} \sigma_j (\sqrt{\sigma_j} \phi_j) \otimes (\sqrt{\sigma_j} \phi_j) = \sum_{j \in J} \sigma_j^2 \phi_j \otimes \phi_j, \tag{8}$$

where the series converges in the Banach space \mathcal{S}_1 (hence in \mathcal{S}_2), see e.g. [15,58,70]. As $\|T\|_{\mathcal{S}_1} = 1$, the positive sequence $(\sigma_j)_{j \in J}$ is summable and sums up to 1. It is clear that the family $(\sqrt{\sigma_j} \phi_j)_{j \in J}$ is an orthonormal basis of the Hilbert space \mathcal{H}_ρ . Conversely, let $(f_j)_{j \in J}$ be an orthonormal basis of \mathcal{H}_ρ of eigenvectors of T with corresponding eigenvalues $(\sigma_j)_{j \in J}$. Define

$$\phi_j(x) = \sigma_j^{-\frac{1}{2}} f_j(x) \quad \forall x \in X.$$

Then, it is not difficult to show that (6), (7) and (8) hold true.

2.4. An analytic characterization of the support

Let Assumption 1 hold true. Collecting the previous results, if \mathcal{H} separates X_ρ , then Theorem 1 gives that

$$X_\rho = \{x \in X \mid F_{X_\rho}(x) = 1\}.$$

The function $F_\rho = F_{X_\rho}$ is defined by (3) in terms of the projection P_ρ , which, in light of Theorem 2, can be characterized using the operator T . Indeed, from the definition of F_ρ and (5) we have

$$F_\rho(x) = \langle P_\rho K_x, K_x \rangle = \langle T^\dagger T K_x, K_x \rangle = \langle \theta(T) K_x, K_x \rangle = \sum_{j \in J} \sigma_j |\phi_j(x)|^2 \tag{9}$$

where T^\dagger is the pseudo-inverse of T and θ is the Heaviside function $\theta(\sigma) = \mathbb{1}_{]0, +\infty[}(\sigma)$ (note that with our definition $\theta(0) = 0$). The above discussion is summarized in the following theorem.

Theorem 3. *If \mathcal{H} satisfies [Assumption 1](#) and separates the support X_ρ of the measure ρ , then*

$$X_\rho = \{x \in X \mid F_\rho(x) = 1\} = \{x \in X \mid \langle T^\dagger T K_x, K_x \rangle = 1\}.$$

As we discussed before, a natural question is whether there exist kernels capable to separate *all* possible closed subsets of X . In a learning scenario, this can be translated into a *universality* property, in the sense that it allows to describe *any* probability distribution and learn consistently its support [29]. Note that in a supervised learning framework a similar role is played by the so-called universal kernels [16,67]. The following section answers positively to the previous question, introducing and studying the concept of completely separating kernels. Interestingly, there are universal kernels in the sense of [16,67] which do not separate all closed subsets of X , as for example the Gaussian kernel.

3. Completely separating reproducing kernel Hilbert spaces

The property defining the class of kernels we are interested in is captured by the following definition.

Definition 2 (*Completely separating kernel*). A reproducing kernel Hilbert space \mathcal{H} satisfying [Assumption 1.a](#)) is called *completely separating* if \mathcal{H} separates all the subsets $C \subset X$ which are closed with respect to the metric d_K defined by (1). In this case, we also say that the corresponding reproducing kernel is completely separating.

The definition of completely separating reproducing kernel Hilbert spaces should be compared with the analogous notion of complete regularity for topological spaces. Indeed, we recall that a topological space is called *completely regular* if, for any closed subset C and any point $x \notin C$, there exists a continuous function f such that $f(x) \neq 0$ and $f(y) = 0$ for all $y \in C$. As we discuss below, completely separating reproducing kernels do exist. For example, for $X = \mathbb{R}^d$ both the Abel kernel $K(x, y) = e^{-\|x-y\|/\sigma}$ and the ℓ_1 -exponential kernel $K(x, y) = e^{-\|x-y\|_1/\sigma}$ are completely separating, where $\|x\|$ is just the Euclidean norm of $x = (x^1, \dots, x^d)$ in \mathbb{R}^d and $\|x\|_1 = \sum_{j=1}^d |x_j|$ is the ℓ_1 -norm. Indeed this follows from [Theorem 4](#) and [Proposition 6](#) below, which give sufficient conditions for a kernel to be completely separating in the case $X = \mathbb{R}^d$. Note that the Gaussian kernel $K(x, y) = e^{-\|x-y\|^2/\sigma^2}$ on \mathbb{R}^d is not completely separating. This is a consequence of the following fact. It is known that the elements of the corresponding reproducing kernel Hilbert space \mathcal{H} are analytic functions, see Corollary 4.44 in [68]. If C is a closed subset of \mathbb{R}^d with non-empty interior and $f \in \mathcal{H}$ is equal to zero on C , then a standard result in complex analysis implies that $f(x) = 0$ for all $x \in \mathbb{R}^d$, hence \mathcal{H} does not separate C .

We end this section with [Proposition 6](#), which gives a simple way to build completely separating kernels in high-dimensional spaces from completely separating kernels in one dimension, the latter usually being easier to characterize.

3.1. Separating properties of translation invariant kernels

The first result studies translation invariant kernels on \mathbb{R}^d , i.e. of the form $K(x, y) = K(x - y)$. We show that if the Fourier transform of the kernel satisfies a suitable growth condition, then the corresponding reproducing kernel Hilbert space is completely separating. As usual, $C(\mathbb{R}^d)$ denotes the space of real continuous functions on \mathbb{R}^d and, for any $p \in [1, +\infty[$, $L^p(\mathbb{R}^d)$ is the space of (equivalence classes of) real functions on \mathbb{R}^d which are p -integrable with respect to the Lebesgue measure dx . We will consider the *real* spaces $L^p_h(\mathbb{R}^d)$ of hermitian complex functions, i.e.

$$L^p_h(\mathbb{R}^d) = \{\phi_1 + i\phi_2 \mid \phi_1, \phi_2 \in L^p(\mathbb{R}^d) \text{ and } \phi_1(-x) = \phi_1(x), \phi_2(-x) = -\phi_2(x)\}.$$

If $\phi \in L^1(\mathbb{R}^d)$, its Fourier transform is the complex hermitian bounded continuous function $\hat{\phi}$ on \mathbb{R}^d given by

$$\hat{\phi}(z) = \int_{\mathbb{R}^d} e^{-2\pi iz \cdot x} \phi(x) dx.$$

If $\phi \in L^2(\mathbb{R}^d)$, we denote by $\hat{\phi} \in L^2_h(\mathbb{R}^d)$ its Fourier-Plancherel transform, obtained extending the above definition on functions $\phi \in L^1(\mathbb{R}^d) \cap L^2(\mathbb{R}^d)$ to a unitary map $L^2(\mathbb{R}^d) \ni \phi \rightarrow \hat{\phi} \in L^2_h(\mathbb{R}^d)$.

Throughout, we assume \mathbb{R}^d to be a metric space with respect to the standard metric $d_{\mathbb{R}^d}$ induced by the Euclidean norm.

We need a preliminary result characterizing a reproducing kernel Hilbert space, whose reproducing kernel is continuous and integrable, as a suitable non-closed subspace of $L^2(\mathbb{R}^d)$. The first part is a converse of Bochner's theorem (Theorem 4.18 in [35]).

Proposition 4. *Let K be a continuous function in $L^1(\mathbb{R}^d)$ such that its Fourier transform \hat{K} is strictly positive. Then the kernel $K(x, y) = K(x - y)$ is positive definite and its corresponding (real) reproducing kernel Hilbert space \mathcal{H} is*

$$\mathcal{H} = \left\{ \phi \in C(\mathbb{R}^d) \cap L^2(\mathbb{R}^d) \mid \int_{\mathbb{R}^d} \hat{K}(z)^{-1} |\hat{\phi}(z)|^2 dz < +\infty \right\} \quad (10)$$

with norm

$$\|\phi\|^2 = \int_{\mathbb{R}^d} \hat{K}(z)^{-1} |\hat{\phi}(z)|^2 dz \quad \forall \phi \in \mathcal{H}. \quad (11)$$

Proof. The integral operator

$$(L_K \phi)(x) = \int_{\mathbb{R}^d} K(x-y)\phi(y) dy = (K * \phi)(x),$$

is well defined and bounded from $L^2(\mathbb{R}^d)$ into $L^2(\mathbb{R}^d)$ since $K \in L^1(\mathbb{R}^d)$. Since L_K is a convolution operator, Fourier transform turns it into the operator of multiplication by the bounded function \hat{K} , that is $\widehat{L_K \phi} = \hat{K} \hat{\phi}$ for all $\phi \in L^2(\mathbb{R}^d)$. It follows that

$$\langle L_K \phi, \phi \rangle_{L^2} = \langle \hat{K} \hat{\phi}, \hat{\phi} \rangle_{L^2} > 0 \quad \forall \phi \in L^2(\mathbb{R}^d) \setminus \{0\}$$

since $\hat{K} > 0$ by assumption, hence L_K is a strictly positive operator. In order to show that K is positive definite, pick a Dirac sequence $(\varphi_n)_{n \geq 1}$ as in Chapter VIII.3 of [45], and, for each $x \in X$, define φ_n^x be equal to $\varphi_n^x(y) = \varphi_n(y-x)$. Fixed $x_1, x_2, \dots, x_N \in \mathbb{R}^d$ and $c_1, c_2, \dots, c_N \in \mathbb{R}$, set $\phi_n = \sum_{i=1}^N c_i \varphi_n^{x_i}$, then

$$0 \leq \langle L_K \phi_n, \phi_n \rangle_{L^2} = \sum_{i,j=1}^N c_i c_j \langle L_K \varphi_n^{x_i}, \varphi_n^{x_j} \rangle_{L^2} \xrightarrow{n \rightarrow \infty} \sum_{i,j=1}^N c_i c_j K(x_j, x_i),$$

where the last equality is due to continuity of K and the usual properties of Dirac sequences. It follows that $\sum_{i,j=1}^N c_i c_j K(x_j, x_i) \geq 0$, i.e. the kernel K is positive definite.

Let \mathcal{H} be the (real) reproducing kernel Hilbert space associated to K . Since the support of the Lebesgue measure is \mathbb{R}^d , Mercer theorem (as stated e.g. in Proposition 6.1 of [15] and the subsequent discussion, or Theorem 2.11 of [70]) shows that $L_K^{1/2}$ is a unitary isomorphism from $L^2(\mathbb{R}^d)$ onto \mathcal{H} . More precisely, for any $\psi \in L^2(\mathbb{R}^d)$ there exists a unique function $\phi \in C(\mathbb{R}^d)$ such that its equivalence class belongs to $L_K^{1/2} \psi \in L^2(\mathbb{R}^d)$, and the correspondence $\psi \mapsto \phi$ is an isometry from $L^2(\mathbb{R}^d)$ onto \mathcal{H} . By further applying the Fourier-Plancherel transform and taking into account that $\hat{\phi}(z) = \sqrt{\hat{K}(z)} \hat{\psi}(z)$ for almost all $z \in \mathbb{R}^d$, one has

$$\|\phi\|_{\mathcal{H}}^2 = \|\psi\|_{L^2(\mathbb{R}^d)}^2 = \|\hat{\psi}\|_{L^2(\mathbb{R}^d)}^2 = \int_{\mathbb{R}^d} \hat{K}(z)^{-1} |\hat{\phi}(z)|^2 dz < +\infty,$$

so that (10) and (11) follow. \square

We now state a sufficient condition on K ensuring that \mathcal{H} is completely separating.

Theorem 4. Let K be a continuous function in $L^1(\mathbb{R}^d)$ such that

$$\hat{K}(z) \geq \frac{a}{(1+b\|z\|^{\gamma_1})^{\gamma_2}} \quad \forall z \in \mathbb{R}^d \quad (12)$$

for some $a, b, \gamma_1, \gamma_2 > 0$. Then,

- i) the translation invariant kernel $K(x, y) = K(x-y)$ is positive definite and continuous;
- ii) the topologies induced by the metric d_K and the Euclidean metric $d_{\mathbb{R}^d}$ coincide on \mathbb{R}^d ;
- iii) the kernel K is completely separating.

Proof. Condition (12) implies that \hat{K} is strictly positive, so item i) follows from Proposition 4. In particular, from (10) we see that, if $\phi \in L^2(\mathbb{R}^d)$ and $\int_{\mathbb{R}^d} (1+b\|z\|^{\gamma_1})^{\gamma_2} |\hat{\phi}(z)|^2 dz$ is finite, then $\phi \in \mathcal{H}$. This implies that $C_c^\infty(\mathbb{R}^d) \subset \mathcal{H}$: indeed, if $\phi \in C_c^\infty(\mathbb{R}^d)$, then $\hat{\phi}$ is a Schwartz function on \mathbb{R}^d (Theorem 3.2 in [66]), hence the last integral is convergent. Functions in $C_c^\infty(\mathbb{R}^d)$ separate every set C which is closed with respect to the metric $d_{\mathbb{R}^d}$ (as it easily follows by suitably translating and dilating the function $\psi \in C_c^\infty(\mathbb{R}^d)$ defined in item (b) p. 19 of [66]), hence \mathcal{H} separates the $d_{\mathbb{R}^d}$ -closed subsets. Items ii) and iii) then follow from Proposition 3. \square

As an application, we show that the Abel kernel is completely separating.

Proposition 5. *Let*

$$K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}, \quad K(x, y) = e^{-\frac{\|x-y\|}{\sigma}}, \quad (13)$$

with $\sigma > 0$. Then K is a positive definite kernel and the corresponding reproducing kernel Hilbert space \mathcal{H} is completely separating for all $d \geq 1$.

Proof. A standard Fourier transform computation gives

$$\hat{K}(z) = \frac{1}{2\pi\sigma} \pi^{-\frac{d+1}{2}} \Gamma\left(\frac{d+1}{2}\right) \left(\frac{1}{4\pi^2\sigma^2} + \|z\|^2\right)^{-\frac{d+1}{2}}, \quad (14)$$

where Γ is Euler gamma function (Theorem 1.14 in [66]). The claim then follows from [Theorem 4](#). \square

Eqs. (10), (11) and (14) show that (up to a rescaling of the norm) the reproducing kernel Hilbert space associated to the Abel kernel (13) is just $W^{(d+1)/2}(\mathbb{R}^d)$, the Sobolev space of order $(d+1)/2$.

3.2. Building separating kernels

The following result gives a way to construct completely separating reproducing kernel Hilbert spaces on high-dimensional spaces.

Proposition 6. *If X_i , $i = 1, 2, \dots, d$, are sets and $K^{(i)}$ are completely separating reproducing kernels on X_i for all $i = 1, 2, \dots, d$, then the product kernel*

$$K((x_1, \dots, x_d), (y_1, \dots, y_d)) = K^{(1)}(x_1, y_1) \cdots K^{(d)}(x_d, y_d)$$

is completely separating on the set $X = X_1 \times X_2 \times \cdots \times X_d$.

Proof. Each set X_i and X are endowed with the metric $d_{K^{(i)}}$ and d_K induced by the corresponding kernels, and \mathcal{H}_i and \mathcal{H} denote the reproducing kernel Hilbert spaces with kernels $K^{(i)}$ and K , respectively. A standard result gives that $\mathcal{H} = \mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_d$ and $K_x = K_{x_1}^{(1)} \otimes \cdots \otimes K_{x_d}^{(d)}$ for all $x = (x_1, \dots, x_d) \in X$ [2]. We claim that the d_K -topology on X is contained in the product topology of the $d_{K^{(i)}}$ -topologies on X_i (actually, it is not difficult to show that the two topologies coincide). Indeed, if $(x_{i,k})_{k \geq 1}$ are sequences in X_i such that $\lim_{k \rightarrow \infty} d_{K^{(i)}}(x_{i,k}, x_i) = 0$ for all $i = 1, \dots, d$, then

$$\begin{aligned} \lim_{k \rightarrow \infty} d_K((x_{1,k}, \dots, x_{d,k}), (x_1, \dots, x_d))^2 &= \lim_{k \rightarrow \infty} \|K_{(x_{1,k}, \dots, x_{d,k})} - K_{(x_1, \dots, x_d)}\|^2 \\ &= \lim_{k \rightarrow \infty} [K^{(1)}(x_{1,k}, x_{1,k}) \cdots K^{(d)}(x_{d,k}, x_{d,k}) - 2K^{(1)}(x_{1,k}, x_1) \cdots K^{(d)}(x_{d,k}, x_d) \\ &\quad + K^{(1)}(x_1, x_1) \cdots K^{(d)}(x_d, x_d)] \\ &= 0, \end{aligned}$$

since $\lim_{k \rightarrow \infty} K^{(i)}(x_{i,k}, x_{i,k}) = \lim_{k \rightarrow \infty} K^{(i)}(x_{i,k}, x_i) = K^{(i)}(x_i, x_i)$. We now prove that \mathcal{H} is completely separating. If $C \subset X$ is d_K -closed and $x = (x_1, \dots, x_d) \in X \setminus C$, since C is also closed in the product topology, for all $i = 1, \dots, d$ there exists an open neighborhood U_i of x_i in X_i such that $U = U_1 \times \cdots \times U_d \subset X \setminus C$. Since each \mathcal{H}_i is completely separating, for all $i = 1, \dots, d$ there exists $f_i \in \mathcal{H}_i$ such that $f_i(x_i) \neq 0$ and $f_i(y_i) = 0$ for all $y_i \in X_i \setminus U_i$. Then the product function $f = f_1 \otimes \cdots \otimes f_d$ is in \mathcal{H} , and satisfies $f(x) \neq 0$ and $f(y) = 0$ for all $y \in C$. \square

As a consequence, the Abel kernel defined by the ℓ_1 -norm

$$K(x, y) = e^{-\frac{\|x-y\|_1}{\sigma}} = \prod_{i=1}^d e^{-\frac{|x_i-y_i|}{\sigma}}, \quad x = (x_1, \dots, x_d), \quad y = (y_1, \dots, y_d)$$

is completely separating since each kernel in the product is positive definite and completely separating by [Proposition 5](#).

4. A spectral approach to learning the support

In this section we study the set estimation problem in the context of learning theory. We fix a triple (X, ρ, K) as in [Section 2](#), and assume throughout that the reproducing kernel K satisfies [Assumption 1](#). We regard X as a metric space with respect to d_K , and continue to denote by X_ρ the support of ρ defined in [Proposition 2](#).

If \mathcal{H} separates X_ρ , [Theorem 3](#) shows that the support X_ρ is the 1-level set of a suitable function F_ρ defined by the integral operator T , and therefore depending on K and ρ . However, the probability distribution ρ is unknown, as we only have a set of i.i.d. points x_1, \dots, x_n sampled from ρ at our disposal. Our task is now to use our sample in order to estimate the set X_ρ .

The definition of T given by [\(4\)](#) suggests that it can be estimated by the data dependent operator

$$T_n = \frac{1}{n} \sum_{i=1}^n K_{x_i} \otimes K_{x_i}. \quad (15)$$

The operator T_n is positive and with finite rank; in particular, $T_n \in \mathcal{S}_1$ and $\|T_n\|_{\mathcal{S}_1} = \text{tr}[T_n] = 1$. We denote by $(\sigma_j^{(n)})_{j \in J_n}$ the strictly positive eigenvalues of T_n (each one repeated according to its multiplicity) and by $(\sqrt{\sigma_j^{(n)}} \phi_j^{(n)})_{j \in J_n}$ the corresponding eigenvectors; note that in the present case the index set J_n is finite. However, though T_n converges to T in all relevant topologies (see [Lemma 1](#) and [Remark 6](#) below), in general $T_n^\dagger T_n$ does not converge to $T^\dagger T$ since T^\dagger may be unbounded, or, equivalently, since 0 may be an accumulation point of the spectrum of T when $\dim \mathcal{H} = \infty$. Hence, the problem of support estimation is ill-posed, and regularization techniques are needed to restore well-posedness and ensure a stable solution. In the following sections, we will show that spectral regularization [\[3,34,48\]](#) can be used to learn the support efficiently from the data.

4.1. Regularized estimators via spectral filtering

An approach which is classical in inverse problems (see [\[34\]](#), and also [\[3,48\]](#) for applications to learning) consists in replacing the pseudo-inverses T_n^\dagger and T^\dagger with some bounded approximations obtained by *filtering out* the components corresponding to the eigenvalues of T_n and T which are smaller than a fixed regularization parameter λ . This is achieved by introducing a suitable *filter function* $g_\lambda : [0, +\infty[\rightarrow [0, +\infty[$ and replacing T_n^\dagger , T with the bounded operators $g_\lambda(T_n)$, $g_\lambda(T)$ defined by spectral calculus. If the function g_λ is sufficiently regular, then convergence of T_n to T implies convergence of $g_\lambda(T_n)$ to $g_\lambda(T)$ in the Hilbert–Schmidt norm. On the other hand, if the regularization parameter λ goes to zero, then $g_\lambda(T)$ converges to T^\dagger in an appropriate sense. We are now going to apply the same idea to our setting. Since we are interested in approximating the orthogonal projection $P_\rho = T^\dagger T = \theta(T)$ rather than the pseudo-inverse T^\dagger , we introduce a low-pass filter r_λ , in a way that the bounded operator $r_\lambda(T)$ is an approximation of $\theta(T)$. In terms of the previously defined function g_λ , this can be achieved by setting $r_\lambda(\sigma) = g_\lambda(\sigma)\sigma$ for all $\sigma \in \mathbb{R}$, so that $r_\lambda(T) = g_\lambda(T)T$. Explicitly, in terms of the spectral decompositions of T_n and T we have

$$r_\lambda(T_n) = \sum_{j \in J_n} r_\lambda(\sigma_j^{(n)}) \left(\sqrt{\sigma_j^{(n)}} \phi_j^{(n)} \right) \otimes \left(\sqrt{\sigma_j^{(n)}} \phi_j^{(n)} \right), \quad r_\lambda(T) = \sum_{j \in J} r_\lambda(\sigma_j) \left(\sqrt{\sigma_j} \phi_j \right) \otimes \left(\sqrt{\sigma_j} \phi_j \right).$$

Note that, since the spectra of T_n and T are both contained in the interval $[0, 1]$, we can assume that the functions g_λ and r_λ are defined on $[0, 1]$. Moreover, as the operators $r_\lambda(T_n)$ and $r_\lambda(T)$ approximate orthogonal projections, it is useful to have the bound $0 \leq r_\lambda(T_n)$, $r_\lambda(T) \leq I$ satisfied for all T_n and T 's, and this can be achieved by choosing the function r_λ such that $0 \leq r_\lambda(\sigma) \leq 1$ for all σ .

As a consequence of the above discussion, the characterization of filter functions giving rise to stable algorithms is captured by the following assumption.

Assumption 2. The family of functions $(r_\lambda)_{\lambda > 0}$, with $r_\lambda : [0, 1] \rightarrow [0, 1]$ for all $\lambda > 0$, has the following properties:

- a) $r_\lambda(0) = 0$ for all $\lambda > 0$;
- b) for all $\sigma > 0$, we have $\lim_{\lambda \rightarrow 0^+} r_\lambda(\sigma) = 1$;
- c) for all $\lambda > 0$, there exists a positive constant L_λ such that

$$|r_\lambda(\sigma) - r_\lambda(\tau)| \leq L_\lambda |\sigma - \tau| \quad \forall \sigma, \tau \in [0, 1].$$

By [Assumption 2.a](#), there exists a function $g_\lambda : [0, 1] \rightarrow [0, +\infty[$ such that $r_\lambda(\sigma) = g_\lambda(\sigma)\sigma$. On the other hand, by [Assumption 2.b](#) we have $\lim_{\lambda \rightarrow 0^+} r_\lambda(\sigma) = \theta(\sigma)$ for all $\sigma \in [0, 1]$. [Assumption 2.c](#) is of technical nature, and will become clear in [Section 5.2](#); here we note that in particular it implies that r_λ is a continuous function for all $\lambda > 0$.

A few examples of filter functions r_λ satisfying [Assumption 2](#) and of corresponding functions g_λ are given in [Table 1](#). It is easy to check that for each of them $L_\lambda = 1/\lambda$. See [\[34\]](#) for further examples.

For a chosen filter, the corresponding regularized empirical estimator of F_ρ is defined by

$$F_n(x) = \langle r_{\lambda_n}(T_n) K_x, K_x \rangle = \sum_{j \in J_n} r_{\lambda_n}(\sigma_j^{(n)}) \sigma_j^{(n)} |\phi_j^{(n)}(x)|^2 \quad (16)$$

Table 1

Examples of filter functions satisfying [Assumption 2](#). For Landweber filter m_λ is an integer such that $\lim_{\lambda \rightarrow 0} m_\lambda = \infty$.

Tikhonov regularization	$r_\lambda(\sigma) = \frac{\sigma}{\sigma + \lambda}$	$g_\lambda(\sigma) = \frac{1}{\sigma + \lambda}$
Spectral cut-off	$r_\lambda(\sigma) = \mathbb{1}_{[\lambda, +\infty)}(\sigma) + \frac{\sigma}{\lambda} \mathbb{1}_{[0, \lambda)}(\sigma)$	$g_\lambda(\sigma) = \frac{1}{\sigma} \mathbb{1}_{[\lambda, +\infty)}(\sigma) + \frac{1}{\lambda} \mathbb{1}_{[0, \lambda)}(\sigma)$
Landweber filter	$r_\lambda(\sigma) = \sigma \sum_{k=0}^{m_\lambda} (1 - \sigma)^k$	$g_\lambda(\sigma) = \sum_{k=0}^{m_\lambda} (1 - \sigma)^k$

where we allow the regularization parameter λ_n to depend on the number of samples n . Note that the functions F_n and F_ρ are continuous on X by continuity of the mapping $x \mapsto K_x$ (see i) of [Proposition 1](#)). In [Section 5](#) we will show that, for an appropriate choice of the sequence $(\lambda_n)_{n \geq 1}$, the estimator F_n converges almost surely to F_ρ uniformly on compact subsets of X . Unfortunately, this does not imply convergence of the 1-level sets of F_n to the 1-level set of F_ρ in any sense (as, for example, with respect to the Hausdorff distance). However, an estimator of X_ρ can be obtained by setting

$$X_n = \{x \in X \mid F_n(x) \geq 1 - \tau_n\}, \quad (17)$$

where $\tau_n > 0$ is an off-set parameter that depends on the sample size n (recall that F_n takes values in $[0, 1]$). In [Section 5](#) we show that, for a suitable choice of the sequence $(\tau_n)_{n \geq 1}$, the closed set X_n is indeed a consistent estimator of the support with respect to the Hausdorff distance.

In the following section we discuss some remarks about the computation of F_n .

4.2. Algorithmic and computational aspects

We show that the computation of F_n (hence of X_n) reduces to a finite-dimensional problem involving the empirical kernel matrix defined by the data. To this purpose, it is useful to introduce the sampling operator

$$S_n : \mathcal{H} \rightarrow \mathbb{R}^n \quad S_n f = \begin{pmatrix} f(x_1) \\ \vdots \\ f(x_n) \end{pmatrix}, \quad (18)$$

which can be interpreted as the restriction operator which evaluates functions in \mathcal{H} on the points of the training set. The transpose of S_n is

$$S_n^\top : \mathbb{R}^n \rightarrow \mathcal{H} \quad S_n^\top \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{pmatrix} = \sum_{i=1}^n \alpha_i K_{x_i},$$

and S_n^\top can be interpreted as the out-of-sample extension operator [\[19,58\]](#). A simple computation shows that

$$T_n = \frac{1}{n} S_n^\top S_n \quad S_n S_n^\top = \mathbf{K}_n \quad (\mathbf{K}_n)_{ij} = K(x_i, x_j).$$

Hence, considering the filter given in the form $r_\lambda(T_n) = g_\lambda(T_n)T_n$, we have

$$r_\lambda(T_n) = g_\lambda\left(\frac{S_n^\top S_n}{n}\right) \frac{S_n^\top S_n}{n} = \frac{1}{n} S_n^\top g_\lambda\left(\frac{S_n S_n^\top}{n}\right) S_n = \frac{1}{n} S_n^\top g_\lambda\left(\frac{\mathbf{K}_n}{n}\right) S_n,$$

where the second equality follows from spectral calculus. Using the definition of the sampling operator, we can consider the n -dimensional vector \mathbf{K}_x defined by

$$\mathbf{K}_x = S_n K_x = \begin{pmatrix} K(x_1, x) \\ \vdots \\ K(x_n, x) \end{pmatrix},$$

and [\(16\)](#) can be written as

$$F_n(x) = \langle r_{\lambda_n}(T_n) K_x, K_x \rangle = \left\langle \frac{1}{n} g_\lambda\left(\frac{\mathbf{K}_n}{n}\right) S_n K_x, S_n K_x \right\rangle = \frac{1}{n} \mathbf{K}_x^* g_{\lambda_n}\left(\frac{\mathbf{K}_n}{n}\right) \mathbf{K}_x, \quad (19)$$

where \mathbf{K}_x^* is the conjugate transpose of \mathbf{K}_x . More explicitly we have

$$F_n(x) = \sum_{i=1}^n \alpha_i(x) K(x, x_i) \quad \alpha_i(x) = \frac{1}{n} \sum_{j=1}^n \left(g_{\lambda_n}\left(\frac{\mathbf{K}_n}{n}\right) \right)_{ij} K(x_j, x). \quad (20)$$

The above equation shows that, while \mathcal{H} could be infinite dimensional, the computation of the estimator reduces to a finite-dimensional problem. Further, though the mathematical definition of the filter is done through spectral calculus, the computations might not require performing an eigen-decomposition. As an example, for Tikhonov regularization $g_{\lambda_n}(\sigma) = \frac{1}{\sigma + \lambda_n}$, so that $g_{\lambda_n}(\frac{\mathbf{K}_n}{n}) = (\frac{\mathbf{K}_n}{n} + \lambda_n)^{-1}$ and the coefficient vector $\alpha(x)$ in (20) is given by

$$\alpha(x) = (\mathbf{K}_n + n\lambda_n)^{-1}\mathbf{K}_x.$$

In the case of the Landweber filter, it is possible to prove that the coefficient vector can be evaluated iteratively by setting $\alpha^0(x) = 0$, and

$$\alpha^t(x) = \alpha^{t-1}(x) + \frac{1}{n}(\mathbf{K}_x - \mathbf{K}_n\alpha^{t-1}(x))$$

for $t = 1, \dots, m_{\lambda_n}$. We refer to [48] for the corresponding algorithm in a supervised framework; see also the discussion in Section 6.3.

We thus see that the estimator corresponding to Tikhonov regularization can be computed via Cholesky decomposition and has complexity of order $O(n^3)$. For Landweber iteration the complexity is $O(n^2m)$, where m is the number of iterations. Finally, the spectral cut-off, or truncated SVD, requires $O(n^3)$ operations to compute the eigen-decomposition of the kernel matrix. Further discussions can be found in [48] and references therein. We end remarking that, in order to test whether N points belong or not to the support, we simply have to repeat the above computation replacing \mathbf{K}_x by an $n \times N$ matrix $\mathbf{K}_{x,N}$, in which each column is a vector \mathbf{K}_x corresponding to a point x in the test set. Note that in this case the coefficients $\alpha(x)$ will also form an $n \times N$ matrix.

5. Error analysis: convergence and stability

In this section we develop an error analysis for the proposed class of estimators. First, we discuss convergence (consistency) and then stability with respect to random sampling in terms of finite sample bounds. We continue to suppose throughout this section that Assumption 1 holds true, and consider X as a metric space with metric d_K .

5.1. Empirical data

We recall that the empirical data are a set of i.i.d. points x_1, \dots, x_n , each one drawn from X with probability ρ . Since we need to study asymptotic properties when the sample size n goes to infinity, we introduce the following probability space

$$\Omega = \{(x_i)_{i \geq 1} \mid x_i \in X \forall i \geq 1\}, \quad (21)$$

endowed with the product σ -algebra $\mathcal{A}_\Omega = \mathcal{A}_X \otimes \mathcal{A}_X \otimes \dots$ and the product probability measure $\mathbb{P} = \rho \otimes \rho \otimes \dots$. We recall that, given an integer n and a topological space M endowed with the σ -algebra of its Borel subsets, an M -valued estimator of size n is a measurable map $\mathcal{E}_n : \Omega \rightarrow M$ depending only on the first n -variables, that is

$$\mathcal{E}_n(\omega) = \xi_n(x_1, \dots, x_n) \quad \omega = (x_i)_{i \geq 1}$$

for some measurable map $\xi_n : X^n \rightarrow M$. The number n is the cardinality of the sampled data. We then have the following facts.

Proposition 7. For all $n \geq 1$

- i) T_n is a S_k -valued estimator for $k = 1, 2$;
- ii) if X is locally compact, then F_n is a $C(X)$ -valued estimator, where $C(X)$ is the space of continuous functions on X with the topology of uniform convergence on compact subsets.

The proof of the above proposition is rather technical, and we defer the interested reader to Appendix A.2 for more details.

Remark 4. In item ii) of Proposition 7, the assumption that X is locally compact is needed to ensure that the topology of uniform convergence on compact subsets is a separable metric topology on $C(X)$, which in turn is essential to prove measurability of the random variable F_n (see the proof of Proposition 16 in Appendix A.2). In many examples, the set X has its own locally compact separable metric d_X . In this case, in order for X to be locally compact metric space also for the metric d_K , it is enough that the kernel K is a d_X -continuous function separating every subset of X which is closed with respect to d_X , as the two topologies induced by d_X and d_K then coincide by item ii) of Proposition 3.

If X is not locally compact (which we will regard as a pathological case), then, in order to have measurability of F_n , one needs to replace the probability measure \mathbb{P} with the outer measure (see the discussion in Section 2 of [43] and in Section 1.7 of [73]).

Remark 5. Statisticians adopt a different notation: the data are described by a family Y_1, Y_2, \dots of random variables taking value in X , each defined on the same probability space $(\Gamma, \mathcal{A}_\Gamma, \mathbb{Q})$, which are i.i.d. according to ρ . An M -valued estimator of size n is then simply a random variable $\xi_n(Y_1, \dots, Y_n)$, where $\xi_n : X^n \rightarrow M$ is a measurable map. The equivalence between the two approaches is made clear by setting $(\Gamma, \mathcal{A}_\Gamma, \mathbb{Q}) \equiv (\Omega, \mathcal{A}_\Omega, \mathbb{P})$ and $Y_i(\omega) = x_i$ for all $\omega = (x_j)_{j \geq 1}$ and $i \geq 1$.

Concentration of measure results for random variables in Hilbert spaces can be used to prove that T_n is an unbiased estimator of T , as stated in the following lemma.

Lemma 1. For $n \geq 1$ and $\delta > 0$,

$$\|T - T_n\|_{\mathcal{S}_2} \leq \frac{2(\delta \vee \sqrt{2\delta})}{\sqrt{n}} \quad (22)$$

with probability at least $1 - 2e^{-\delta}$. Furthermore

$$\lim_{n \rightarrow \infty} \frac{\sqrt{n}}{\log n} \|T - T_n\|_{\mathcal{S}_2} = 0 \quad \text{almost surely.} \quad (23)$$

Proof. The result is known, but we report its short proof. For all $i \geq 1$ define the random variables $Z_i : \Omega \rightarrow \mathcal{S}_2$ as

$$Z_i(\omega) = K_{x_i} \otimes K_{x_i} \quad \omega = (x_j)_{j \geq 1} \in \Omega.$$

The fact that Z_i is measurable follows from Lemma 5 in Appendix A.2. Then, for all $i \geq 1$, we have $\|Z_i\|_{\mathcal{S}_2} \leq 1$ -almost surely, $\mathbb{E}[Z_i] = T$, and clearly $\mathbb{E}[\|Z_i\|_{\mathcal{S}_2}^2] \leq 1$. The first result follows easily applying Lemma 8 in Appendix A.4 and simplifying the right hand side of (A.4), and the second is a consequence of Lemma 9 in Appendix A.4. \square

Remark 6. Note that (23) and Theorem 2.19 in [64] imply that

$$\lim_{n \rightarrow \infty} \|T - T_n\|_{\mathcal{S}_1} = 0 \quad \text{almost surely.}$$

5.2. Consistency

We now choose a family of filter functions $(r_\lambda)_{\lambda > 0}$ and study the convergence of the associated estimators F_n and X_n introduced in Section 4.

5.2.1. Consistency of F_n

We begin proving convergence of the functions F_n defined in (16) to the function F_ρ in (9). We introduce the map $G_\lambda : X \rightarrow \mathbb{R}$ defined by

$$G_\lambda(x) = \langle r_\lambda(T)K_x, K_x \rangle \quad \forall x \in X,$$

which can be seen as the *infinite sample* analogue of F_n . Clearly, G_λ is a continuous function. For all sets $C \subset X$, we then have the following splitting of the error into two parts, the *sample error* and the *approximation error*

$$\sup_{x \in C} |F_n(x) - F_\rho(x)| \leq \underbrace{\sup_{x \in C} |F_n(x) - G_{\lambda_n}(x)|}_{\text{sample error}} + \underbrace{\sup_{x \in C} |G_{\lambda_n}(x) - F_\rho(x)|}_{\text{approximation error}}. \quad (24)$$

In order to prove consistency, we need to show that the left hand side goes to 0 as the sequence of regularization parameters $(\lambda_n)_{n \geq 1}$ tends to 0. This will be done separately for the approximation and the sample errors in the next two propositions.

Proposition 8. Under Assumption 2.b), if the sequence $(\lambda_n)_{n \geq 1}$ is such that $\lim_{n \rightarrow \infty} \lambda_n = 0$, then, for any compact subset $C \subset X$,

$$\lim_{n \rightarrow \infty} \sup_{x \in C} |G_{\lambda_n}(x) - F_\rho(x)| = 0.$$

Proof. Assumption 2.b) and $\lim_{n \rightarrow \infty} \lambda_n = 0$ imply that the sequence of non-negative functions $(r_{\lambda_n})_{n \geq 1}$ is bounded by 1 and converges pointwisely to the Heaviside function θ on the interval $[0, 1]$. Spectral theorem ensures that, for all $x \in C$,

$$\lim_{n \rightarrow \infty} r_{\lambda_n}(T)K_x = \theta(T)K_x. \quad (25)$$

Given $\epsilon > 0$, by compactness of C there exists a finite covering of C by balls of radius ϵ , namely $C \subset \bigcup_{i=1}^m B(x_i, \epsilon)$. By (25) there exists n_0 such that

$$\max_{i \in \{1, \dots, m\}} \|r_{\lambda_n}(T)K_{x_i} - \theta(T)K_{x_i}\| \leq \epsilon \quad \forall n \geq n_0.$$

Hence, for all $n \geq n_0$, we have

$$\begin{aligned} \sup_{x \in C} |G_{\lambda_n}(x) - F_\rho(x)| &= \sup_{x \in C} |(r_{\lambda_n}(T) - \theta(T))K_x| \\ &\leq \sup_{x \in C} \|K_x\| \sup_{x \in C} \|(r_{\lambda_n}(T) - \theta(T))K_x\| \\ &\leq \max_{i \in \{1, \dots, m\}} \sup_{x \in B(x_i, \epsilon)} \|(r_{\lambda_n}(T) - \theta(T))K_{x_i} + (r_{\lambda_n}(T) - \theta(T))(K_x - K_{x_i})\| \\ &\leq \max_{i \in \{1, \dots, m\}} \sup_{x \in B(x_i, \epsilon)} (\|(r_{\lambda_n}(T) - \theta(T))K_{x_i}\| + \|r_{\lambda_n}(T) - \theta(T)\|_\infty \|K_x - K_{x_i}\|) \\ &\leq \epsilon + \epsilon \sup_{\sigma \in [0, 1]} |r_{\lambda_n}(\sigma) - \theta(\sigma)| = 3\epsilon, \end{aligned}$$

where $\|K_x - K_{x_i}\| < \epsilon$ for all $x \in B(x_i, \epsilon)$ since $\|K_x - K_{x_i}\| = d_K(x, x_i)$, and, because $|r_{\lambda_n}(\sigma)| \leq 1$, $|\theta(\sigma)| \leq 1$, $\sup_{\sigma \in [0, 1]} |r_{\lambda_n}(\sigma) - \theta(\sigma)| \leq 2$. \square

Convergence to zero of the sample error follows from (23) and the next proposition.

Proposition 9. For all sets $C \subset X$ we have

$$\sup_{x \in C} |F_n(x) - G_{\lambda_n}(x)| \leq \|r_{\lambda_n}(T_n) - r_{\lambda_n}(T)\|_{\mathcal{S}_2}. \quad (26)$$

In particular, if [Assumption 2.c](#)) holds, then

$$\sup_{x \in C} |F_n(x) - G_{\lambda_n}(x)| \leq L_{\lambda_n} \|T_n - T\|_{\mathcal{S}_2}. \quad (27)$$

Proof. For all $x \in X$, we have the bound

$$\begin{aligned} |F_n(x) - G_{\lambda_n}(x)| &= |(r_{\lambda_n}(T_n) - r_{\lambda_n}(T))K_x| \\ &\leq \|r_{\lambda_n}(T_n) - r_{\lambda_n}(T)\|_\infty \|K_x\|^2 \\ &\leq \|r_{\lambda_n}(T_n) - r_{\lambda_n}(T)\|_{\mathcal{S}_2}, \end{aligned}$$

which proves (26). [Assumption 2.c](#)) and Theorem 8.1 in [9] (see also [Lemma 7](#) in [Appendix A.3](#) for a simple unpublished proof due to A. Maurer) imply that

$$\|r_{\lambda_n}(T_n) - r_{\lambda_n}(T)\|_{\mathcal{S}_2} \leq L_{\lambda_n} \|T_n - T\|_{\mathcal{S}_2}.$$

Inequality (27) then follows. \square

The above results can be combined in the following theorem, showing that, if the sequence λ_n is suitably chosen, then F_n converges almost surely to F_ρ with respect to the topology of uniform convergence on compact subsets of X .

Theorem 5. Under [Assumption 2](#), if the sequence $(\lambda_n)_{n \geq 1}$ is such that

$$\lim_{n \rightarrow \infty} \lambda_n = 0 \quad \text{and} \quad \sup_{n \geq 1} \frac{L_{\lambda_n} \log n}{\sqrt{n}} < +\infty, \quad (28)$$

then, for every compact subset $C \subset X$,

$$\lim_{n \rightarrow \infty} \sup_{x \in C} |F_n(x) - F_\rho(x)| = 0 \quad \text{almost surely.} \quad (29)$$

Proof. We show convergence to zero of both the two terms in the right hand side of inequality (24), thus implying (29). By (27), we have

$$\sup_{x \in C} |F_n(x) - G_{\lambda_n}(x)| \leq L_{\lambda_n} \|T_n - T\|_{\mathcal{S}_2} = \frac{L_{\lambda_n} \log n}{\sqrt{n}} \frac{\sqrt{n} \|T_n - T\|_{\mathcal{S}_2}}{\log n} \leq M \frac{\sqrt{n} \|T_n - T\|_{\mathcal{S}_2}}{\log n},$$

where $M = \sup_{n \geq 1} (L_{\lambda_n} \log n) / \sqrt{n}$ is finite by (28). Then (23) implies that the first term in the right hand side of inequality (24) converges to zero almost surely. Since the second term goes to zero by [Proposition 8](#), the claim follows. \square

5.2.2. Consistency of X_n

As already remarked above, uniform convergence of F_n to F_ρ on compact subsets *does not* imply convergence of the level sets of F_n to the corresponding level sets of F_ρ in any sense (as, for example, with respect to the Hausdorff distance among compact subsets). For this reason, we introduce a family of threshold parameters $(\tau_n)_{n \geq 1}$ and define the estimator X_n of the set X_ρ as in (17).

We define a data dependent parameter τ_n as the function on Ω

$$\tau_n(\omega) = 1 - \min_{1 \leq i \leq n} [F_n(\omega)](x_i) \quad \omega = (x_i)_{i \geq 1}, \quad (30)$$

where we wrote explicitly the dependence of F_n on the training set $\omega \in \Omega$. Since F_n takes values in $[0, 1]$, clearly $\tau_n(\omega) \in [0, 1]$.

Proposition 10. *Suppose the metric space X is compact. Then, under Assumption 2, the function τ_n is an \mathbb{R} -valued estimator. Moreover, if the sequence $(\lambda_n)_{n \geq 1}$ satisfies (28), we have*

$$\lim_{n \rightarrow \infty} \tau_n = 0 \quad \text{almost surely.}$$

Proof. The proof that τ_n is an \mathbb{R} -valued estimator is of technical nature, and we postpone it to Proposition 17 in Appendix A.2.

Here we prove that $\lim_{n \rightarrow \infty} \tau_n = 0$ with probability 1. By Theorem 5, we can find an event $E_1 \subset \Omega$ with $\mathbb{P}(E_1) = 1$ such that $\lim_{n \rightarrow \infty} \sup_{x \in X} |F_n(x) - F_\rho(x)| = 0$ on E_1 . Moreover, for the event $E_2 = \{x_i \in X_\rho \text{ for all } i \geq 1\}$, we clearly have $\mathbb{P}(E_2) = 1$ by definition of X_ρ and \mathbb{P} . If $\omega \in E_1 \cap E_2$ and $\epsilon > 0$ is fixed, then there exists $n_0 \geq 1$ (possibly depending on ω and ϵ) such that for all $n \geq n_0$ $|[F_n(\omega)](x) - F_\rho(x)| \leq \epsilon$ for all $x \in X$. Since $F_\rho(x) = 1$ for all $x \in X_\rho$ by definition and $x_1, \dots, x_n \in X_\rho$, it follows that $|[F_n(\omega)](x_i) - 1| \leq \epsilon$ for all $1 \leq i \leq n$, that is

$$0 \leq 1 - [F_n(\omega)](x_i) \leq \epsilon \quad \forall i \in \{1, 2, \dots, n\},$$

so that $0 \leq \tau_n(\omega) \leq \epsilon$. Thus, $\lim_{n \rightarrow \infty} \tau_n(\omega) = 0$, and, since $\mathbb{P}(E_1 \cap E_2) = 1$, the sequence $(\tau_n)_{n \geq 1}$ goes to zero with probability 1. \square

The following is the central result of this section. It shows that, assuming X is compact and for the above choice of the sequence $(\tau_n)_{n \geq 1}$, the Hausdorff distance between X_n and X_ρ goes to zero with probability 1. Here we recall that the Hausdorff distance between two subsets $A, B \subset X$ is

$$d_H(A, B) = \max \left\{ \sup_{a \in A} d_K(a, B), \sup_{b \in B} d_K(b, A) \right\},$$

where $d_K(x, Y) = \inf_{y \in Y} d_K(x, y)$.

Theorem 6. *Suppose the metric space X is compact. Under Assumption 2, if \mathcal{H} separates the set X_ρ and the sequence $(\lambda_n)_{n \geq 1}$ satisfies (28), for the choice of the threshold parameters $(\tau_n)_{n \geq 1}$ given in (30) we have*

$$\lim_{n \rightarrow \infty} d_H(X_n, X_\rho) = 0 \quad \text{almost surely.}$$

We devote the rest of this section to proof of the above theorem. For simplicity, we split it into a few lemmas.

Lemma 2. *Under the hypotheses of Theorem 6, we have*

$$\lim_{n \rightarrow \infty} \sup_{x \in X_n} d_K(x, X_\rho) = 0 \quad \text{almost surely.} \quad (31)$$

Proof. Let E be the event $E = \{\lim_{n \rightarrow \infty} \tau_n = 0\}$. Then, $\mathbb{P}(E) = 1$ by Proposition 10. We fix $\omega \in E$, and suppose by contradiction that at such ω the limit (31) does not hold. Then (depending on ω) there exists $\epsilon > 0$ such that for all k there is $n_k \geq k$ satisfying the inequality $\sup_{x \in X_{n_k}} d_K(x, X_\rho) \geq 2\epsilon$. Hence there is $z_k \in X_{n_k}$ such that

$$d_K(z_k, X_\rho) \geq \epsilon \quad \text{for all } x \in X_\rho. \quad (32)$$

Since X is compact, possibly passing to a subsequence we can assume that the sequence $(z_k)_{k \geq 1}$ converges to a limit $z \in X$. We claim that $z \in X_\rho$. Indeed, if k is sufficiently large, then we have

$$\begin{aligned} |F_\rho(z) - 1| &\leq |F_\rho(z) - F_\rho(z_k)| + |F_\rho(z_k) - F_{n_k}(z_k)| + |F_{n_k}(z_k) - 1| \\ &\leq |F_\rho(z) - F_\rho(z_k)| + \sup_{x \in X} |F_\rho(x) - F_{n_k}(x)| + \tau_{n_k}, \end{aligned}$$

where $|F_{n_k}(z_k) - 1| \leq \tau_{n_k}$ is due to the fact that $z_k \in X_{n_k}$, so that

$$1 + \tau_{n_k} \geq 1 \geq F_{n_k}(z_k) \geq 1 - \tau_{n_k}.$$

As n_k goes to ∞ , we have $\sup_{x \in X} |F_\rho(x) - F_{n_k}(x)| \rightarrow 0$ by [Theorem 5](#); moreover, since F_ρ is continuous in z and τ_{n_k} goes to zero, the above inequality for $|F_\rho(z) - 1|$ gives $F_\rho(z) = 1$. Since \mathcal{H} separates X_ρ , this implies $z \in X_\rho$. However, [\(32\)](#) implies that $d_K(z, x) \geq \epsilon$ for all $x \in X_\rho$, which is the desired contradiction. \square

The proof that $\sup_{x \in X_n} d_K(x, X_\rho)$ goes to zero as $n \rightarrow \infty$ requires a further technical lemma, see [\[36, Lemma 6.1\]](#). In its statement, for all $n \geq 1$ and $x \in X$, we denote by $\xi_{1,n}(x)$ the nearest neighbor of x in the training set $\{x_1, \dots, x_n\}$, i.e.

$$\xi_{1,n}(x) = \arg \min_{x_1, x_2, \dots, x_n} d_K(x_i, x).$$

Lemma 3. For all $x \in X_\rho$,

$$\lim_{n \rightarrow \infty} d_K(\xi_{1,n}(x), x) = 0 \quad \text{almost surely.}$$

Proof. Given $x \in X_\rho$, fix $\epsilon > 0$ and, denoted by $B(x, \epsilon)$ the closed ball with center x and radius ϵ , set $p = \rho(B(x, \epsilon))$. By definition of the support and the fact that ρ is a probability measure, $0 < p \leq 1$. Furthermore

$$\begin{aligned} \mathbb{P}(d_K(\xi_{1,n}(x), x) > \epsilon) &= \mathbb{P}(x_i \notin B(x, \epsilon) \forall i = 1, \dots, n) \\ &\text{(by independence of the } x_i\text{'s)} = \prod_{i=1}^n \mathbb{P}(x_i \notin B(x, \epsilon)) \\ &\text{(since the } x_i\text{'s are identically distributed)} = \prod_{i=1}^n (1 - \rho(B(x, \epsilon))) \\ &= (1 - p)^n. \end{aligned}$$

Since $0 \leq 1 - p < 1$, the series $\sum_n (1 - p)^n$ converges, so that Borel–Cantelli lemma yields

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} \bigcap_{m=n}^{\infty} \{d_K(\xi_{1,m}(x), x) \leq \epsilon\}\right) = 1.$$

Since this holds for all $\epsilon > 0$, we have

$$\mathbb{P}\left(\bigcap_{k=1}^{\infty} \bigcup_{n=1}^{\infty} \bigcap_{m=n}^{\infty} \left\{d_K(\xi_{1,m}(x), x) \leq \frac{1}{k}\right\}\right) = 1,$$

and the lemma follows. \square

Lemma 4. Under the hypotheses of [Theorem 5](#), if the metric space X is compact, then

$$\lim_{n \rightarrow \infty} \sup_{x \in X_\rho} d_K(x, X_n) = 0 \quad \text{almost surely.} \tag{33}$$

Proof. Choose a denumerable dense family $\{z_j\}_{j \in J}$ in X_ρ . By [Lemma 3](#) there exists an event E with probability 1 such that

$$\lim_{n \rightarrow +\infty} d_K(\xi_{1,n}(z_j), z_j) = 0 \quad \forall j \in J \tag{34}$$

on E . We claim that the limit [\(33\)](#) holds on E . Observe that, by definition of τ_n , $x_i \in X_n$ for all $1 \leq i \leq n$, and

$$\sup_{x \in X_\rho} d_K(x, X_n) \leq \sup_{x \in X_\rho} \min_{1 \leq i \leq n} d_K(x, x_i) = \sup_{x \in X_\rho} d_K(\xi_{1,n}(x), x),$$

so that it is enough to show that $\lim_{n \rightarrow +\infty} \sup_{x \in X_\rho} d_K(\xi_{1,n}(x), x) = 0$.

Fix $\epsilon > 0$. Since X_ρ is compact, there is a finite subset $J_\epsilon \subset J$ such that $\{B(z_j, \epsilon)\}_{j \in J_\epsilon}$ is a finite covering of X_ρ . We claim that

$$\sup_{x \in X_\rho} d_K(\xi_{1,n}(x), x) \leq \max_{j \in J_\epsilon} d_K(\xi_{1,n}(z_j), z_j) + \epsilon. \tag{35}$$

Indeed, fixed $x \in X_\rho$, there exists an index $j \in J_\epsilon$ such that $x \in B(z_j, \epsilon)$. By definition of $\xi_{1,n}$, clearly

$$d_K(\xi_{1,n}(x), x) \leq d_K(\xi_{1,n}(z_j), x),$$

so that by the triangular inequality we get

$$\begin{aligned} d_K(\xi_{1,n}(x), x) &\leq d_K(\xi_{1,n}(z_j), x) \leq d_K(\xi_{1,n}(z_j), z_j) + d_K(z_j, x) \\ &\leq d_K(\xi_{1,n}(z_j), z_j) + \epsilon \\ &\leq \max_{j \in J_\epsilon} d_K(\xi_{1,n}(z_j), z_j) + \epsilon. \end{aligned}$$

Taking the sup over X_ρ we get the claim.

Since J_ϵ is finite, by (34)

$$\lim_{n \rightarrow +\infty} \max_{j \in J_\epsilon} d_K(\xi_{1,n}(z_j), z_j) = 0,$$

hence (35) yields

$$\limsup_{n \rightarrow \infty} \sup_{x \in X_\rho} d_K(\xi_{1,n}(x), x) \leq \epsilon.$$

Since ϵ is arbitrary, we get $\lim_{n \rightarrow +\infty} \sup_{x \in X_\rho} d_K(\xi_{1,n}(x), x) = 0$, and this concludes the proof \square

The proof of [Theorem 6](#) follows easily combining the previous lemmas.

Proof of Theorem 6. As $d_H(X_n, X_\rho) = \max\{\sup_{x \in X_n} d_K(x, X_\rho), \sup_{x \in X_\rho} d_K(x, X_n)\}$, the theorem follows combining [Lemmas 2 and 4](#). \square

We conclude this section with some comments. First, if \mathcal{H} does not separate X_ρ , then the statement of [Theorem 6](#) continues to be true provided that the support X_ρ is replaced by the level set $\{x \in X \mid F_\rho(x) = 1\}$. Note that, the Hausdorff distance d_H has been defined with respect to the metric d_K induced by the kernel, however, if the set X has its own metric d_X making it compact and the hypotheses of [Proposition 3](#) are satisfied, then [Theorem 6](#) implies convergence of X_n to X_ρ also with respect to the Hausdorff distance associated to d_X . Finally, we remark that in [Theorem 6](#) convergence of X_n to X_ρ does not depend on any a priori assumption on the probability ρ .

5.3. Finite sample bounds and stability of random sampling

In order to prove stability of our algorithms under random sampling and determine their convergence rates, we need to specify suitable a priori assumptions on the class of problems to be considered. In the present section, a detailed analysis of the convergence rates of F_n to F_ρ will be carried out for the case of the Tikhonov filter $r_\lambda(\sigma) = \sigma / (\sigma + \lambda)$. The techniques in [\[14\]](#) should allow to derive similar results for filters other than Tikhonov.

For all $\lambda > 0$ we define

$$\mathcal{N}(\lambda) = \text{tr}[(T + \lambda)^{-1}T] = \sum_{j \in J} \frac{\sigma_j}{\sigma_j + \lambda},$$

which is finite since T is a trace class operator. The above quantity is related to the degrees of freedom of the estimator [\[38\]](#). Here, we recall that \mathcal{N} is a decreasing function of λ and $\lim_{\lambda \rightarrow 0^+} \mathcal{N}(\lambda) = N$, where N is the dimension of the range of T .

The a priori conditions we consider in the present paper are given by the following two assumptions, which involve both the reproducing kernel K and the probability measure ρ (compare with [\[12,13\]](#)).

Assumption 3. We assume that

a) there exist $b \in [0, 1]$ and $D_b \geq 1$ such that

$$\sup_{\lambda > 0} \mathcal{N}(\lambda) \lambda^b \leq D_b^2; \tag{36}$$

b) there exist $0 < s \leq 1$ and a constant $C_s > 0$ such that $P_\rho K_x \in \text{ran } T^{s/2}$ for all $x \in X$, and

$$\sup_{x \in X} \|T^{-\frac{s}{2}} P_\rho K_x\|^2 \leq C_s. \tag{37}$$

The above conditions are classical in the theory of inverse problems and have been recently considered in supervised learning. Before showing how they allow to derive a finite sample bound on the error $\sup_{x \in X} |F_n(x) - F_\rho(x)|$, we add some comments. First, [Assumption 3.a](#)) is related to the level of ill-posedness of the problem [\[34\]](#) and can be interpreted as a condition specifying the *aspect ratio* of the range of T . Since $0 < \lambda \mathcal{N}(\lambda) < \text{tr}[T] = 1$, inequality [\(36\)](#) is always satisfied with the choice $b = 1$ and $D_1 = 1$, so that in this case we are not imposing any a priori assumption. If $\dim \text{ran } T = N < \infty$, the best choice is $b = 0$ and $D_0 = \sqrt{N}$; otherwise, if $\dim \text{ran } T = \infty$, then necessarily $b > 0$. In the latter case, a sufficient condition to have $b < 1$ is to assume a decay rate $\sigma_j \sim j^{-1/b}$ on the eigenvalues of T (see [Proposition 3](#) of [\[13\]](#)).

Coming to [Assumption 3.b](#)), first of all we remark that it is always satisfied when $\dim \text{ran } T$ is finite with the choice $s = 1$ and $C_1 = \max_{j \in J} 1/\sigma_j$. In the general case, [Assumption 3.b](#)) can be expressed by the following equivalent condition

$$\sum_{j \in J} \sigma_j^{1-s} |\phi_j(x)|^2 \leq C_s \quad \forall x \in X, \quad (38)$$

where $(\phi_j, \sigma_j)_{j \in J}$ are the eigenvectors and eigenvalues of L_K , which were defined in [Section 2.3](#) (see in particular [\(7\)](#) for the definition of the functions ϕ_j outside the set X_ρ). Clearly, the higher is s , the stronger is the assumption.

Note that in particular inequality [\(38\)](#) holds true if there exists a constant¹ $\kappa > 0$ such that $\sup_{x \in X} |\phi_j(x)| \leq \kappa$ for all $j \in J$, and $s \in]0, 1[$ is chosen to make the series $\sum_{j \in J} \sigma_j^{1-s}$ finite. In this case, it is quite easy to give conditions on the eigenvalues $(\sigma_j)_{j \in J}$ assuring that both [Assumptions 3.a](#)) and [3.b](#)) are satisfied. For example, if $\sigma_j \sim j^{-1/b}$ for some $0 < b < 1$, then [\(36\)](#) holds true with this choice of b , and [\(37\)](#) is satisfied for any $0 < s < 1 - b$.

Remark 7. Setting $\beta = 1 - s \in [0, 1[$, condition [\(38\)](#) is equivalent to the fact that for all $x, y \in X$ the series

$$K_\rho^\beta(x, y) = \sum_{j \in J} \sigma_j^\beta \phi_j(y) \phi_j(x) \quad (39)$$

converges absolutely to a bounded reproducing kernel K_ρ^β . Convergence of the series [\(39\)](#) was studied e.g. in [\[70\]](#), where it is proved that, if the sequence of powers $(\sigma_j^\beta)_{j \in J}$ is summable, there exists a ρ -null set N such that [\(39\)](#) converges absolutely on $(X \setminus N) \times (X \setminus N)$ (see [\[70, Proposition 4.4\]](#)). We remark that this weaker fact is not sufficient in our setting: indeed, on the one hand it does not imply that the series [\(39\)](#) (or, equivalently, [\(38\)](#)) converges on all of X , and on the other it does not guarantee that such series is uniformly bounded, two conditions which however are both needed in the proof of [Theorem 7](#) below to get uniform estimates on the whole set X . A direction of future work is to study the geometric nature of the above conditions when X is a metric space or a Euclidean space and X_ρ a Riemannian submanifold.

The following theorem provides the finite sample bound on the error $\sup_{x \in X} |F_n(x) - F_\rho(x)|$.

Theorem 7. Suppose $r_\lambda(\sigma) = \sigma / (\sigma + \lambda)$. If [Assumption 3](#) holds and we choose

$$\lambda_n = \left(\frac{1}{n} \right)^{\frac{1}{2s+b+1}},$$

then, for $n \geq 1$ and $\delta > 0$, we have

$$\sup_{x \in X} |F_n(x) - F_\rho(x)| \leq (C_s \vee (D_b(2\delta \vee \sqrt{2\delta}))) \left(\frac{1}{n} \right)^{\frac{s}{2s+b+1}} \quad (40)$$

with probability at least $1 - 2e^{-\delta}$.

We postpone the proof to the end of the current section and add here some comments. The above finite sample bound quantifies the stability of the estimator with respect to random sampling. Equivalently, if we set the right hand term of the inequality to ϵ and solve for $n = n(\epsilon, \delta)$, we obtain the sample complexity of the problem, i.e. how many samples are needed in order to achieve the maximum error ϵ with confidence $1 - 2e^{-\delta}$. As remarked before, [Assumption 3.a](#)) is verified for $b = 1$ by any reproducing kernel. In this limit case our result gives a rate $n^{-s/(2s+2)}$, comparable with the one that can be obtained inserting [\(27\)](#) and [\(41\)](#) below into inequality [\(24\)](#), with $\|T_n - T\|$ bounded by [\(22\)](#).

Note that, if $\dim \text{ran } T = N < \infty$, choosing $b = 0$, $D_0 = \sqrt{N}$, $s = 1$ and $C_1 = \max_{j \in J} 1/\sigma_j$, the rate in [\(40\)](#) becomes $n^{-1/3}$.

The proof of [Theorem 7](#) follows the ideas in [\[13\]](#) and is based on refined estimates of the sample and approximation errors. The techniques in [\[14\]](#) should allow to derive similar results for filters beyond the Tikhonov one.

¹ As it happens for example for reproducing kernels on $X = [0, 2\pi]^d$ which are invariant under translations, when ρ is the Lebesgue measure on $[0, 2\pi]^d$.

Proposition 11. If *Assumption 3.a*) holds true, then, for $n \geq 1$ and $\delta > 0$, we have

$$\sup_{x \in X} |F_n(x) - G_{\lambda_n}(x)| \leq \left(\frac{\delta}{n\lambda_n} + \sqrt{\frac{2\delta\mathcal{N}(\lambda_n)}{n\lambda_n}} \right)$$

with probability at least $1 - 2e^{-\delta}$.

Proof. Consider the following decomposition

$$\begin{aligned} r_{\lambda_n}(T) - r_{\lambda_n}(T_n) &= (T + \lambda_n)^{-1}T - (T_n + \lambda_n)^{-1}T_n \\ &= (T + \lambda_n)^{-1}T - (T + \lambda_n)^{-1}T_n + (T + \lambda_n)^{-1}T_n - (T_n + \lambda_n)^{-1}T_n \\ &= (T + \lambda_n)^{-1}(T - T_n) + (T + \lambda_n)^{-1}[(T_n + \lambda_n) - (T + \lambda_n)](T_n + \lambda_n)^{-1}T_n \\ &= (T + \lambda_n)^{-1}(T - T_n) + (T + \lambda_n)^{-1}(T_n - T)(T_n + \lambda_n)^{-1}T_n \\ &= (T + \lambda_n)^{-1}(T - T_n)[I - (T_n + \lambda_n)^{-1}T_n] \\ &= \lambda_n(T + \lambda_n)^{-1}(T - T_n)(T_n + \lambda_n)^{-1}. \end{aligned}$$

It is easy to see that $\|(T_n + \lambda_n)^{-1}\|_\infty \leq \lambda_n^{-1}$, hence

$$\|r_{\lambda_n}(T) - r_{\lambda_n}(T_n)\|_{\mathcal{S}_2} \leq \lambda_n \|(T + \lambda_n)^{-1}(T - T_n)\|_{\mathcal{S}_2} \|(T_n + \lambda_n)^{-1}\|_\infty \leq \|(T + \lambda_n)^{-1}(T - T_n)\|_{\mathcal{S}_2}.$$

Then, from [Lemma 10](#) in [Appendix A.4](#) we have that

$$\|(T + \lambda_n I)^{-1}(T - T_n)\|_{\mathcal{S}_2} \leq \left(\frac{\delta}{n\lambda_n} + \sqrt{\frac{2\delta\mathcal{N}(\lambda_n)}{n\lambda_n}} \right),$$

with probability at least $1 - 2e^{-\delta}$, so that the result follows by [\(26\)](#). \square

Proposition 12. If *Assumption 3.b*) holds true, then

$$\sup_{x \in X} |G_\lambda(x) - F_\rho(x)| \leq \lambda^s C_s. \tag{41}$$

Proof. Since $\theta(\sigma) - r_\lambda(\sigma) = \lambda/(\sigma + \lambda)$ for all $\sigma > 0$, we have

$$\begin{aligned} |G_\lambda(x) - F_\rho(x)| &= |((r_\lambda(T) - \theta(T))K_x, K_x)| = |((r_\lambda(T) - \theta(T))P_\rho K_x, P_\rho K_x)| \\ &= \lambda \|(T + \lambda)^{-\frac{1}{2}} P_\rho K_x\|^2, \end{aligned}$$

as $P_\rho K_x \in \ker T^\perp$. Since by assumption $P_\rho K_x \in \text{ran } T^{s/2}$ for some $0 < s \leq 1$, spectral calculus and the bound $\sigma^s/(\sigma + \lambda) \leq \lambda^{s-1}$ give the inequality

$$\|(T + \lambda)^{-\frac{1}{2}} P_\rho K_x\|^2 = \|(T + \lambda)^{-1} T^{\frac{s}{2}} T^{-\frac{s}{2}} P_\rho K_x\|^2 \leq \lambda^{s-1} \|T^{-\frac{s}{2}} P_\rho K_x\|^2,$$

so that

$$|G_\lambda(x) - F_\rho(x)| \leq \lambda^s \|T^{-\frac{s}{2}} P_\rho K_x\|^2 \leq \lambda^s C_s$$

for all $x \in X$. \square

We are now ready to prove the main result.

Proof of Theorem 7. The choice $\lambda_n = n^{-1/(2s+b+1)}$ is the one that set the contributions of the sample and approximation errors in [\(24\)](#) to be equal. Indeed, we begin by simplifying the bound on the sample error. If $\lambda \geq n^{-1}$, then $n\lambda \geq \sqrt{n\lambda^{b+1}}$ for all $0 < b \leq 1$, so that

$$\frac{\delta}{n\lambda} + \sqrt{\frac{2\delta\mathcal{N}(\lambda)}{n\lambda}} = \frac{\delta}{n\lambda} + \sqrt{\frac{2\delta\mathcal{N}(\lambda)\lambda^b}{n\lambda^{b+1}}} \leq D_b(\delta \vee \sqrt{2\delta}) \left(\frac{1}{n\lambda} + \frac{1}{\sqrt{n\lambda^{b+1}}} \right) \leq \frac{2D_b(\delta \vee \sqrt{2\delta})}{\sqrt{n\lambda^{b+1}}},$$

where we used the definition of D_b (and the fact that $D_b \geq 1$). Then, by the above inequality and [Propositions 11 and 12](#), inequality [\(24\)](#) gives

$$\sup_{x \in X} |F_n(x) - F_\rho(x)| \leq C_s \lambda^s + \frac{2D_b(\delta \vee \sqrt{2\delta})}{\sqrt{n\lambda^{b+1}}}. \quad (42)$$

If we set the contributions of the sample and approximation errors to be equal, the choice for λ is

$$\lambda = \left(\frac{1}{n}\right)^{\frac{1}{2s+b+1}}.$$

It is easy to see that $\lambda \geq n^{-1}$ for all values of s, b , so that from (42) we have

$$\sup_{x \in X} |F_n(x) - F_\rho(x)| \leq (C_s \vee (2D_b(\delta \vee \sqrt{2\delta}))) \left(\frac{1}{n}\right)^{\frac{s}{2s+b+1}}. \quad \square$$

5.4. The kernel PCA filter

A natural choice for the spectral filter r_λ would be the regularization defined by kernel PCA [62], that corresponds to truncating the generalized inverse of the kernel matrix at some cut-off parameter λ . The corresponding filter function is

$$r_\lambda(\sigma) = \begin{cases} 1 & \sigma \geq \lambda, \\ 0 & \sigma < \lambda. \end{cases}$$

The above filter does not satisfy the Lipschitz condition 2.c) in [Assumption 2](#), so that the bound (27) for the sample error $\sup_{x \in X} |F_n(x) - G_{\lambda_n}(x)|$ does not hold in this case.² However, we can still achieve an estimate by employing inequality (A.1) in [Appendix A.3](#). To this aim, with a slight abuse of the notation, here we count the eigenvalues of T and T_n without their multiplicities and we list them in decreasing order. Furthermore, for any $\lambda > 0$ we set $\sigma_{j(\lambda)}$ and $\sigma_{k(\lambda)}^{(n)}$ as the smallest eigenvalues of T and T_n which are greater or equal to λ , i.e.

$$\sigma_1 > \sigma_2 > \dots > \sigma_{j(\lambda)} \geq \lambda > \sigma_{j(\lambda)+1} \quad \sigma_1^{(n)} > \sigma_2^{(n)} > \dots > \sigma_{k(\lambda)}^{(n)} \geq \lambda > \sigma_{k(\lambda)+1}^{(n)}.$$

Inequality (A.1) implies that

$$\|r_\lambda(T_n) - r_\lambda(T)\|_{\mathcal{S}_2} \leq \frac{\|T_n - T\|_{\mathcal{S}_2}}{\min\{\sigma_{j(\lambda)} - \sigma_{k(\lambda)+1}^{(n)}, \sigma_{k(\lambda)}^{(n)} - \sigma_{j(\lambda)+1}\}} \leq \frac{\|T_n - T\|_{\mathcal{S}_2}}{\min\{\sigma_{j(\lambda)} - \lambda, \lambda - \sigma_{j(\lambda)+1}\}},$$

and inequality (26) for the sample error then reads

$$\sup_{x \in C} |F_n(x) - G_{\lambda_n}(x)| \leq \frac{\|T_n - T\|_{\mathcal{S}_2}}{\min\{\sigma_{j(\lambda_n)} - \lambda_n, \lambda_n - \sigma_{j(\lambda_n)+1}\}}.$$

By [Lemma 1](#), in order to have convergence to 0 of the right hand side of this expression we need to choose the sequence $(\lambda_n)_{n \geq 1}$ such that

$$\sup_{n \geq 1} \frac{\log n}{\sqrt{n} \min\{\sigma_{j(\lambda_n)} - \lambda_n, \lambda_n - \sigma_{j(\lambda_n)+1}\}} < \infty.$$

Since the gap $\sigma_{j(\lambda)} - \sigma_{j(\lambda)+1}$ can have any arbitrary rate of convergence to zero as $\lambda \rightarrow 0^+$, we thus see that there exists no distribution independent choice of $(\lambda_n)_{n \geq 1}$ ensuring the convergence to zero of the above bound.

Note that $r_\lambda(T)$ is the projection $P_{j(\lambda)}$ onto the sum of the eigenspaces of the first $j(\lambda)$ eigenvalues of T and $r_\lambda(T_n)$ is the projection $P_{k(\lambda)}^{(n)}$ onto the sum of the eigenspaces of the first $k(\lambda)$ eigenvalues of T . If $(M_n)_{n \geq 1}$ is any strictly increasing sequence with $M_n \in \mathbb{N}$ for all n , we can consider the following distribution dependent choice $\lambda_n = (\sigma_{M_n} + \sigma_{M_n+1})/2$. Then we have

$$\|P_{M_n}^{(n)} - P_{M_n}\|_{\mathcal{S}_2} = \|r_{\lambda_n}(T_n) - r_{\lambda_n}(T)\|_{\mathcal{S}_2} \leq \frac{2\|T_n - T\|_{\mathcal{S}_2}}{\sigma_{M_n} - \sigma_{M_n+1}},$$

which recovers a known result about kernel PCA (see for example [77]). Furthermore, if we have that $\|T_n - T\|_{\mathcal{S}_2} < (\sigma_{M_n} - \sigma_{M_n+1})/2$, then we obtain $\|P_{M_n}^{(n)} - P_{M_n}\|_{\mathcal{S}_2} < 1$, hence $\dim \text{ran } P_{M_n}^{(n)} = \dim \text{ran } P_{M_n}$.

The following result extends [Theorem 5](#) to the case of kernel PCA, at the price of having a distribution dependent choice of the cut-off sequence $(M_n)_{n \geq 1}$.

² Note that, by [Proposition 16](#) in [Appendix A.2](#), if X is locally compact, then F_n defined in (16) still is a $C(X)$ -valued estimator.

Theorem 8. *If the sequence of natural numbers $(M_n)_{n \geq 1}$ is strictly increasing and such that*

$$\sup_{n \geq 1} \frac{\log n}{\sqrt{n}(\sigma_{M_n} - \sigma_{M_{n+1}})} < +\infty$$

and we define the sequence $(\lambda_n)_{n \geq 1}$ as

$$\lambda_n = \frac{\sigma_{M_n} + \sigma_{M_{n+1}}}{2},$$

then, for every compact subset $C \subset X$,

$$\lim_{n \rightarrow \infty} \sup_{x \in C} |F_n(x) - F_\rho(x)| = 0 \quad \text{almost surely.}$$

Proof. By the above discussion and inequality (26),

$$\sup_{x \in C} |F_n(x) - G_{\lambda_n}(x)| \leq \frac{2\|T_n - T\|_{\mathcal{S}_2}}{\sigma_{M_n} - \sigma_{M_{n+1}}} \leq \frac{\sqrt{n}\|T_n - T\|_{\mathcal{S}_2}}{\log n} \sup_{n \geq 1} \frac{2 \log n}{\sqrt{n}(\sigma_{M_n} - \sigma_{M_{n+1}})}.$$

Convergence to 0 of the sample error then follows from (23). Combining this fact and Proposition 8 into inequality (24), the claim then follows. \square

6. Some perspectives

In this section we discuss some different perspectives to our approach and suggest some possible extensions.

6.1. Connection to Mercer theorem

We start discussing some connections between our analytical characterization of the support of ρ and Mercer theorem [50]. With the notations of Section 2.3, the fact that the family $(\sqrt{\sigma_j}\phi_j)_{j \in J}$ is an orthonormal basis of $P_\rho\mathcal{H}$ and the reproducing property give the relation

$$\langle P_\rho K_y, K_x \rangle = \sum_{j \in J} \sigma_j \phi_j(x) \phi_j(y) \quad \forall x, y \in X, \quad (43)$$

where the series converges absolutely. Note that in this expression the eigenfunctions ϕ_j of L_K are defined outside X_ρ through the extension equation (7). Restricting (43) to $x, y \in X_\rho$, we obtain

$$K(x, y) = \sum_{j \in J} \sigma_j \phi_j(x) \phi_j(y) \quad \forall x, y \in X_\rho,$$

which is nothing else than Mercer theorem [68]. In particular, taking $x = y$, this formula implies that $\sum_{j \in J} \sigma_j |\phi_j(x)|^2 = K(x, x)$ for all $x \in X_\rho$. On the other hand, the assumption that the reproducing kernel separates X_ρ precisely ensures that

$$\sum_{j \in J} \sigma_j |\phi_j(x)|^2 \neq K(x, x) \quad \forall x \notin X_\rho.$$

(Recall that, if K separates X_ρ , then X_ρ is the 1-level set of the function $F_\rho = \sum_{j \in J} \sigma_j |\phi_j|^2$.)

6.2. A feature space point of view

In machine learning, kernel methods are often described in terms of a corresponding feature map [74]. This point of view highlights the linear structure of the Hilbert space and often provides a more geometric interpretation.

We recall that a feature map associated to a reproducing kernel is a map $\Psi : X \rightarrow \mathcal{F}$, where \mathcal{F} is a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{F}}$, satisfying $K(x, y) = \langle \Psi(y), \Psi(x) \rangle_{\mathcal{F}}$. While every map Ψ from X into a Hilbert space \mathcal{F} defines a reproducing kernel, it is also possible to prove that each kernel has an associated feature map (and in fact many). Indeed, given K , the natural assignment is $\mathcal{F} \equiv \mathcal{H}$ and $\Psi(x) \equiv K_x$. Such a choice is also minimal, in the sense that, if we make a different choice of \mathcal{F} and Ψ , then there exists an isometry $W : \mathcal{H} \rightarrow \mathcal{F}$ such that $\Psi(x) = WK_x \forall x \in X$ – see for example Proposition 2.4 of [15] or Theorem 4.21 of [68], noticing that both papers deal with the transpose $W^\top : \mathcal{F} \rightarrow \mathcal{H}$.

We next review some of the concepts introduced in Section 2 in terms of feature maps. For the sake of comparison we assume that $\|\Psi(x)\|_{\mathcal{F}} = 1$ for all $x \in X$ (this corresponds to the normalization Assumption 1.d)), we let \mathcal{F}_C be the closure of the linear span of the set $\{\Psi(x) \mid x \in C\}$, and define




Fig. 2. The sets X and the support X_ρ are mapped into the feature space \mathcal{F} , by the feature map Ψ . Here we take $\mathcal{F}_\rho = \mathcal{F}_{X_\rho}$ to be a linear space passing through the origin. The image of the support with respect to the feature map is given by the intersection of the image of X with \mathcal{F}_ρ . By the separating property, a point x belongs to the support if and only if the distance between $\Psi(x)$ and \mathcal{F}_ρ is zero.

$$d_{\mathcal{F}}(\Psi(x), \mathcal{F}_C) = \inf_{f \in \mathcal{F}_C} \|\Psi(x) - f\|_{\mathcal{F}}.$$

It is easy to see that the definition of separating kernel has the following equivalent and natural analogue in the context of feature maps.

Definition 3. We say that a feature map Ψ separates a subset $C \subset X$ if

$$d_{\mathcal{F}}(\Psi(x), \mathcal{F}_C) = 0 \iff x \in C.$$

The above definition is equivalent to [Definition 1](#) since $d_{\mathcal{F}}(\Psi(x), \mathcal{F}_C) = \|\Psi(x) - Q_C \Psi(x)\|_{\mathcal{F}}$, where Q_C is the orthogonal projection onto \mathcal{F}_C . Then, according to [Definition 3](#), a point $x \in C$ if and only if $\|\Psi(x) - Q_C \Psi(x)\|_{\mathcal{F}}^2 = 0$. Since $\Psi(x) = WK_x$ $\forall x \in X$ and $Q_C W = W P_C$, this is equivalent to

$$0 = \|\Psi(x) - Q_C \Psi(x)\|_{\mathcal{F}}^2 = \|K_x - P_C K_x\|^2 = K(x, x) - F_C(x).$$

[Theorem 1](#) then implies that [Definitions 1 and 3](#) are equivalent. We thus see that the separating property has a clear geometric interpretation in the feature space: the set $\Psi(C)$ is the intersection of the closed subspace \mathcal{F}_C , i.e. a linear manifold in \mathcal{F} , and $\Psi(X)$ – see [Fig. 2](#).

In the above interpretation, the estimator we propose for the support then stems from the following observation: given a training set x_1, \dots, x_n , we classify a new point x as belonging to the estimator X_n of X_ρ if the distance of $\Psi(x)$ to the linear span of $\{\Psi(x_1), \dots, \Psi(x_n)\}$ is sufficiently small.

6.3. Inverse problems and empirical risk minimization

Here we suggest a simple interpretation of the estimator F_n and stress the connection with the supervised setting. We regard the sampled data x_1, \dots, x_n as a training set of positive examples, so that each point $x_i \in X_\rho$ -almost surely; the new datum is the point $x \in X$, and we evaluate the estimator F_n at x . We label the examples according to the similarity function K by setting

$$y_i(x) = K(x_i, x) \equiv (\mathbf{K}_x)_i \quad i = 1, \dots, n.$$

If K satisfies [Assumption 1](#), then, since $K(x, x) = 1$ and K is d_K -continuous, the function y_i is close to 1 whenever x_i is close to x . The interpolation problem

$$\text{find } f \in \mathcal{H} \text{ such that } f(x_i) = y_i(x) \quad \forall i \in \{1, \dots, n\} \iff S_n f = \mathbf{K}_x$$

(where S_n is defined in [\(18\)](#)) is ill-posed. To restore well-posedness we can consider the corresponding least square problem (empirical risk minimization problem)

$$\min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n |f(x_i) - y_i(x)|^2 \iff \min_{f \in \mathcal{H}} \frac{1}{n} \|S_n f - \mathbf{K}_x\|_{\mathbb{R}^n}^2,$$

or in fact its regularized version

$$\min_{f \in \mathcal{H}} \left(\frac{1}{n} \sum_{i=1}^n |f(x_i) - y_i(x)|^2 + \lambda \|f\|^2 \right) \iff \min_{f \in \mathcal{H}} \left(\frac{1}{n} \|S_n f - \mathbf{K}_x\|_{\mathbb{R}^n}^2 + \lambda \|f\|^2 \right),$$

where $\lambda > 0$ is the regularization parameter (Tikhonov regularization). It is known [34] that the minimum of the above expression is achieved by $f \equiv f_n^\lambda$, with

$$f_n^\lambda = \frac{1}{n} g_\lambda \left(\frac{S_n^\top S_n}{n} \right) S_n^\top y,$$

where g_λ is the function $g_\lambda(\sigma) = 1/(\sigma + \lambda)$.

More generally, Tikhonov regularization can be replaced by spectral regularization induced by a different choice of the filter g_λ ; the corresponding regularized solution f_n^λ is still given by the previous equation, but the function g_λ appearing in it is now completely arbitrary. Comparing with (19), we see that $f_n^{\lambda n}(x) = F_n(x)$. Eq. (17) has then the following interpretation: a new point x is estimated to be a positive example (that is, to belong to the support X_ρ) if and only if $f_n^{\lambda n}(x) \geq 1 - \tau$, where τ is a threshold parameter.

The above discussion suggests several extensions and variations of our method, obtained considering more general penalized empirical risk minimization functionals of the form

$$\min_{f \in \mathcal{H}} \left(\frac{1}{n} \sum_{i=1}^n V(y_i(x), f(x_i)) + \lambda R(f) \right),$$

where:

- V is a (regression) loss function measuring the approximation property of f , for example the logistic loss or a robust loss such as the one used in support vector machine regression. Our theoretical analysis does not carry on to other loss functions and different mathematical concepts from empirical process theory are probably needed;
- R is a regularizer measuring the complexity of a function $f \in \mathcal{H}$. For example, one can consider the case where the kernel is given by a dictionary of atoms $f_\gamma : X \rightarrow \mathbb{R}$, with $\gamma \in \Gamma$, such that $\sum_{\gamma \in \Gamma} |f_\gamma(x)|^2 = 1$, so that we have $K(x, y) = \sum_{\gamma \in \Gamma} f_\gamma(x) f_\gamma(y)$ and, hence, $f = \sum_{\gamma \in \Gamma} w_\gamma f_\gamma$, with $w = (w_\gamma)_{\gamma \in \Gamma} \in \ell_2(\Gamma)$. In this setting, Tikhonov regularization corresponds to the choice $R(f) = \sum_{\gamma \in \Gamma} |w_\gamma|^2$, but other norms, such as the ℓ_1 norm $\sum_{\gamma \in \Gamma} |w_\gamma|$, can also be considered.

7. Empirical analysis

In this section we describe some preliminary experiments aimed at testing the properties and the performances of the proposed methods both on simulated and real data. We only discuss spectral algorithms induced by Tikhonov regularization to contrast the general method to some current state of the art algorithms. Note that while computations can be made more efficient in several ways, we consider a simple algorithmic protocol and leave a more refined computational study for future work. Recall that Tikhonov regularization defines an estimator $F_n(x) = \mathbf{K}_x^* (\mathbf{K}_n + n\lambda)^{-1} \mathbf{K}_x$, and a point x is labeled as belonging to the support X_ρ if $F_n(x) \geq 1 - \tau$. The computational cost for the algorithm is, in the worst case, of order n^3 – like standard regularized least squares – for training, and order Nn^2 if we have to predict the value of F_n at N test points. In practice, one has to choose a good value for the regularization parameter λ and this requires computing multiple solutions, a so-called *regularization path*. As noted in [57], if we form the inverse using the eigen-decomposition of the kernel matrix the price of computing the full regularization path is essentially the same as that of computing a single solution (note that the cost of the eigen-decomposition of \mathbf{K}_n is also of order n^3 , though the constant is worse). This is the strategy that we consider in the following. In our experiments we considered two datasets: the MNIST³ dataset and the CBCL⁴ face database. For the digits we considered a reduced set consisting of a training set of 5000 images and a test set of 1000 images. In the first experiment we trained on 500 images for the digit 3 and tested on 200 images of digits 3 and 8. Each experiment consists of training on one class and testing on two different classes and was repeated for 20 trials over different training set choices. For all our experiments we considered the Abel kernel. Note that in this case the algorithm requires to choose 3 parameters: the regularization parameter λ , the kernel width σ and the threshold τ . In supervised learning cross validation is typically used for parameter tuning, but cannot be used in our setting since support estimation is an unsupervised problem. Then, we considered the following heuristics. The kernel width is chosen as the median of the distribution of distances of the k -th nearest neighbor of each training set point for $k = 10$. Fixed the kernel width, we choose the regularization parameter in correspondence of the maximum curvature in the eigenvalue behavior – see Fig. 3 – the rationale being that after this value the eigenvalues are relatively small.

For comparison we considered a Parzen window density estimator and one-class SVM (1CSVM) as implemented by [11]. For the Parzen window estimator we used the same kernel of the spectral algorithm, that is the Laplacian kernel, and also the same width. Given a kernel width, an estimate of the probability distribution is computed and can be used to estimate the support by fixing a threshold τ' . For the one-class SVM we considered the Gaussian kernel, so that we have to fix the

³ <http://yann.lecun.com/exdb/mnist/>.

⁴ <http://cbcl.mit.edu/>.

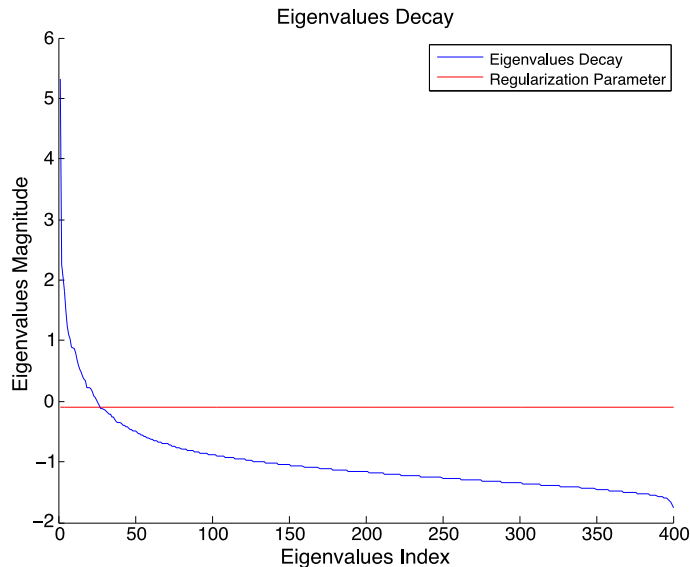


Fig. 3. Decay of the eigenvalues of the kernel matrix ordered in decreasing magnitude and corresponding regularization parameter in logarithmic scale.




Fig. 4. ROC curves for the different estimator in three different tasks: digit 9 vs 4 (left), digit 1 vs 7 (center), CBCL (right).

Table 2

Average and standard deviation of the AUC for the different estimators on the considered tasks.

	3 vs 8	8 vs 3	1 vs 7	9 vs 4	CBCL
Spectral	0.837 ± 0.006	0.783 ± 0.003	0.9921 ± 0.0005	0.865 ± 0.002	0.868 ± 0.002
Parzen	0.784 ± 0.007	0.766 ± 0.003	0.9811 ± 0.0003	0.724 ± 0.003	0.878 ± 0.002
1CSVM	0.790 ± 0.006	0.764 ± 0.003	0.9889 ± 0.0002	0.753 ± 0.004	0.882 ± 0.002

kernel width and a regularization parameter ν . We fixed the kernel width to be the same used by our estimator and set $\nu = 0.9$. For the sake of comparison, also for one-class SVM we considered a varying offset τ'' . The performance is evaluated computing ROC curve (and the corresponding AUC value) for varying values of the thresholds τ, τ', τ'' . The ROC curves on the different tasks are reported (for one of the trials) in Fig. 4. The mean and standard deviation of the AUC for the three methods is reported in Table 2. Similar experiments were repeated considering other pairs of digits, see Table 2. Also in the case of the CBCL datasets we considered a reduced dataset consisting of 472 images for training and other 472 for test. On the different test performed on the MNIST data the spectral algorithm always achieves results which are better – and often substantially better – than those of the other methods. On the CBCL dataset SVM provides the best result, but spectral algorithm still provides a competitive performance.

Remark 8. We remark that, although binary classification data sets are used in the experiments, the considered set-up is that of a one-class classification problem. Indeed, the training and tuning of the algorithms are performed using only examples of one class and the other class is only considered for testing. Accordingly, the proposed methods are compared to state of the art algorithms for one-class classification.

Acknowledgments

A.T. gratefully acknowledges the financial support of the Italian Ministry of Education, University and Research (FIRB project RBFR10COAQ). L.R. gratefully acknowledges the financial support of the Italian Ministry of Education, University and Research (FIRB project RBFR12M3AC).

Appendix A. Auxiliary proofs

In this section we give the proofs of a few technical results needed in the paper.

A.1. Normalizing a kernel

The next result shows that, if K is a reproducing kernel which is non-zero on the diagonal, then it can be normalized, and its normalized version separates the same sets. When $K(x, x) = 0$ for some $x \in X$, then clearly this result still holds replacing the set X with $X \setminus X_0$ and considering the restriction of K to $(X \setminus X_0) \times (X \setminus X_0)$, where $X_0 = \{x \in X \mid K(x, x) = 0\}$.

Proposition 13. *Assume that $K(x, x) > 0$ for all $x \in X$. Then, the reproducing kernel K' on X , given by*

$$K'(x, y) = \frac{K(x, y)}{\sqrt{K(x, x)K(y, y)}} \quad \forall x, y \in X,$$

is normalized and separates the same sets as K .

Proof. Clearly K is a kernel of positive type. Denote by \mathcal{H}' the reproducing kernel Hilbert space with kernel K' , and define the feature map $\Psi : X \rightarrow \mathcal{H}$, $\Psi(x) = K_x / \|K_x\|$. It is simple to check that $\langle \Psi(y), \Psi(x) \rangle = K'(x, y)$ and $\Psi(X)^\perp = \{0\}$, so that the map $\Psi_* : \mathcal{H} \rightarrow \mathcal{H}'$

$$(\Psi_* f)(x) = \langle f, \Psi(x) \rangle$$

is a unitary operator with $K'_x = \Psi_*(\Psi(x))$ [15]. Clearly, for any $f \in \mathcal{H}$ and $x \in X$

$$\langle \Psi_* f, K'_x \rangle = \langle \Psi_* f, \Psi_* \Psi(x) \rangle = \frac{\langle f, K_x \rangle}{\|K_x\|}.$$

The above equality shows that \mathcal{H} and \mathcal{H}' separate the same sets. \square

A.2. Analytic results

In this section, we suppose that the kernel K satisfies [Assumption 1](#), and endow the set X with the metric d_K induced by K . Measurability of a map taking values in a topological space will be always understood with respect to the Borel σ -algebra of such space. The next simple lemma will be used frequently.

Lemma 5. *For all $k = 1, 2$, the map*

$$\xi : X \rightarrow \mathcal{S}_k, \quad \xi(x) = K_x \otimes K_x$$

is continuous and measurable. Moreover, if $Z_i : \Omega \rightarrow \mathcal{S}_k$ is given by

$$Z_i(\omega) = K_{x_i} \otimes K_{x_i} \quad \omega = (x_j)_{j \geq 1},$$

then Z_i is measurable for all $i \geq 1$.

Proof. The map $x \mapsto K_x$, is continuous from X into \mathcal{H} by item i) in [Proposition 1](#). Since $\xi(x) = K_x \otimes K_x$, continuity of ξ follows at once. By item v) in [Proposition 1](#), ξ is then a measurable map, hence Z_i is such. \square

We recall some basic properties of the operator T defined by the kernel. The next result is known (see for example [\[26\]](#)), but we report a short proof for completeness.

Proposition 14. *The \mathcal{S}_1 -valued map ξ defined in [Lemma 5](#) is Bochner-integrable with respect to ρ , and its integral*

$$T = \int_X K_x \otimes K_x d\rho(x)$$

is a positive trace class operator on \mathcal{H} , with $\|T\|_{\mathcal{S}_1} = \text{tr}[T] = 1$.

Proof. The map ξ is bounded because $\|K_x \otimes K_x\|_{\mathcal{S}_1} = \text{tr}[K_x \otimes K_x] = K(x, x) = 1$ and measurable by [Lemma 5](#). Therefore, ξ is a Bochner-integrable \mathcal{S}_1 -valued map, and its integral T is a trace class operator. As $\xi(x)$ is a positive operator for all x , so is T . In particular, $\|T\|_{\mathcal{S}_1} = \text{tr}[T]$, and $\text{tr}[T] = \int_X \text{tr}[K_x \otimes K_x] d\rho(x) = 1$. \square

Now, we come to the proof of [Proposition 7](#). We will split it into the proofs of [Propositions 15 and 16](#) below.

Lemma 6. For all $k = 1, 2$, the map

$$\check{T}_n : X^n \rightarrow \mathcal{S}_k, \quad \check{T}_n(x_1, \dots, x_n) = \frac{1}{n} \sum_{i=1}^n K_{x_i} \otimes K_{x_i}$$

is continuous and measurable.

Proof. Evident by [Lemma 5](#). \square

Proposition 15. For all $n \geq 1$, the map T_n defined in [\(15\)](#) is a \mathcal{S}_k -valued estimator for $k = 1, 2$.

Proof. We have

$$T_n(\omega) = \check{T}_n(x_1, \dots, x_n) \quad \omega = (x_i)_{i \geq 1},$$

hence T_n is measurable by [Lemma 6](#). \square

For the next proposition we recall that the topology of uniform convergence on compact subsets of X is generated by the following basis of open sets $U_{f, \epsilon, C} \subset C(X)$

$$U_{f, \epsilon, C} = \left\{ g \in C(X) \mid \sup_{x \in C} |f(x) - g(x)| < \epsilon \right\} \quad f \in C(X), \epsilon > 0, C \subset X \text{ compact.}$$

Proposition 16. Suppose X is locally compact. Let $(r_\lambda)_{\lambda > 0}$ be a family of functions $r_\lambda : [0, 1] \rightarrow [0, 1]$ such that each r_λ is upper semicontinuous. Then, for any sequence of positive numbers $(\lambda_n)_{n \geq 1}$ and all $n \geq 1$, the map F_n defined in [\(16\)](#) is a $C(X)$ -valued estimator, where $C(X)$ is the space of continuous functions on X with the topology of uniform convergence on compact subsets.

Proof. Throughout the proof, $n \geq 1$ will be fixed. Let $(\varphi_k)_{k \geq 1}$ be a decreasing sequence of continuous functions $\varphi_k : [0, 1] \rightarrow [0, 1]$ such that $\varphi_k(\sigma) \downarrow r_{\lambda_n}(\sigma)$ for all $\sigma \in [0, 1]$ (such sequence exists by [\(12.7.8\)](#) of [\[31\]](#)). Then, by [Lemma 6](#) and continuity of the functional calculus (see e.g. [Problem 126](#) in [\[37\]](#)), for all $k \geq 1$ the map

$$\varphi_k(\check{T}_n) : X^n \rightarrow \mathcal{S}_0, \quad [\varphi_k(\check{T}_n)](x_1, \dots, x_n) = \varphi_k(\check{T}_n(x_1, \dots, x_n))$$

is continuous from X^n into the Banach space \mathcal{S}_0 of the bounded operators on \mathcal{H} with the uniform operator norm. Thus, for all $x \in X$, the real function $(x_1, \dots, x_n) \mapsto \langle [\varphi_k(\check{T}_n)](x_1, \dots, x_n) K_x, K_x \rangle$ is continuous on X^n , hence is measurable by item v) of [Proposition 1](#). By spectral calculus and dominated convergence theorem, for all $\omega = (x_i)_{i \geq 1}$

$$\langle r_{\lambda_n}(T_n(\omega)) K_x, K_x \rangle = \langle r_{\lambda_n}(\check{T}_n(x_1, \dots, x_n)) K_x, K_x \rangle = \lim_{k \rightarrow \infty} \langle [\varphi_k(\check{T}_n)](x_1, \dots, x_n) K_x, K_x \rangle.$$

It then follows that, for each $x \in X$, the real function $\omega \mapsto \langle r_{\lambda_n}(T_n(\omega)) K_x, K_x \rangle$ is measurable on Ω , being the pointwise limit of measurable functions.

We now prove that the map $F_n : \omega \mapsto (x \mapsto \langle r_{\lambda_n}(T_n(\omega)) K_x, K_x \rangle)$ is measurable from Ω into the space $C(X)$. By [M2](#), p. 115 in [\[45\]](#), this is equivalent to the measurability of the subsets $F_n^{-1}(U) \subset \Omega$ for all open sets $U \subset C(X)$. Since X is a locally compact separable metric space, the topology of uniform convergence on compact subsets is a separable metric topology on $C(X)$ by [\(12.14.6.2\)](#) in [\[31\]](#). By separability of $C(X)$, each open set $U \subset C(X)$ then is the denumerable union of sets of the neighborhood basis $\{U_{f, \epsilon, C} \mid f \in C(X), \epsilon > 0, C \subset X \text{ compact}\}$. Hence, it is enough to show that $F_n^{-1}(U_{f, \epsilon, C})$ is measurable for all f, ϵ and C . We have

$$F_n^{-1}(U_{f, \epsilon, C}) = \left\{ \omega \in \Omega \mid \sup_{x \in C} |f(x) - \langle r_{\lambda_n}(T_n(\omega)) K_x, K_x \rangle| < \epsilon \right\}.$$

By separability of X , there exists a countable set $C_0 \subset C$ such that $\overline{C_0} = C$. A continuity argument then shows that

$$\begin{aligned} F_n^{-1}(U_{f, \epsilon, C}) &= \bigcap_{k \geq 1} \left\{ \omega \in \Omega \mid \sup_{x \in C} |f(x) - \langle r_{\lambda_n}(T_n(\omega)) K_x, K_x \rangle| \leq \epsilon - \frac{1}{k} \right\} \\ &= \bigcap_{k \geq 1} \bigcap_{x \in C_0} \left\{ \omega \in \Omega \mid |f(x) - \langle r_{\lambda_n}(T_n(\omega)) K_x, K_x \rangle| \leq \epsilon - \frac{1}{k} \right\}. \end{aligned}$$

Since each set $\{\omega \in \Omega \mid |f(x) - \langle r_{\lambda_n}(T_n(\omega)) K_x, K_x \rangle| \leq \epsilon - 1/k\}$ is measurable in Ω , measurability of the countable intersection $F_n^{-1}(U_{f, \epsilon, C})$ then follows. \square

We conclude this section with the proof of measurability of the threshold parameters $(\tau_n)_{n \geq 1}$ defined in (30).

Proposition 17. *Suppose X is locally compact. Let $(r_\lambda)_{\lambda > 0}$ be a family of functions $r_\lambda : [0, 1] \rightarrow [0, 1]$ such that each r_λ is upper semicontinuous. Then, for any sequence of positive numbers $(\lambda_n)_{n \geq 1}$ and all $n \geq 1$, the map τ_n defined in (30) is an \mathbb{R} -valued estimator.*

Proof. As F_n depends only on (x_1, \dots, x_n) , it is clear that so does τ_n . It remains to show measurability of τ_n .

Given $i \geq 1$, the map $\omega \mapsto x_i$ is measurable by definition of the product σ -algebra \mathcal{A}_Ω on Ω . Moreover, for any $n \geq 1$, the map F_n is measurable from Ω into $C(X)$ by Proposition 16. Therefore, the map $\Theta_1 : \Omega \rightarrow C(X) \times X$, with $\Theta_1(\omega) = (F_n(\omega), x_i)$, is measurable when $C(X) \times X$ is endowed with the product σ -algebra of the Borel σ -algebras of $C(X)$ and X , respectively.

Since X is locally compact, the map $\Theta_2 : C(X) \times X \rightarrow \mathbb{R}$, with $\Theta_2(f, x) = f(x)$, is jointly continuous by [41, Theorem 5, p. 223] and the discussion following it. Thus, Θ_2 is measurable with respect to the Borel σ -algebras of $C(X) \times X$ and \mathbb{R} .

The metric spaces X and $C(X)$ are both separable (for $C(X)$, this is (12.14.6.2) in [31]). By [32, Proposition 4.1.7], the product σ -algebra of the Borel σ -algebras of $C(X)$ and X then coincides with the Borel σ -algebra of $C(X) \times X$. Thus, the composition map $\Phi_i = \Theta_2 \Theta_1$, which is $\Phi_i(\omega) = [F_n(\omega)](x_i)$, is measurable.

Finally, the map $m(t_1, \dots, t_n) \mapsto \min_{1 \leq i \leq n} t_i$ is continuous from \mathbb{R}^n into \mathbb{R} , so that $\tau_n = 1 - m(\Phi_1, \Phi_2, \dots, \Phi_n)$ is measurable. \square

A.3. A useful inequality

The following proof of inequality (A.2) below is due to A. Maurer.⁵

Lemma 7. *Suppose S and T are two symmetric Hilbert–Schmidt operators on \mathcal{H} with spectrum contained in the interval $[a, b]$, and let $(\sigma_j)_{j \in J}$ and $(\tau_k)_{k \in K}$ be the eigenvalues of S and T , respectively. Given a function $r : [a, b] \rightarrow \mathbb{R}$, if the constant*

$$L = \sup_{j \in J, k \in K} \left| \frac{r(\sigma_j) - r(\tau_k)}{\sigma_j - \tau_k} \right| \quad (\text{with } 0/0 \equiv 0)$$

is finite, then

$$\|r(S) - r(T)\|_{\mathcal{S}_2} \leq L \|S - T\|_{\mathcal{S}_2}. \quad (\text{A.1})$$

In particular, if r is a Lipschitz function with Lipschitz constant L_r , then

$$\|r(S) - r(T)\|_{\mathcal{S}_2} \leq L_r \|S - T\|_{\mathcal{S}_2}. \quad (\text{A.2})$$

Proof. Let $(f_j)_{j \in J}$ and $(g_k)_{k \in K}$ be the orthonormal bases of eigenvectors of S and T corresponding to the eigenvalues $(\sigma_j)_{j \in J}$ and $(\tau_k)_{k \in K}$, respectively, which here we list repeated accordingly to their multiplicity. We have

$$\begin{aligned} \|r(S) - r(T)\|_{\mathcal{S}_2}^2 &= \sum_{j,k} |\langle (r(S) - r(T))f_j, g_k \rangle|^2 = \sum_{j,k} (r(\sigma_j) - r(\tau_k))^2 |\langle f_j, g_k \rangle|^2 \\ &\leq L^2 \sum_{j,k} (\sigma_j - \tau_k)^2 |\langle f_j, g_k \rangle|^2 = L^2 \sum_{j,k} |\langle (S - T)f_j, g_k \rangle|^2 \\ &= L^2 \|S - T\|_{\mathcal{S}_2}^2, \end{aligned}$$

which is (A.1). \square

A.4. Concentration of measure results

We will use the following standard concentration inequality for Hilbert space random variables (see Theorem 8.6 in [53], and [54]). Let \mathcal{V} be a separable Hilbert space and $(\Omega, \mathcal{A}_\Omega, \mathbb{P})$ a probability space. Suppose that Y_1, Y_2, \dots is a sequence of independent \mathcal{V} -valued random variables $Y_i : \Omega \rightarrow \mathcal{V}$. If $\mathbb{E}[\|Y_i\|_{\mathcal{V}}^m] \leq (1/2)m!B^2L^{m-2} \forall m \geq 2$, then, for all $n \geq 1$ and $\epsilon > 0$,

⁵ <http://www.andreas-maurer.eu>.

$$\mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^n Y_i\right\|_{\mathcal{V}} > \epsilon\right) \leq 2e^{-\frac{n\epsilon^2}{B^2+L\epsilon+B\sqrt{B^2+2L\epsilon}}}. \quad (\text{A.3})$$

We will need in particular the next two straightforward consequences of this inequality.

Lemma 8. *If Z_1, Z_2, \dots is a sequence of i.i.d. \mathcal{V} -valued random variables, such that $\|Z_i\|_{\mathcal{V}} \leq M$ -almost surely, $\mathbb{E}[Z_i] = \mu$ and $\mathbb{E}[\|Z_i\|_{\mathcal{V}}^2] \leq \sigma^2$ for all i , then, for all $n \geq 1$ and $\delta > 0$,*

$$\left\|\frac{1}{n}\sum_{i=1}^n Z_i - \mu\right\|_{\mathcal{V}} \leq \frac{M\delta}{n} + \sqrt{\frac{2\sigma^2\delta}{n}} \quad (\text{A.4})$$

with probability at least $1 - 2e^{-\delta}$.

Proof. Let $Y_i = Z_i - \mu$. Then $\|Y_i\|_{\mathcal{V}} \leq 2M$ and $\mathbb{E}[\|Y_i\|_{\mathcal{V}}^2] \leq \mathbb{E}[\|Z_i\|_{\mathcal{V}}^2] = \sigma^2$. Moreover, for all i and $m \geq 2$ $\mathbb{E}[\|Y_i\|_{\mathcal{V}}^m] \leq \sigma^2(2M)^{m-2} \leq (1/2)m!\sigma^2 M^{m-2}$, where the last inequality follows since $2^{m-2} \leq m!/2$. Then,

$$\mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^n Z_i - \mu\right\|_{\mathcal{V}} > \epsilon\right) = \mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^n Y_i\right\|_{\mathcal{V}} > \epsilon\right) \leq 2e^{-\frac{n\epsilon^2}{\sigma^2+M\epsilon+\sigma\sqrt{\sigma^2+2M\epsilon}}} = 2e^{-\frac{\sigma^2 n}{M^2} g\left(\frac{M\epsilon}{\sigma^2}\right)} = 2e^{-\delta},$$

where $g(t) = t^2/(1+t+\sqrt{1+2t})$.

Since $g^{-1}(t) = t + \sqrt{2t}$, by solving the equation $(\sigma^2 n/M^2)g(M\epsilon/\sigma^2) = \delta$ we have

$$\epsilon = \frac{\sigma^2}{M} \left(\frac{M^2\delta}{n\sigma^2} + \sqrt{\frac{2M^2\delta}{n\sigma^2}} \right) = \frac{M\delta}{n} + \sqrt{\frac{2\sigma^2\delta}{n}}. \quad \square$$

The above result and Borel–Cantelli lemma imply that

$$\lim_{n \rightarrow \infty} \left\|\frac{1}{n}\sum_{i=1}^n Z_i - \mu\right\|_{\mathcal{V}} = 0$$

almost surely. In the paper we actually need a slightly stronger result which is given in the following lemma.

Lemma 9. *If Z_1, Z_2, \dots is a sequence of i.i.d. \mathcal{V} -valued random variables, such that $\|Z_i\|_{\mathcal{V}} \leq M$ -almost surely, then we have*

$$\lim_{n \rightarrow \infty} \frac{\sqrt{n}}{\log n} \left\|\frac{1}{n}\sum_{i=1}^n Z_i - \mu\right\|_{\mathcal{V}} = 0$$

almost surely.

Proof. We continue with the notations in the proof of Lemma 8. By (A.3), for all $\epsilon > 0$ we have

$$\mathbb{P}\left(\frac{\sqrt{n}}{\log n} \left\|\frac{1}{n}\sum_{i=1}^n Z_i - \mu\right\|_{\mathcal{V}} > \epsilon\right) = \mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^n Y_i\right\|_{\mathcal{V}} > \epsilon \frac{\log n}{\sqrt{n}}\right) \leq 2e^{-A(n,\epsilon)} = 2\left(\frac{1}{n}\right)^{\frac{A(n,\epsilon)}{\log n}},$$

with

$$A(n, \epsilon) = \frac{\epsilon^2 \log^2 n}{\sigma^2 + M\epsilon \frac{\log n}{\sqrt{n}} + \sigma\sqrt{\sigma^2 + 2M\epsilon \frac{\log n}{\sqrt{n}}}}.$$

It follows that

$$\sum_{n \geq 1} \mathbb{P}\left(\frac{\sqrt{n}}{\log n} \left\|\frac{1}{n}\sum_{i=1}^n Z_i - \mu\right\|_{\mathcal{V}} > \epsilon\right) \leq 2 \sum_{n \geq 1} \left(\frac{1}{n}\right)^{\frac{A(n,\epsilon)}{\log n}}.$$

For all $\epsilon > 0$, $\lim_{n \rightarrow \infty} A(n, \epsilon)/\log n = +\infty$, so that the series $\sum_{n \geq 1} n^{-A(n,\epsilon)/\log n}$ is convergent, and Borel–Cantelli lemma gives the result. \square

The following inequality is given in [13] and we report its proof for completeness.

Lemma 10. If *Assumption 1* holds true, then for all $\delta > 0$ we have

$$\|(T + \lambda)^{-1}(T - T_n)\|_{\mathcal{S}_2} \leq \left(\frac{\delta}{n\lambda} + \sqrt{\frac{2\delta\mathcal{N}(\lambda)}{n\lambda}} \right)$$

with probability at least $1 - 2e^{-\delta}$.

Proof. Let $(\Omega, \mathcal{A}_\Omega, \mathbb{P})$ be the probability space defined at the beginning of Section 5.1. For all $i \geq 1$ we define the random variable $Y_i : \Omega \rightarrow \mathcal{S}_2$ as

$$Y_i(\omega) = (T + \lambda)^{-1}(K_{x_i} \otimes K_{x_i}) \quad \omega = (x_j)_{j \geq 1},$$

which is measurable by Lemma 5. Then, we have $\|Y_i\|_{\mathcal{S}_2} \leq 1/\lambda$ -almost surely, $\mathbb{E}[Y_i] = (T + \lambda)^{-1}T$, $(1/n) \sum_{i=1}^n Y_i = (T + \lambda)^{-1}T_n$ and

$$\begin{aligned} \mathbb{E}[\|Y_i\|_{\mathcal{S}_2}^2] &= \int_{\Omega} \text{tr}[Y_i(\omega)^* Y_i(\omega)] d\mathbb{P}(\omega) = \int_{\mathcal{X}} \text{tr}[(T + \lambda)^{-2}(K_x \otimes K_x)] d\rho(x) \\ &= \text{tr}[(T + \lambda)^{-2}T] \leq \|(T + \lambda)^{-1}\|_{\infty} \text{tr}[(T + \lambda)^{-1}T] \leq \frac{\mathcal{N}(\lambda)}{\lambda}, \end{aligned}$$

where we have bounded the operator norm $\|(T + \lambda)^{-1}\|_{\infty}$ by $1/\lambda$. The result follows applying Lemma 8. \square

References

- [1] W.K. Allard, G. Chen, M. Maggioni, Multiscale geometric methods for data sets II: Geometric multi-resolution analysis, *Appl. Comput. Harmon. Anal.* 32 (2012) 435–462.
- [2] N. Aronszajn, Theory of reproducing kernels, *Trans. Amer. Math. Soc.* 68 (1950) 337–404.
- [3] F. Bauer, S. Pereverzev, L. Rosasco, On regularization algorithms in learning theory, *J. Complexity* 23 (2007) 52–72.
- [4] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation, *Neural Comput.* 15 (2003) 1373–1396.
- [5] M. Belkin, P. Niyogi, V. Sindhwani, Manifold regularization: a geometric framework for learning from labeled and unlabeled examples, *J. Mach. Learn. Res.* 7 (2006) 2399–2434.
- [6] M. Belkin, P. Niyogi, Towards a theoretical foundation for Laplacian-based manifold methods, *J. Comput. System Sci.* 74 (2008) 1289–1308.
- [7] C. Berg, J. Christensen, P. Ressel, *Harmonic Analysis on Semigroups*, Springer-Verlag, New York, 1984.
- [8] G. Biau, B. Cadre, D. Mason, B. Pelletier, Asymptotic normality in density support estimation, *Electron. J. Probab.* 14 (2009) 2617–2635.
- [9] M. Birman, M. Solomyak, Double operator integrals in a Hilbert space, *Integral Equations Operator Theory* 47 (2003) 131–168.
- [10] O. Bousquet, S. Boucheron, G. Lugosi, Theory of classification: A survey of recent advances, *ESAIM Probab. Stat.* 9 (2004) 323–375.
- [11] S. Canu, Y. Grandvalet, V. Guigue, A. Rakotomamonjy, SVM and Kernel Methods Matlab Toolbox. Perception Systèmes et Information, INSA de Rouen, Rouen, France, 2005.
- [12] A. Caponnetto, Optimal rates for regularization operators in learning theory, Technical report, MIT (CSAIL), Cambridge, MA, 2006.
- [13] A. Caponnetto, E. De Vito, Optimal rates for the regularized least-squares algorithm, *Found. Comput. Math.* 7 (2007) 331–368.
- [14] A. Caponnetto, Y. Yao, Cross-validation based adaptation for regularization operators in learning theory, *Anal. Appl.* 8 (2010) 161–183.
- [15] C. Carmeli, E. De Vito, A. Toigo, Vector valued reproducing kernel Hilbert spaces of integrable functions and Mercer theorem, *Anal. Appl.* 4 (2006) 377–408.
- [16] C. Carmeli, E. De Vito, A. Toigo, V. Umanità, Vector valued reproducing kernel Hilbert spaces and universality, *Anal. Appl.* 8 (2010) 19–61.
- [17] V. Chandola, A. Banerjee, V. Kumar, Anomaly detection: A survey, *ACM Comput. Surv.* 41 (2009) 1–58.
- [18] G. Chen, A. Little, M. Maggioni, L. Rosasco, Some recent advances in multiscale geometric analysis of point clouds, in: *Wavelets and Multiscale Analysis*, Birkhäuser/Springer, New York, 2011, pp. 199–225.
- [19] R. Coifman, S. Lafon, Geometric harmonics: a novel tool for multiscale out-of-sample extension of empirical functions, *Appl. Comput. Harmon. Anal.* 21 (2006) 31–52.
- [20] R. Coifman, S. Lafon, A. Lee, M. Maggioni, B. Nadler, S. Warner, S. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: diffusion maps, *Proc. Natl. Acad. Sci. USA* 102 (2005) 7426–7431.
- [21] R. Coifman, S. Lafon, A. Lee, M. Maggioni, B. Nadler, S. Warner, S. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: multiscale methods, *Proc. Natl. Acad. Sci. USA* 102 (2005) 7432–7438.
- [22] F. Cucker, S. Smale, On the mathematical foundations of learning, *Bull. Aust. Math. Soc.* 39 (2002) 1–49.
- [23] A. Cuevas, R. Fraiman, A plug-in approach to support estimation, *Ann. Statist.* 25 (1997) 2300–2312.
- [24] A. Cuevas, R. Fraiman, Set estimation, in: *New Perspectives in Stochastic Geometry*, Oxford Univ. Press, Oxford, 2010, pp. 374–397.
- [25] A. Cuevas, A. Rodríguez-Casal, Set estimation: an overview and some recent developments, in: *Recent Advances and Trends in Nonparametric Statistics*, Elsevier B.V., Amsterdam, 2003, pp. 251–264.
- [26] C. De Mol, E. De Vito, L. Rosasco, Elastic-net regularization in learning theory, *J. Complexity* 25 (2009) 201–230.
- [27] V. De Silva, J.B. Tenenbaum, J.C. Langford, A global geometric framework for nonlinear dimensionality reduction, *Science* 290 (2000) 2319–2323.
- [28] E. De Vito, L. Rosasco, A. Toigo, Spectral regularization for support estimation, in: *Adv. Neural Inf. Process. Syst.*, vol. 24, MIT Press, Cambridge, MA, 2010, pp. 487–495.
- [29] L. Devroye, L. Györfi, G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, Springer, New York, 1996.
- [30] L. Devroye, G.L. Wise, Detection of abnormal behavior via nonparametric estimation of the support, *SIAM J. Appl. Math.* 38 (1980) 480–488.
- [31] J. Dieudonné, *Treatise on Analysis*, vol. II, Academic Press, New York, San Francisco and London, 1976.
- [32] R.M. Dudley, *Real Analysis and Probability*, Cambridge University Press, Cambridge, 2002.
- [33] L. Dümbgen, G. Walther, Rates of convergence for random approximations of convex sets, *Adv. in Appl. Probab.* 28 (1996) 384–393.
- [34] H.W. Engl, M. Hanke, A. Neubauer, *Regularization of inverse problems*, Kluwer Academic Publishers Group, Dordrecht, 1996.
- [35] G. Folland, *A Course in Abstract Harmonic Analysis*, CRC Press, Boca Raton, FL, 1995.
- [36] L. Györfi, M. Kohler, A. Krzyzak, H. Walk, *A Distribution-Free Theory of Nonparametric Regression*, Springer-Verlag, New York, Berlin, Paris, 2002.

- [37] P. Halmos, *A Hilbert Space Problem Book*, Springer-Verlag, New York, 1982.
- [38] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning*, Springer, New York, 2001.
- [39] H. Hoffmann, Kernel PCA for novelty detection, *Pattern Recognit.* 40 (2007) 863–874.
- [40] M. Kazhdan, M. Bolitho, H. Hoppe, Poisson surface reconstruction, in: *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*, Eurographics Association, Aire-la-Ville, 2006, pp. 61–70.
- [41] J.L. Kelley, *General Topology*, reprint of the 1955 edition [Van Nostrand, Toronto, Ont.], *Grad. Texts in Math.*, vol. 27, Springer-Verlag, New York–Berlin, 1975.
- [42] R. Kolluri, J. Shewchuk, J. O'Brien, Spectral surface reconstruction from noisy point clouds, in: *Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, ACM, New York, 2004, pp. 11–21.
- [43] V. Koltchinskii, *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*, Springer, Heidelberg, 2011.
- [44] A.P. Korostel'ev, A.B. Tsybakov, *Minimax Theory of Image Reconstruction*, Springer-Verlag, New York, 1993.
- [45] S. Lang, *Real and Functional Analysis*, Springer-Verlag, New York, 1993.
- [46] G. Lerman, T. Zhang, Probabilistic recovery of multiple subspaces in point clouds by geometric ℓ_p minimization, arXiv:1002.1994v2, 2010.
- [47] D. Levin, Mesh-independent surface interpolation, in: *Geometric Modeling for Scientific Visualization*, Springer-Verlag, 2003, pp. 37–49.
- [48] L. Lo Gerfo, L. Rosasco, F. Odone, E. De Vito, A. Verri, Spectral algorithms for supervised learning, *Neural Comput.* 20 (2008) 1873–1897.
- [49] A. Maurer, M. Pontil, K -dimensional coding schemes in Hilbert spaces, *IEEE Trans. Inform. Theory* 56 (2010) 5839–5846.
- [50] J. Mercer, Functions of positive and negative type and their connection with the theory of integral equations, *Philos. Trans. R. Soc. A* 209 (1909) 415–446.
- [51] H. Narayanan, S. Mitter, Sample complexity of testing the manifold hypothesis, in: *Adv. Neural Inf. Process. Syst.*, vol. 24, MIT Press, Cambridge, MA, 2010, pp. 1786–1794.
- [52] P. Niyogi, S. Smale, S. Weinberger, A topological view of unsupervised learning from noisy data, *SIAM J. Comput.* 40 (2011) 646–663.
- [53] I. Pinelis, Optimum bounds for the distributions of martingales in Banach spaces, *Ann. Probab.* 22 (1994) 1679–1706.
- [54] I. Pinelis, Correction: optimum bounds for the distributions of martingales in Banach spaces, *Ann. Probab.* 27 (1999) 2119.
- [55] T. Poggio, S. Smale, The mathematics of learning: dealing with data, *Notices Amer. Math. Soc.* 50 (2003) 537–544.
- [56] M. Reitzner, Random polytopes and the Efron–Stein jackknife inequality, *Ann. Probab.* 31 (2003) 2136–2166.
- [57] R. Rifkin, R. Lippert, Notes on regularized least squares, Technical report, Massachusetts Institute of Technology, MIT (CSAIL), Cambridge, MA, 2007.
- [58] L. Rosasco, M. Belkin, E. De Vito, On learning with integral operators, *J. Mach. Learn. Res.* 11 (2010) 905–934.
- [59] S.T. Roweis, L.K. Saul, Nonlinear dimensionality reduction by locally linear embedding, *Science* 290 (2000) 2323–2326.
- [60] B. Schölkopf, J. Giesen, S. Spalinger, Kernel methods for implicit surface modeling, in: *Advances in Neural Information Processing Systems*, vol. 17, MIT Press, Cambridge, MA, 2005, pp. 1193–1200.
- [61] B. Schölkopf, J. Platt, J. Shawe-Taylor, A. Smola, R. Williamson, Estimating the support of a high-dimensional distribution, *Neural Comput.* 13 (2001) 1443–1471.
- [62] B. Schölkopf, A. Smola, K.-R. Müller, Nonlinear component analysis as a kernel eigenvalue problem, *Neural Comput.* 10 (1998) 1299–1319.
- [63] C.D. Scott, R.D. Nowak, Learning minimum volume sets, *J. Mach. Learn. Res.* 7 (2006) 665–704.
- [64] B. Simon, *Trace Ideals and Their Applications*, Cambridge University Press, Cambridge, 1979.
- [65] S. Smale, D.X. Zhou, Geometry of probability spaces, *Constr. Approx.* 30 (2009) 311–323.
- [66] E.M. Stein, G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*, Princeton University Press, Princeton, NJ, 1971.
- [67] I. Steinwart, On the influence of the kernel on the consistency of support vector machines, *J. Mach. Learn. Res.* 2 (2002) 67–93.
- [68] I. Steinwart, A. Christmann, *Support Vector Machines*, Springer, New York, 2008.
- [69] I. Steinwart, D. Hush, C. Scovel, A classification framework for anomaly detection, *J. Mach. Learn. Res.* 6 (2005) 211–232.
- [70] I. Steinwart, C. Scovel, Mercer's theorem on general domains: on the interaction between measures, kernels, and RKHSs, *Constr. Approx.* 35 (2012) 363–417.
- [71] V.Y. Tikhonov, A.N. Arsenin, *Solutions of Ill-Posed Problems*, Winston, New York, 1977.
- [72] A.B. Tsybakov, On nonparametric estimation of density level sets, *Ann. Statist.* 25 (1997) 948–969.
- [73] A. van der Vaart, J. Wellner, *Weak Convergence and Empirical Processes*, Springer-Verlag, New York, 1996.
- [74] V.N. Vapnik, *Statistical Learning Theory*, John Wiley & Sons Inc., New York, 1998.
- [75] R. Vert, J.-P. Vert, Consistency and convergence rates of one-class SVMs and related algorithms, *J. Mach. Learn. Res.* 7 (2006) 817–854.
- [76] R. Vidal, Y. Ma, S. Shankar, Generalized principal component analysis (GPCA), *IEEE Trans. Pattern Anal.* 27 (2005) 1945–1959.
- [77] L. Zwald, G. Blanchard, On the convergence of eigenspaces in kernel principal component analysis, in: *Adv. Neural Inf. Process. Syst.*, vol. 18, MIT Press, Cambridge, MA, 2006, pp. 1649–1656.