

Support Vector Machine For Functional Data Classification

Fabrice Rossi ^{a,*}, Nathalie Villa ^b,

^a*Projet AxIS, INRIA-Rocquencourt, Domaine de Voluceau, Rocquencourt,
B.P. 105, 78153 Le Chesnay Cedex, France*

^b*Equipe GRIMM - Université Toulouse Le Mirail, 5 allées A. Machado, 31058
Toulouse cedex 1 - FRANCE*

* Corresponding author:

Fabrice Rossi

Projet AxIS

INRIA Rocquencourt

Domaine de Voluceau, Rocquencourt, B.P. 105

78153 LE CHESNAY CEDEX – FRANCE

Tel: (33) 1 39 63 54 45

Fax: (33) 1 39 63 58 92

Email addresses: Fabrice.Rossi@inria.fr (Fabrice Rossi),
villa@univ-tlse2.fr (Nathalie Villa).

Support Vector Machine For Functional Data Classification

Abstract

In many applications, input data are sampled functions taking their values in infinite dimensional spaces rather than standard vectors. This fact has complex consequences on data analysis algorithms that motivate modifications of them. In fact most of the traditional data analysis tools for regression, classification and clustering have been adapted to functional inputs under the general name of Functional Data Analysis (FDA). In this paper, we investigate the use of Support Vector Machines (SVMs) for functional data analysis and we focus on the problem of curves discrimination. SVMs are large margin classifier tools based on implicit non linear mappings of the considered data into high dimensional spaces thanks to kernels. We show how to define simple kernels that take into account the functional nature of the data and lead to consistent classification. Experiments conducted on real world data emphasize the benefit of taking into account some functional aspects of the problems.

Key words: Functional Data Analysis, Support Vector Machine, Classification, Consistency

1 Introduction

In many real world applications, data should be considered as discretized functions rather than as standard vectors. In these applications, each observation corresponds to a mapping between some conditions (that might be implicit) and the observed response. A well studied example of those functional data is given by spectrometric data (see section 6.3): each spectrum is a function that maps the wavelengths of the illuminating light to the corresponding absorbances (the responses) of the studied sample. Other natural examples can be found in voice recognition area (see sections 6.1 and 6.2) or in meteorological problems, and more generally, in multiple time series analysis where each observation is a complete time series.

The direct use of classical models for this type of data faces several difficulties: as the inputs are discretized functions, they are generally represented by high dimensional vectors whose coordinates are highly correlated. As a consequence, classical methods lead to ill-posed problems, both on a theoretical point of view (when working in functional spaces that have infinite dimension) and on a practical one (when working with the discretized functions). The goal of Functional Data Analysis (FDA) is to use, in data analysis algorithms, the underlying functional nature of the data: many data analysis methods have been adapted to functions (see [27] for a comprehensive introduction to functional data analysis and a review of linear methods). While the original papers on FDA focused on linear methods such as Principal Component Analysis [9, 7, 8, 2] and the linear model [26, 14, 16], non linear models have been studied extensively in the recent years. This is the case, for instance, of most neural network models [12, 28, 29, 30].

In the present paper, we adapt Support Vector Machines (SVMs, see e.g. [39, 6]) to functional data classification (the paper extends results from [31, 40]). We show in particular both the practical and theoretical advantages of using functional kernels, which are kernels that take into account the functional nature of the data. On a practical point of view, those kernels allow to take advantage of the expert knowledge on the data. On the theoretical point of view, a specific type of functional kernels allows the construction of a consistent training procedure for functional SVMs.

The paper is organized as follow: section 2 presents the functional data classification and why it generally leads to ill-posed problems. Section 3 provides a short introduction to SVMs and explains why their generalization to FDA can lead to particular problems. Section 4 describes several functional kernels and explains how they can be practically computed while section 5 presents a consistency result for some of them. Finally, section 6 illustrates the various approaches presented in the paper on real data sets.

2 Functional Data Analysis

2.1 Functional Data

To simplify the presentation, this article focuses on functional data for which each observation is described by one function from \mathbb{R} to \mathbb{R} . Extension to the case of several real valued functions is straightforward. More formally, if μ denotes a finite positive Borel measure on \mathbb{R} , an observation is an element of $L^2(\mu)$, the Hilbert space of μ -square-integrable real valued functions defined on \mathbb{R} . In some situations, additional regularity assumptions (e.g., existence of derivatives) will be needed.

However, almost all the developments of this paper are not specific to functions and use only the Hilbert space structure of $L^2(\mu)$. We will therefore denote \mathcal{X} an arbitrary Hilbert space and $\langle \cdot, \cdot \rangle$ the corresponding inner product. Additional assumptions on \mathcal{X} will be given on a case by case basis. As stated above, the most common situation will of course be $\mathcal{X} = L^2(\mu)$ with $\langle u, v \rangle = \int uv d\mu$.

2.2 Data analysis methods for Hilbert spaces

It should be first noted that many data analysis algorithms can be written so as to apply, at least on a theoretical point of view, to arbitrary Hilbert spaces. This is obviously the case, for instance, for distance-based algorithms such as the k -nearest neighbor method. Indeed, this algorithm uses only the fact that distances between observations can be calculated. Obviously, it can be applied to Hilbert spaces using the distance induced by the inner product. This is also the case of methods directly based on inner products such as multi-layer perceptrons (see [32, 33, 38] for a presentation of multi-layer perceptrons with almost arbitrary input spaces, including Hilbert spaces).

However, functional spaces have infinite dimension and a basic transposition of standard algorithms introduces both theoretical and practical difficulties. In fact, some simple problems in \mathbb{R}^d become ill-posed in \mathcal{X} when the space has infinite dimension, even on a theoretical point of view.

Let us consider for instance the linear regression model in which a real valued target variable Y is modeled by $E(Y|X) = H(X)$ where H is a linear continuous operator defined on the input space. When X has values in \mathbb{R}^d (i.e., $\mathcal{X} = \mathbb{R}^d$), H can be easily estimated thanks to a least square method that leads to the inversion of the covariance matrix of X . In practice, problems might appear when d is not small compared to N , the number of available

examples, and regularization techniques should be used (e.g., ridge regression [19]). When X has values in a Hilbert space, the problem is ill-posed because the covariance operator of X is not one-to-one in general and even in particular cases where it is, it has no continuous inverse (see [4]).

To overcome the infinite dimensional problem, most of FDA methods so far have been constructed thanks to two general principles: *filtering* and *regularization*. In the filtering approach, the idea is to use representation methods that allow to work in finite dimension (see for instance [4] for the functional linear model and [3] for a functional k -nearest neighbor method). In the regularization approach, the complexity of the solution is constrained thanks to smoothness constraints. For instance, building a linear model in a Hilbert space consists in finding a function $h \in L^2(\mu)$ such that $E(Y|X) = \langle h, X \rangle$. In the regularization approach, h is chosen among smooth candidates (for instance twice derivable functions with minimal curvature), see e.g. [16, 24, 5]. Other examples of the regularization approach include smooth Principal Component Analysis [25] and penalized Canonical Component Analysis [21]. A comparison of filtering and regularization approaches for a semi-parametric model used in curve discrimination can be found in [13].

Using both approaches, a lot of data analysis algorithms have been successfully adapted to functional data. Our goal in the present paper is to study the case of Support Vector Machines (SVM), mainly thanks to a filtering approach.

3 Support Vector Machines for FDA

3.1 Support Vector Machines

We give, in this section, a very brief presentation of Support Vector Machines (SVMs) that is needed for the definition of their functional versions. We refer the reader to e.g. [6] for a more comprehensive presentation. As stated in section 2.1, \mathcal{X} denotes an arbitrary Hilbert space. Our presentation of SVM departs from the standard introduction because it assumes that the observations belong to \mathcal{X} rather than to a \mathbb{R}^d . This will make clear that the definition of SVM on arbitrary Hilbert spaces is not the difficult part in the construction of functional SVM. We will discuss problems related to the functional nature of the data in section 3.2.

Our goal is to classify data into two predefined classes. We assume given a learning set, i.e. N examples $(x_1, y_1), \dots, (x_N, y_N)$ which are i.i.d. realizations of the random variable pair (X, Y) where X has values in \mathcal{X} and Y in $\{-1, 1\}$, i.e. Y is the class label for X which is the observation.

3.1.1 Hard margin SVM

The principle of SVM is to perform an affine discrimination of the observations with maximal margin, that is to find an element $w \in \mathcal{X}$ with a minimum norm and a real value b , such that $y_i(\langle w, x_i \rangle + b) \geq 1$ for all i . To do so, we have to solve the following quadratic programming problem:

$$(P_0) \min_{w,b} \langle w, w \rangle, \text{ subject to } y_i(\langle w, x_i \rangle + b) \geq 1, 1 \leq i \leq N.$$

The classification rule associated to (w, b) is simply $f(x) = \text{sign}(\langle w, x \rangle + b)$. In this situation (called hard margin SVM), we request the rule to have zero error on the learning set.

3.1.2 Soft margin SVM

In practice, the solution provided by problem (P_0) is not very satisfactory. Firstly, perfectly linearly separable problems are quite rare, partly because non linear problems are frequent, but also because noise can turn a linearly separable problem into a non separable one. Secondly, choosing a classifier with maximal margin does not prevent overfitting, especially in very high dimensional spaces (see e.g. [17] for a discussion about this point).

A first step to solve this problem is to allow some classification errors on the learning set. This is done by replacing (P_0) by its soft margin version, i.e., by the problem:

$$(P_C) \min_{w,b,\xi} \langle w, w \rangle + C \sum_{i=1}^N \xi_i, \\ \text{subject to } y_i(\langle w, x_i \rangle + b) \geq 1 - \xi_i, 1 \leq i \leq N, \\ \xi_i \geq 0, 1 \leq i \leq N.$$

Classification errors are allowed thanks to the slack variables ξ_i . The C parameter acts as an inverse regularization parameter. When C is small, the cost of violating the hard margin constraints, i.e., the cost of having some $\xi_i > 0$ is small and therefore the constraint on w dominates. On the contrary, when C is large, classification errors dominate and (P_C) gets closer to (P_0) .

3.1.3 Non linear SVM

As noted in the previous section, some classification problems don't have a satisfactory linear solution but have a non linear one. Non linear SVMs are obtained thanks to a transformation of the original data. Assume given an Hilbert space \mathcal{H} (and denote $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ the corresponding inner product) and a function ϕ from \mathcal{X} to \mathcal{H} (this function is called a *feature map*). A linear SVM

in \mathcal{H} can be constructed on the data set $(\phi(x_1), y_1), \dots, (\phi(x_N), y_N)$. If ϕ is a non linear mapping, the classification rule $f(x) = \text{sign}(\langle w, \phi(x) \rangle_{\mathcal{H}} + b)$ is also non linear.

In order to obtain the linear SVM in \mathcal{H} one has to solve the following optimization problem:

$$\begin{aligned} (P_{C,\mathcal{H}}) \min_{w,b,\xi} & \langle w, w \rangle_{\mathcal{H}} + C \sum_{i=1}^N \xi_i, \\ \text{subject to} & \quad y_i(\langle w, \phi(x_i) \rangle_{\mathcal{H}} + b) \geq 1 - \xi_i, \quad 1 \leq i \leq N, \\ & \quad \xi_i \geq 0, \quad 1 \leq i \leq N. \end{aligned}$$

It should be noted that this feature mapping allows to define SVM on almost arbitrary input spaces.

3.1.4 Dual formulation and Kernels

Solving problems (P_C) or $(P_{C,\mathcal{H}})$ might seem very difficult at first, because \mathcal{X} and \mathcal{H} are arbitrary Hilbert spaces and can therefore have very high or even infinite dimension (when \mathcal{X} is a functional space for instance). However, each problem has a dual formulation. More precisely, (P_C) is equivalent to the following optimization problem (see [22]):

$$\begin{aligned} (D_C) \max_{\alpha} & \sum_{i=1}^N \alpha_i - \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j \langle x_i, x_j \rangle, \\ \text{subject to} & \quad \sum_{i=1}^N \alpha_i y_i = 0, \\ & \quad 0 \leq \alpha_i \leq C, \quad 1 \leq i \leq N. \end{aligned}$$

This result applies to the original problem in which data are not mapped into \mathcal{H} , but also to the mapped data, i.e., $(P_{C,\mathcal{H}})$ is equivalent to a problem $(D_{C,\mathcal{H}})$ in which the x_i are replaced by $\phi(x_i)$ and in which the inner product of \mathcal{H} is used. This leads to:

$$\begin{aligned} (D_{C,\mathcal{H}}) \max_{\alpha} & \sum_{i=1}^N \alpha_i - \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j \langle \phi(x_i), \phi(x_j) \rangle_{\mathcal{H}}, \\ \text{subject to} & \quad \sum_{i=1}^N \alpha_i y_i = 0, \\ & \quad 0 \leq \alpha_i \leq C, \quad 1 \leq i \leq N. \end{aligned}$$

Solving $(D_{C,\mathcal{H}})$ rather than $(P_{C,\mathcal{H}})$ has two advantages. The first positive aspect is that $(D_{C,\mathcal{H}})$ is an optimization problem in \mathbb{R}^N rather than in \mathcal{H} which can have infinite dimension (the same is true for \mathcal{X}).

The second important point is linked to the fact that the optimal classification rule can be written $f(x) = \text{sign}(\sum_{i=1}^N \alpha_i y_i \langle \phi(x_i), \phi(x) \rangle_{\mathcal{H}} + b)$. This means that both the optimization problem and the classification rule do not make

direct use of the transformed data, i.e. of the $\phi(x_i)$. All the calculations are done through the inner product in \mathcal{H} , more precisely thanks to the values $\langle \phi(x_i), \phi(x_j) \rangle_{\mathcal{H}}$. Therefore, rather than choosing directly \mathcal{H} and ϕ , one can provide a so called *Kernel function* K such that $K(x_i, x_j) = \langle \phi(x_i), \phi(x_j) \rangle_{\mathcal{H}}$ for a given pair (\mathcal{H}, ϕ) .

In order that K corresponds to an actual inner product in a Hilbert space, it has to fulfill some conditions. K has to be symmetric and positive definite, that is for every N , x_1, \dots, x_N in \mathcal{X} and $\alpha_1, \dots, \alpha_N$ in \mathbb{R} , $\sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j K(x_i, x_j) \geq 0$. If K satisfies those conditions, according to Moore-Aronszajn theorem [1], there exists a Hilbert space \mathcal{H} and feature map ϕ such that $K(x_i, x_j) = \langle \phi(x_i), \phi(x_j) \rangle_{\mathcal{H}}$.

3.2 The case of functional data

The short introduction to SVM proposed in the previous section has clearly shown that defining linear SVM for data in a functional space is as easy as for data in \mathbb{R}^d , because we only assumed that the input space was a Hilbert space. Thanks to the dual formulation of the optimization problem (P_C) , a software implementation of linear SVM on functional data is even possible, by relying on numerical quadrature methods to calculate the requested integrals (inner product in $L^2(\mu)$, cf section 4.3).

However, the functional nature of the data has some effects. It should be first noted that in infinite dimensional Hilbert spaces, the hard margin problem (P_0) has always a solution when the input data are in general positions, i.e., when N observations span a N dimensional subspace of \mathcal{X} . A very naive solution would therefore consists in avoiding soft margins and non linear kernels. This would not give very interesting results in practice because of the lack of regularization (see [17] for some examples in very high dimension spaces, as well as section 6.1).

Moreover, the linear SVM with soft margin can also lead to bad performances. It is indeed well known (see e.g. [18]) that problem (P_C) is equivalent to the following unconstrained optimization problem:

$$(R_\lambda) \min_{w, b} \frac{1}{N} \sum_{i=1}^N \max(0, 1 - y_i(\langle w, x_i \rangle + b)) + \lambda \langle w, w \rangle,$$

with $\lambda = \frac{1}{CN}$. This way of viewing (P_C) emphasizes the regularization aspect (see also [34, 35, 11]) and links the SVM model to ridge regression [19]. As shown in [15], the penalization used in ridge regression behaves poorly with functional data. Of course, the loss function used by SVM (the *hinge loss*, i.e., $h(u, v) = \max(0, 1 - uv)$) is different from the quadratic loss used in

ridge regression and therefore no conclusion can be drawn from experiments reported in [15]. However they show that we might expect bad performances with the linear SVM applied directly to functional data. We will see in sections 6.1 and 6.2 that the efficiency of the ridge regularization seems to be linked with the actual dimension of the data: it does not behave very well when the number of discretization points is very big and thus leads to approximate the ridge penalty by a dot product in a very high dimensional space (see also section 4.3).

It is therefore interesting to consider non linear SVM for functional data, by introducing adapted kernels. As pointed out in e.g. [11], $(P_{C,\mathcal{H}})$ is equivalent to

$$(R_{\lambda,\mathcal{H}}) \min_{f \in \mathcal{H}} \frac{1}{N} \sum_{i=1}^N \max(0, 1 - y_i f(x_i)) + \lambda \langle f, f \rangle_{\mathcal{H}}.$$

Using a kernel corresponds therefore both to replace a linear classifier by a non linear one, but also to replace the ridge penalization by a penalization induced by the kernel which might be more adapted to the problem (see [35] for links between regularization operators and kernels).

4 Kernels for FDA

4.1 Classical kernels

Many standard kernels for \mathbb{R}^d data are based on the Hilbert structure of \mathbb{R}^d and can therefore be applied to any Hilbert space. This is the case for instance of the Gaussian kernel (based on the norm in \mathcal{X}) and of the polynomial kernels (based on the inner product in \mathcal{X}). Obviously, the only practical difficulty consists in implementing the calculations needed in \mathcal{X} so as to evaluate the chosen kernel (the problem also appears for the plain linear “kernel”, i.e. when no feature mapping is done). Section 4.3 discusses this point.

4.2 Using the functional nature of the data

While the functional version of the standard kernels can provide an interesting library of kernels, they do not take advantage of the functional nature of the data (they use only the Hilbert structure of $L^2(\mu)$). Kernels that use the fact that we are dealing with functions are nevertheless quite easy to define.

A standard method consists in introducing kernels that are made by a composition of a simple feature map with a standard kernel. More formally, we use

a transformation operator P from \mathcal{X} to another space \mathcal{D} on which a kernel K is defined. The actual kernel Q on \mathcal{X} is defined as $Q(u, v) = K(P(u), P(v))$ (if K is a kernel, then so is Q).

4.2.1 Functional transformations

In some application domains, such as chemometrics, it is well known that the shape of a spectrum (which is a function) is sometimes more important than its actual mean value. Several transformations can be proposed to deal with this kind of data. For instance, if μ is a finite measure (i.e., $\mu(\mathbb{R}) < \infty$), a centering transformation can be defined as the following mapping from $L^2(\mu)$ to itself:

$$C(u) = u - \frac{1}{\mu(\mathbb{R})} \int u d\mu.$$

A normalization mapping can also be defined:

$$N(u) = \frac{1}{\|C(u)\|} C(u).$$

If the functions are smooth enough, i.e., if we restrict ourselves to a Sobolev space $W^{s,2}$, then some derivative transformations can be used: the Sobolev space $W^{s,2}$, also denoted H^s , is the Hilbert space of functions which have L^2 derivatives up to the order s (in the sense of the distribution theory). For instance, with $s \geq 2$, we can use the second derivative that allows to focus on the curvature of the functions: this is particularly useful in near infrared spectrometry (see e.g., [28, 30], and section 6.3).

4.2.2 Projections

Another type of transformations can be used in order to define adapted kernels. The idea is to reduce the dimensionality of the input space, that is to apply the standard filtering approach of FDA. We assume given a d -dimensional subspace V_d of \mathcal{X} and an orthonormal basis of this space denoted $\{\Psi_j\}_{j=1,\dots,d}$. We define the transformation P_{V_d} as the orthogonal projection on V ,

$$P_{V_d}(x) = \sum_{j=1}^d \langle x, \Psi_j \rangle \Psi_j.$$

$(V_d, \langle \cdot, \cdot \rangle_{\mathcal{X}})$ is isomorphic to $(\mathbb{R}^d, \langle \cdot, \cdot \rangle_{\mathbb{R}^d})$ and therefore one can use a standard \mathbb{R}^d SVM on the vector data $(\langle x, \Psi_1 \rangle, \dots, \langle x, \Psi_d \rangle)$. This means that K can be any kernel adapted to vector data.

Obviously, this approach is not restricted to functional data, but the choice of V_d can be directed by expert knowledge on the considered functions and

we can then consider that it takes advantage of the functional nature of the data. We outline here two possible solutions based on orthogonal basis and on B-spline basis.

If \mathcal{X} is separable, it has a Hilbert basis, i.e., a complete orthonormal system $\{\Psi_j\}_{j \geq 1}$. Therefore one can define V_d as the space spanned by $\{\Psi_j\}_{j=1, \dots, d}$. The choice of the basis can be based on expert considerations. Good candidates include Fourier basis and wavelet basis. If the signal is known to be non stationary, a wavelet based representation might for instance give better results than a Fourier representation. Once the basis is chosen, an optimal value for d can be derived from the data, as explained in 5, in such a way that the obtained SVM has some consistency properties. Moreover, this projection approach gives good results in practice (see section 6.1).

Another solution is to choose a projection space that has interesting practical properties, for instance a spline space with its associated B-spline bases. Spline functions regularity can be chosen *a priori* so as to enforce expert knowledge on the functions. For instance, near infrared spectra are smooth because of the physical properties of the light transmission (and reflection). By using a spline representation of the spectra, we replace original unconstrained observations by C^k approximations (k depends on what kind of smoothness hypothesis can be done). This projection can also be combined with a derivative transformation operation (as proposed in section 4.2.1).

4.3 Functional data in practice

In practice, the functions $(x_i)_{1 \leq i \leq N}$ are never perfectly known. It is therefore difficult to implement exactly the functional kernels described in this section.

The best situation is the one in which d discretization points have been chosen in \mathbb{R} , $(t_k)_{1 \leq k \leq d}$, and each function x_i is described by a vector of \mathbb{R}^d , $(x_i(t_1), \dots, x_i(t_d))$. In this situation, a simple solution consists in assuming that standard operations in \mathbb{R}^d (linear combinations, inner product and norm) are good approximations of their counterparts in the considered functional space. When the sampling is regular, this is equivalent to applying standard SVMs to the vector representation of the functions (see section 6 for real world examples of this situation). When the sampling is not regular, integrals should be approximated thanks to a quadrature method that will take into account the relative position of the sampling points.

In some application domains, especially medical ones (e.g., [20]), the situation is not as good. Each function is in general badly sampled: the number and the location of discretization points depend on the function and therefore a simple vector model is not anymore possible. A possible solution in this context

consists in constructing a approximation of x_i based on its observation values (thanks to e.g., B-splines) and then to work with the reconstructed functions (see [27, 30] for details).

The function approximation tool used should be simple enough to allow easy implementation of the requested operations. This is the case for instance for B-splines that allow in addition derivative calculations and an easy implementation of the kernels described in section 4.2.1. It should be noted that spline approximation is different from projection on a spline subspace. Indeed each sampled function could be approximated on a different B-spline basis, whereas the projection operator proposed in section 4.2.2 requests an unique projection space and therefore the same B-spline basis for each input function. In other words, the spline approximation is a convenient way of representing functions (see section 6.3 for an application to real world data), whereas the spline projection corresponds to a data reduction technique. Both aspects can be combined.

5 Consistency of functional SVM

5.1 Introduction

In this section we study one of the functional kernel described above and show that it can be used to define a consistent classifier for functional data. We introduce first some notations and definitions.

Our goal is to define a training procedure for functional SVM such that the asymptotic generalization performances of the constructed model is optimal. We define as usual the generalization error of a classifier f by the probability of misclassification:

$$Lf = \mathbb{P}(f(X) \neq Y).$$

The minimal generalization error is the Bayes error achieved by the optimal classifier f^* given by

$$f^*(x) = \begin{cases} 1 & \text{when } \mathbb{P}(Y = 1 \mid X = x) > 1/2 \\ -1 & \text{otherwise.} \end{cases}$$

We denote $L^* = Lf^*$ the optimal Bayes error. Of course, the closer the error of a classifier is from L^* , the better its generalization ability is.

Suppose that we are given a learning sample of size N defined as in section 3.1. A learning procedure is an algorithm which allows the construction, from

this learning sample, of a classification rule f_N chosen in a set of admissible classifiers. This algorithm is said to be consistent if

$$Lf_N \xrightarrow{N \rightarrow +\infty} L^*.$$

It should be noted that when the data belong to \mathbb{R}^d , SVMs don't always provide consistent classifiers. Some sufficient conditions have been given in [37]: the input data must belong to a compact subset of \mathbb{R}^d , the regularization parameter (C in $(P_{C,\mathcal{H}})$) for N observations must be equal to $C_N = N^{\beta-1}$, for a β which satisfies some constraints linked with the kernel, and the kernel must be *universal* [36]. If ϕ is the feature map associated to a kernel K , the kernel is universal if the set of all the functions of the form $x \mapsto \langle w, \phi(x) \rangle$ for $w \in \mathcal{H}$ is dense in the set of all continuous functions defined on the considered compact subset. In particular, the Gaussian kernel with any $\sigma > 0$ is universal for all compact subsets of \mathbb{R}^d .

5.2 A learning algorithm for functional SVM

The general methodology proposed in [3] allows to turn (with some adaptations) a consistent algorithm for data in \mathbb{R}^d into a consistent algorithm for data in \mathcal{X} , a separable Hilbert space. We describe in this section the adapted algorithm based on SVM.

The methodology proposed in [3] is based on projection operators described in section 4.2.2, more precisely on the usage of a Hilbert basis of \mathcal{X} . In order to build a SVM classifier thanks to N examples, one need to choose from the data several parameters (in addition to the weights $\{\alpha_i\}_{1 \leq i \leq N}$ and b in problem $(D_{C,\mathcal{H}})$):

- (1) the projection size parameter d , i.e., the dimension of the subset V_d on which the functions are projected before being submitted to the SVM (recall that V_d is the space spanned by $\{\Psi_j\}_{j=1,\dots,d}$);
- (2) C , the regularization parameter;
- (3) the fully specified kernel K , that is the type of the universal kernel (Gaussian, exponential, etc.) but also the parameter of this kernel such as σ for the Gaussian kernel $K(u, v) = e^{-\sigma^2 \|u-v\|^2}$.

Let us denote \mathcal{A} the set of lists of parameters to explore (see section 5.3 for practical examples). Following [3] we use a validation approach to choose the best list of parameters $a \in \mathcal{A}$ and in fact the best classifier.

The data are split into two sets: a training set $\{(x_i, y_i), i = 1, \dots, l_N\}$ and a validation set $\{(x_i, y_i), i = l_N + 1, \dots, N\}$. For each fixed list a of parameters, the training set $\{(x_i, y_i), i = 1, \dots, l_N\}$ is used to calculate the SVM classification

rule $f_a(x) = \text{sign} \left(\sum_{i=1}^{l_N} \alpha_i^* y_i K(P_{V_d}(x), P_{V_d}(x_i)) + b^* \right)$ where $(\{\alpha_i^*\}_{1 \leq i \leq l_N}, b^*)$ is the solution of $(D_{C, \mathcal{H}})$ applied to the projected data $\{P_{V_d}(x_i), i = 1, \dots, l_N\}$ (please note that everything should be indexed by a , for instance one should write K_a rather than K).

The validation set is used to select the optimal value of a in \mathcal{A} , a^* , according to estimation of the generalization error based on a penalized empirical error, that is, we define

$$a^* = \arg \min_{a \in \mathcal{A}} \hat{L}f_a + \frac{\lambda_a}{\sqrt{N - l_N}},$$

where

$$\hat{L}f_a = \frac{1}{N - l_N} \sum_{n=l_N+1}^N \mathbb{I}_{\{f_a(x_n) \neq y_n\}},$$

and λ_a is a penalty term used to avoid the selection of the most complex models (i.e., the one with the highest d in general). The classifier f_N is then chosen as $f_N = f_{a^*}$.

5.3 Consistency

Under some conditions on \mathcal{A} , the algorithm proposed in the previous section is consistent. We assume given a fixed Hilbert basis of the separable Hilbert space \mathcal{X} , $\{\Psi_j\}_{j \geq 1}$. When the dimension of the projection space V_d is chosen, a fully specified kernel K has to be chosen in a finite set of kernels, \mathcal{J}_d . The regularization parameter C can be chosen in a bounded interval of the form $[0, C_d]$, for instance thanks to the algorithm proposed in [17] that allows to calculate the validation performances for all values of C in a finite time. Therefore, the set \mathcal{A} can be written $\cup_{d \geq 1} \{d\} \times \mathcal{J}_d \times [0, C_d]$. An element of \mathcal{A} is a triple $a = (d, K, C)$ that specifies the projection operator P_{V_d} , the kernel K (including all its parameters) and the regularization constant C .

Let us first define, for all $\epsilon > 0$, $\mathcal{N}(\mathcal{H}, \epsilon)$ the covering number of the Hilbert space \mathcal{H} which is the minimum number of balls with radius ϵ that are needed to cover the whole space \mathcal{H} (see e.g., chapter 28 of [10]). Note that in SVM, as \mathcal{H} is induced by a kernel K , this number is closely related to the kernel; in this case, we will then denote the covering number $\mathcal{N}(K, \epsilon)$. For example, Gaussian kernels are known to induce feature spaces with covering number of the form $\mathcal{O}(\epsilon^{-d})$ where d is the dimension of the input space (see [37]).

Then we have:

Theorem 1 *We assume that X takes its values in a bounded subspace of the separable Hilbert space \mathcal{X} . We suppose that,*

$$\begin{aligned} \forall d \geq 1, \quad \mathcal{J}_d \text{ is a finite set,} \\ \exists K_d \in \mathcal{J}_d \text{ such that: } K_d \text{ is universal,} \\ \exists \alpha_d > 0 : \mathcal{N}(K_d, \epsilon) = \mathcal{O}(\epsilon^{-\alpha_d}), \\ \mathcal{C}_d > 1, \end{aligned}$$

and that

$$\sum_{d \geq 1} |\mathcal{J}_d| e^{-2\lambda_d^2} < +\infty,$$

and finally that

$$\begin{aligned} \lim_{N \rightarrow +\infty} l_N = +\infty & \qquad \lim_{N \rightarrow +\infty} N - l_N = +\infty \\ \lim_{N \rightarrow +\infty} \frac{l_N \log(N - l_N)}{N - l_N} = 0. \end{aligned}$$

Then, the functional SVM $f_N = f_{a^*}$ chosen as described in section 5.2 (where a^* is optimal in $\mathcal{A} = \bigcup_{d \geq 1} \{d\} \times \mathcal{J}_d \times [0, \mathcal{C}_d]$) is consistent that is:

$$Lf_N \xrightarrow{N \rightarrow +\infty} L^*.$$

The proof of this result is given in Appendix A. It is close from the proof given in [3] except that in [3] the proof follows from an oracle inequality given for a finite grid search model. The grid search is adapted to the classifier used in the paper (a k -nearest neighbor method), but not to our setting. Our result includes the search for a parameter C which can belong to an infinite and non countable set; this can be done by the use of the shatter coefficient of a particular class of linear classifiers which provides the behavior of the classification rule on a set of $N - l_N$ observations (see [10]).

As pointed out before, the Gaussian kernel satisfies the hypothesis of the theorem. Therefore, if \mathcal{I}_d contains a Gaussian kernel for all d , then consistency of the whole procedure is guaranteed. Other non universal kernels can of course be included in the search for the optimal model.

6 Applications

We present, in this section, several applications of the functional SVM models described before to real world data. The first two applications illustrate the

consistent methodology introduced in section 5.2: one has an input variable with a high number of discretization points and the second have much less discretization points. Those applications show that more benefits are obtained from the functional approach when the data can be reasonably considered as functions, that is when the number of discretization points is higher than the number of observations.

The last application deals with spectrometric data and allows to show how a functional transformation (derivative calculation) can improve the efficiency of SVMs. For this application, we do not use the consistent methodology but a projection on a spline space that permits easy derivative calculations.

6.1 *Speech recognition*

We first illustrate in this section the consistent learning procedure given in section 5. We compare it to the original procedure based on k -nn described in [3]. In practice, the only difference between the approaches is that we use a SVM whereas [3] uses a k -NN.

The problems considered in [3] consist in classifying speech samples¹. There are three problems with two classes each: classifying “yes” against “no”, “boat” against “goat” and “sh” against “ao”. For each problem, we have 100 functions. Table 1 gives the sizes of the classes for each problem.

Problem	Class 1	Class -1
yes/no	48	52
boat/goat	55	45
sh/ao	42	48

Table 1
Sizes of the classes

Each function is described by a vector in \mathbb{R}^{8192} which corresponds to a digitized speech frame. The goal of this benchmark is to compare data processing methods that make minimal assumptions on the data: no prior knowledge is used to preprocess the data.

In order to directly compare to results from [3], performances of the algorithms are assessed thanks to a leave-one-out procedure: 99 functions are used as the learning set (to which the split sample procedure is applied to choose SVM) and the remaining function provide a test example.

¹ Data are available at <http://www.math.univ-montp2.fr/~biau/bbwdata.tgz>

While the procedure described in 5.2 allows to choose most of the parameters, both the basis $\{\Psi_j\}_{j \geq 1}$ and the penalty term λ_d can be freely chosen. To focus on the improvement provided by SVM over k -nn, we have used the same elements as [3]. As the data are temporal patterns, [3] relies on the Fourier basis (moreover, the Fast Fourier Transform allows an efficient calculation of the coordinates of the data on the basis). The penalty term is 0 for all d below 100 and a high value (for instance 1000) for $d > 100$. This allows to only evaluate the models for $d \leq 100$ because the high value of λ_d for higher d prevents the corresponding models to be chosen, regardless of their performances. As pointed out in [3], this choice appears to be safe as most of the dimensions then selected are much smaller than 50.

The last free parameter is the split between the training set and the validation set. As in [3] we have used the first 50 examples for training and the remaining 49 for validation. We report the error rate for each problem and several methods in table 2. The first two columns have been reproduced from

Problem	k-nn	QDA	Gaussian SVM (projection)	linear SVM (direct)
yes/no	10%	7%	10%	58%
boat/goat	21%	35%	8%	46%
sh/ao	16%	19%	12%	47%

Table 2

Error rate for all methods (leave-one out)

[3]. QDA corresponds to Quadratic Discriminant Analysis performed, as for k -NN and SVM, on the projection of the data onto a finite dimensional subspace induced by the Fourier basis. The ‘‘Gaussian SVM’’ column corresponds to the functional kernel obtained thanks to the projection using the Fourier basis combined to a Gaussian kernel in \mathbb{R}^d . The ‘‘linear SVM’’ corresponds to the direct application of the procedure described in 3.1.2, without any prior projection. This is in fact the plain linear SVM directly applied to the original data.

The most obvious fact is that the plain linear kernel gives very poor performances, especially compared to the functional kernel on projections: its results are sometimes worse than the rule that affects any observation to the dominating class. This shows that the ridge regularization of problem (R_λ) is not adapted to functional data, a fact that was already known in the context of linear discriminant analysis [15]. As expected, the functional SVM performs generally better than k -nn and QDA, but the training times of the methods are not comparable.

6.2 Using wavelet basis

In order to investigate the limitation of the direct use of the linear SVM, we have applied them to another speech recognition problem. We studied a part of TIMIT database which was used in [15]². The data are log-periodograms corresponding to recording phonemes of 32 ms duration (the length of each log-periodogram is 256). We have chosen to restrict ourselves to classifying “aa” against “ao”, because this is the most difficult sub-problem in the database. The database is a multi-speaker database. There are 325 speakers in the training set and 112 in the test set. We have 519 examples for “aa” in the training set (759 for “ao”) and 176 in the test set (263 for “ao”). We use the split sample approach to choose the parameters on the training set (50% of the training examples are used for validation) and we report the classification error on the test set.

Here, we do not use a Fourier basis as the functions are already represented in a frequency form. As the data are very noisy, we decided to use a hierarchical wavelet basis (see e.g., [23]). We used the same penalty term as in 6.1. The error rate on the test set is reported in table 3. It appears that functional kernels are

Functional Gaussian SVM	Functional linear SVM	Linear SVM
22%	19.4%	20%

Table 3

Error rate for all methods on the test set

not as useful here as in the previous example, as a linear SVM applied directly to the discretized functions (in \mathbb{R}^{256}) performs as well as a linear SVM on the wavelet coefficients. A natural explanation is that the actual dimension of the input space (256) is smaller than the number of training examples (639) which means that evaluating the optimal coefficients of the SVM is less difficult than in the previous example. Therefore, the additional regularization provided by reducing the dimension with a projection onto a small dimensional space is not really useful in this context.

6.3 Spectrometric data set

We study in this section spectrometric data from food industry³. Each observation is the near infrared absorbance spectrum of a meat sample (finely chopped), recorded on a Tecator Infracore Food and Feed Analyser (we have

² Data are available at <http://www-stat.stanford.edu/~tibs/ElemStatLearn/datasets/phoneme.data>

³ Data are available on statlib at <http://lib.stat.cmu.edu/datasets/tecator>

215 spectra). More precisely, an observation consists in a 100 channel spectrum of absorbances in the wavelength range 850–1050 nm (see figure 1). The classification problem consists in separating meat samples with a high fat content (more than 20%) from samples with a low fat content (less than 20%).

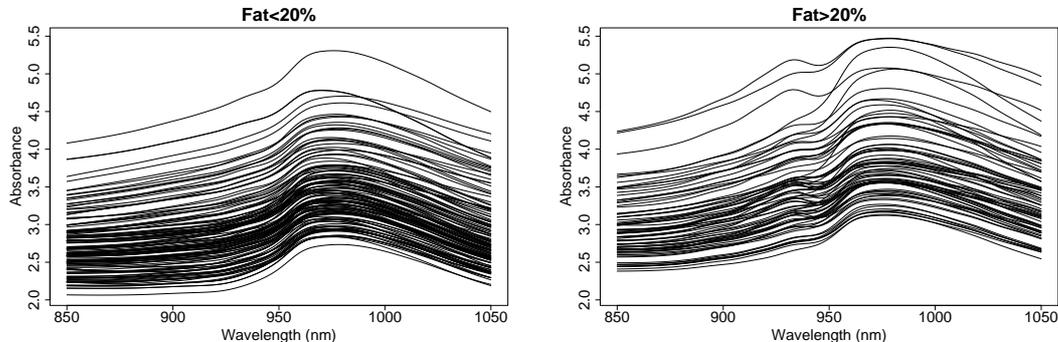


Figure 1. Spectra for both classes

It appears on figure 1 that high fat content spectra have sometimes two local maxima rather than one: we have therefore decided to focus on the curvature of the spectra, i.e., to use the second derivative. The figure 2 shows that there is more differences between the second derivatives of each class than between the original curves.

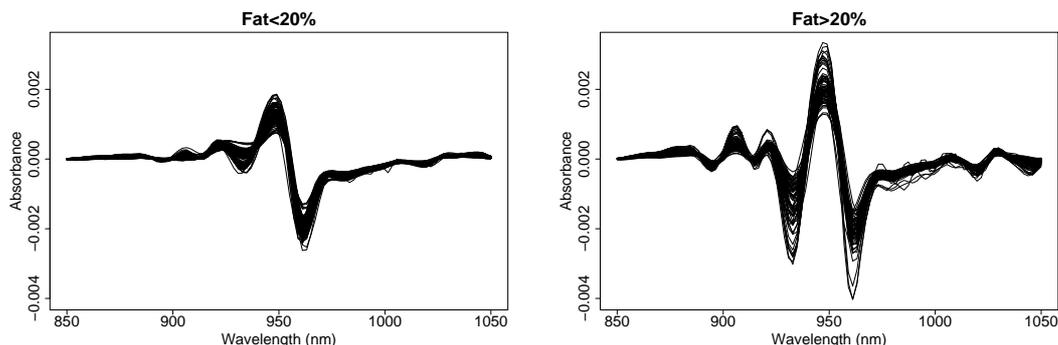


Figure 2. Second derivatives of the spectra for both classes

The data set is split into 120 spectra for learning and 95 spectra for testing. The problem is used to compare standard kernels (linear and Gaussian kernels) to a derivative based kernel. We do not use here the consistent procedure as we choose a fixed spline subspace to represent the functions so as to calculate their second derivative. However, the parameters C and σ are still chosen by a split sample approach that divides the 120 learning samples into 60 spectra for learning and 60 spectra for validation. The dimension of the spline subspace is obtained thanks to a leave-one-out procedure applied to the whole set of input functions, without taking into account classes (see [30] for details).

The performances depend of course on the random split between learning and test. We have therefore repeated this splitting 250 times (as we do not select an optimal projection dimension, the procedure is much faster than the one

used for both previous experiments). Table 4 gives the mean error rate of those experiments on the test set.

Kernel	mean test error
Linear	3.38%
Linear on second derivatives	3.28%
Gaussian	7.5%
Gaussian on second derivatives	2.6%

Table 4
Mean test error rate for all methods

The results show that the problem is less difficult than the previous ones. Nevertheless, it also appears that a functional transformation improves the results: the use of a Gaussian kernel on second derivatives gives significantly better results than the use of an usual kernel (linear or Gaussian) on the original data (*t*-test results). The relatively bad performances of the Gaussian kernel on plain data can be explained by the fact that a direct comparison of spectra based on their $L^2(\mu)$ norm is in general dominated by the mean value of those spectra which is not a good feature for classification in spectrometric problems. The linear kernel is less sensitive to this problem and is not really improved by the derivative operator. In the Gaussian case, the use of a functional transformation introduces expert knowledge (i.e., curvature is a good feature for some spectrometric problems) and allows to overcome most of the limitations of the original kernel.

7 Conclusion

In this paper, we have shown how to use Support Vector Machines (SVMs) for functional data classification. While plain linear SVMs could be used directly on functional data, we have shown the benefits of using adapted functional kernels. We have indeed defined projection based kernels that provide a consistent learning procedure for functional SVMs. We have also introduced transformation based kernels that allow to take into account expert knowledge (such as the fact that the curvature of a function can be more discriminant than its values in some applications). Both types of kernels have been tested on real world problems. The experiments gave very satisfactory results and showed that for some types of functional data, the performances of SVM based classification can be improved by using kernels that make use of the functional nature of the data.

A Proofs

In order to simplify the notations, we denote $l = l_N$ when N is obvious. We also denote $X^{(d)} = P_{V_d}(X)$ and $x_i^{(d)} = P_{V_d}(x_i)$.

The proof of the consistency result of [3] is based on an oracle. We demonstrate a similar inequality: for N large enough,

$$Lf_{a^*} - L^* \leq \inf_{d \geq 1} \left[L_d^* - L^* + \inf_{C \in \mathcal{I}_d, K \in \mathcal{J}_d} (Lf_a - L_d^*) + \frac{\lambda_d}{\sqrt{m}} \right] + \sqrt{\frac{32(l+1) \log m}{m}} + 128\Delta \sqrt{\frac{1}{32m(l+1) \log m}} \quad (\text{A.1})$$

where $m = N - l$, $\Delta \equiv \sum_{d \geq 1} |\mathcal{J}_d| e^{-\lambda_d^2/32} < +\infty$ and L_d^* is the Bayes error for the projected problem, i.e. $L_d^* = \inf_{f: \mathbb{R}^d \rightarrow \{-1,1\}} \mathbb{P}(f(X^{(d)}) \neq Y)$.

Following [3], we see that the definition of $a^* = (d^*, K^*, C^*)$ leads to,

$$\widehat{L}f_{a^*} + \frac{\lambda_{d^*}}{\sqrt{m}} \leq \widehat{L}f_a + \frac{\lambda_d}{\sqrt{m}}$$

for all $a = (d, C, K)$ in $\mathcal{A} = \cup_{d \geq 1} \{d\} \times \mathcal{J}_d \times [0, \mathcal{C}_d]$. Then, for all $\epsilon > 0$,

$$\begin{aligned} \mathbb{P} \left(Lf_{a^*} - \widehat{L}f_a > \frac{\lambda_d}{\sqrt{m}} + \epsilon \right) &\leq \mathbb{P} \left(Lf_{a^*} - \widehat{L}f_{a^*} > \frac{\lambda_{d^*}}{\sqrt{m}} + \epsilon \right) \\ &\leq \sum_{d \geq 1} \mathbb{P} \left(Lf_{(d, C^*, K^*)} - \widehat{L}f_{(d, C^*, K^*)} > \frac{\lambda_d}{\sqrt{m}} + \epsilon \right) \\ &\leq \sum_{d \geq 1, K \in \mathcal{J}_d} \mathbb{P} \left(Lf_{(d, C^*, K)} - \widehat{L}f_{(d, C^*, K)} > \frac{\lambda_d}{\sqrt{m}} + \epsilon \right) \end{aligned} \quad (\text{A.2})$$

In [3], the right part of the inequality is bounded by the use of the union bound on \mathcal{A} . Here, $[0, \mathcal{C}_d]$ is not countable and thus we can not do the same. We will then use the generalization capability of a set of linear classifiers via its shatter coefficient. Actually, when d and K are set, $f_{(d, C^*, K)}$ is an affine discrimination function built from the observation projections and the kernel K . More precisely, we have:

$$\text{for all } x \text{ in } \mathcal{X}, \quad f_a(x^{(d)}) = \sum_{n=1}^l \alpha_n^* y_n K(x_n^{(d)}, x^{(d)}) + b^*.$$

Then, f_a has the form $b + f$ where f is chosen in the set of functions spanned by $\{K(x_1^{(d)}, \cdot), \dots, K(x_l^{(d)}, \cdot)\}$. Let us denote by $\mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)})$ this set of classifiers and, for all f in $\mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)})$, we introduce $L^l f = \mathbb{P}(f(X^{(d)}) \neq$

$Y \mid (x_1, y_1), \dots, (x_l, y_l)$. By Theorem 12.6 in [10], we then have, for all $\nu > 0$,

$$\mathbb{P} \left(\sup_{f \in \mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)})} |\widehat{L}f - L^l f| > \nu \mid (x_1, y_1), \dots, (x_l, y_l) \right) \leq 8\mathcal{S}(\mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)}), m) e^{-m\nu^2/32},$$

where $\mathcal{S}(\mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)}), m)$ is the shatter coefficient of $\mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)})$, that is the maximum number of different subsets of l points that can be separated by the set of classifiers $\mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)})$. As $\mathcal{S}(\mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)}), m) \leq m^{l+1}$ we can show that, for all $(d, K) \in \mathbb{N}^* \times \mathcal{J}_d$,

$$\begin{aligned} & \mathbb{P} \left(Lf_{(d, C^*, K)} - \widehat{L}f_{(d, C^*, K)} > \frac{\lambda_d}{\sqrt{m}} + \epsilon \right) \\ &= \mathbb{E} \left[\mathbb{P} \left(Lf_{(d, C^*, K)} - \widehat{L}f_{(d, C^*, K)} > \frac{\lambda_d}{\sqrt{m}} + \epsilon \mid (x_1, y_1), \dots, (x_l, y_l) \right) \right] \\ &\leq \mathbb{E} \left[\mathbb{P} \left(\sup_{f \in \mathcal{F}_K(x_1^{(d)}, \dots, x_l^{(d)})} |\widehat{L}f - L^l f| > \frac{\lambda_d}{\sqrt{m}} + \epsilon \mid (x_1, y_1), \dots, (x_l, y_l) \right) \right] \\ &\leq 8m^{l+1} e^{-\lambda_d^2/32} e^{-m\epsilon^2/32}. \end{aligned} \tag{A.3}$$

Combining (A.2) and (A.3), we finally see that

$$\mathbb{P} \left(Lf_{a^*} - \widehat{L}f_a > \frac{\lambda_d}{\sqrt{m}} + \epsilon \right) \leq 8\Delta m^{l+1} e^{-m\epsilon^2/32}.$$

If Z is a positive random variable, we have obviously

$$\mathbb{E}(Z) \leq \mathbb{E}(Z \mathbb{I}_{\{Z>0\}}) = \int_0^{+\infty} \mathbb{P}(Z \geq \epsilon) d\epsilon.$$

For $Z = Lf_{a^*} - \widehat{L}f_a - \frac{\lambda_d}{\sqrt{m}}$, this leads, for all a in $\cup_a \{d\} \times \mathcal{I}_d \times \mathcal{J}_d$, to

$$Lf_{a^*} \leq \mathbb{E}(\widehat{L}f_a) + \frac{\lambda_d}{\sqrt{m}} + \int_0^{+\infty} \mathbb{P} \left(Lf_{a^*} - \widehat{L}f_a > \frac{\lambda_d}{\sqrt{m}} + \epsilon \right) d\epsilon.$$

Finally, following [3], for all $u > 0$,

$$\begin{aligned} \int_0^{+\infty} \mathbb{P} \left(Lf_{a^*} - \widehat{L}f_a > \frac{\lambda_d}{\sqrt{m}} + \epsilon \right) d\epsilon &\leq \int_0^u 1 d\epsilon + \int_u^{+\infty} 8\Delta m^{l+1} e^{-m\epsilon^2/32} d\epsilon \\ &\leq u + 128\Delta m^{l+1} \int_u^{+\infty} \left(\frac{1}{16} + \frac{1}{m\epsilon^2} \right) e^{-m\epsilon^2/32} d\epsilon \end{aligned}$$

and then

$$Lf_{a^*} \leq \mathbb{E}(\widehat{L}f_a) + \frac{\lambda_d}{\sqrt{m}} + u + \frac{128\Delta m^l}{u} e^{-mu^2/32},$$

if we set $u = \sqrt{\frac{32(l+1)\log m}{m}}$ and by the equality $\mathbb{E}(\widehat{L}f_a) = Lf_a$, we deduce that, for all a in \mathcal{A} ,

$$Lf_{a^*} \leq Lf_a + \frac{\lambda_d}{\sqrt{m}} + \sqrt{\frac{32(l+1)\log m}{m}} + 128\Delta\sqrt{\frac{1}{32(l+1)\log m}}$$

which finally proves oracle (A.1).

We conclude thanks to the following steps:

- (1) $\lim_{m \rightarrow +\infty} \sqrt{\frac{32(l+1)\log m}{m}} + 128\Delta\sqrt{\frac{1}{32m(l+1)\log m}} = 0$ from the assumptions of Theorem 1;
- (2) Lemma 5 in [3] shows that $L_d^* - L^* \xrightarrow{d \rightarrow +\infty} 0$;
- (3) Let $\epsilon > 0$. If we take a d_0 such that, for all $d \geq d_0$, $L_d^* - L^* \leq \epsilon$. To conclude, we finally have to prove that

$$\inf_{(C,K) \in \mathcal{I}_{d_0} \times \mathcal{J}_{d_0}} Lf_{(d_0,C,K)} - L_{d_0}^* \xrightarrow{N \rightarrow +\infty} 0.$$

Let K_{d_0} in \mathcal{J}_{d_0} be a universal kernel with covering number of the form $\mathcal{O}(\epsilon^{-\alpha_{d_0}})$ and let $C_l(d_0)$ in \mathcal{I}_{d_0} be a regularization parameter such that $C_l(d_0) = l^{\beta_{d_0}-1}$ for a $\beta_{d_0} \in]0, \frac{1}{\alpha_{d_0}}[$. By the assumption $\mathcal{C}_{d_0} > 1$, such a $C_l(d_0)$ can always be found in $[0, \mathcal{C}_{d_0}]$ for all $l \geq 1$. As X is bounded in \mathcal{X} , by definition of $x \rightarrow x^{(d)}$, $X^{(d)}$ takes its values in a bounded set of \mathbb{R}^d , that is in a compact subset of \mathbb{R}^d . From Theorem 2 in [37], the SVM with regularization parameter $C_l(d_0)$ are consistent which concludes the proof.

References

- [1] N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, May 1950.
- [2] Philippe Besse and Jim Ramsay. Principal component analysis of sampled curves. *Psychometrika*, 51:285–311, 1986.
- [3] Gérard Biau, Florentina Bunea, and Marten Wegkamp. Functional classification in hilbert spaces. *IEEE Transactions on Information Theory*, 2005. To be published.
- [4] Hervé Cardot, Frédéric Ferraty, and Pascal Sarda. Functional linear model. *Statist. & Prob. Letters*, 45:11–22, 1999.
- [5] Hervé Cardot, Frédéric Ferraty, and Pascal Sarda. Spline estimators for the functional linear model. *Statistica Sinica*, 13:571–591, 2003.
- [6] Nello Cristianini and John Shawe-Taylor. *An Introduction to Support Vector Machines*. Cambridge University Press, Cambridge, UK, 2000.

- [7] J. Dauxois and A. Pousse. *Les analyses factorielles en calcul des probabilités et en statistiques : essai d'étude synthétique*. Thèse d'état, Université Paul Sabatier, Toulouse, 1976.
- [8] J. Dauxois, A. Pousse, and Y. Romain. Asymptotic theory for the principal component analysis of a vector of random function: some applications to statistical inference. *Journal of Multivariate Analysis*, pages 136–154, 1982.
- [9] J.C. Deville. Méthodes statistiques et numériques de l'analyse harmonique. *Annales de l'INSEE*, 15:3–97, Janvier–Avril 1974.
- [10] Luc Devroye, László Györfi, and Gábor Lugosi, editors. *A Probabilistic Theory of Pattern Recognition*, volume 21 of *Applications of Mathematics*. Springer, 1996.
- [11] T. Evgeniou, M. Pontil, and T. Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13(1):1–50, 2000.
- [12] Louis Ferré and Nathalie Villa. Multi-layer neural network with functional inputs: an inverse regression approach. *Submitted to Scandinavian Journal of Statistics*, 2004.
- [13] Louis Ferré and Nathalie Villa. Discrimination de courbes par régression inverse fonctionnelle. *Revue de Statistique Appliquée*, LIII(1):39–57, 2005.
- [14] I.E. Frank and J. H. Friedman. A statistical view of some chemometrics regression tools. *Technometrics*, 35:109–148, 1993.
- [15] T. Hastie, A. Buja, and R. Tibshirani. Penalized discriminant analysis. *Annals of Statistics*, 23:73–102, 1995.
- [16] T. Hastie and C. Mallows. A discussion of "a statistical view of some chemometrics regression tools" by i.e. frank and j.h. friedman. *Technometrics*, 35:140–143, 1993.
- [17] Trevor Hastie, Saharon Rosset, Robert Tibshirani, and Ji Zhu. The entire regularization path for the support vector machine. *Journal of Machine Learning Research*, 5:1391–1415, October 2004.
- [18] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, 2001.
- [19] A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- [20] Gareth M. James and Trevor J. Hastie. Functional linear discriminant analysis of irregularly sampled curves. *Journal of the Royal Statistical Society Series B*, 63:533–550, 2001.
- [21] S. Leurgans, R. Moyeed, and B. Silverman. Canonical correlation analysis when the data are curves. *Journal of the Royal Statistical Society B*, 55(3):725–740, 1993.
- [22] Chih-Jen Lin. Formulations of support vector machines: a note from an optimization point of view. *Neural Computation*, 2(13):307–317, 2001.
- [23] Stéphane Mallat. Multiresolution approximation and wavelet orthonormal bases of l_2 . *Transaction of the American Mathematical Society*,

- 315:69–87, September 1989.
- [24] B. D. Marx and P. H. Eilers. Generalized linear regression on sampled signals with penalized likelihood. In R. Hatzinger A. Forcina, G. M. Marchetti and G. Galmacci, editors, *Statistical Modelling. Proceedings of the 11th International workshop on Statistical Modelling*, Orvieto, 1996.
 - [25] S. Pezzulli and B. Silverman. On smoothed principal components analysis. *Computational Statistics*, 8:1–16, 1993.
 - [26] J. O. Ramsay and C. J. Dalzell. Some tools for functional data analysis (with discussion). *Journal of the Royal Statistical Society Series B*, 53:539–572, 1991.
 - [27] Jim Ramsay and Bernard Silverman. *Functional Data Analysis*. Springer Series in Statistics. Springer Verlag, June 1997.
 - [28] Fabrice Rossi and Brieuc Conan-Guez. Functional multi-layer perceptron: a nonlinear tool for functional data analysis. *Neural Networks*, 18(1):45–60, January 2005.
 - [29] Fabrice Rossi, Brieuc Conan-Guez, and Aïcha El Golli. Clustering functional data with the som algorithm. In *Proceedings of ESANN 2004*, pages 305–312, Bruges, Belgium, April 2004.
 - [30] Fabrice Rossi, Nicolas Delannay, Brieuc Conan-Guez, and Michel Verleysen. Representation of functional data in neural networks. *Neurocomputing*, 64:183–210, March 2005.
 - [31] Fabrice Rossi and Nathalie Villa. Classification in hilbert spaces with support vector machines. In *Proceedings of ASMDA 2005*, pages 635–642, Brest, France, May 2005.
 - [32] Irwin W. Sandberg. Notes on weighted norms and network approximation of functionals. *IEEE Transactions on Circuits and Systems–I: Fundamental Theory and Applications*, 43(7):600–601, July 1996.
 - [33] Irwin W. Sandberg and Lilian Xu. Network approximation of input-output maps and functionals. *Circuits Systems Signal Processing*, 15(6):711–725, 1996.
 - [34] Alexander Smola and Bernhard Schölkopf. On a kernel-based method for pattern recognition, regression, approximation and operator inversion. *Algorithmica*, 22(1064):211–231, 1998.
 - [35] Alexander Smola, Bernhard Schölkopf, and Klaus-Robert Müller. The connection between regularization operators and support vector kernels. *Neural Networks*, 11:637–649, 1998.
 - [36] Ingo Steinwart. On the influence of the kernel on the consistency of support vector machines. *Journal of Machine Learning Research*, 2:67–93, November 2001.
 - [37] Ingo Steinwart. Support vector machines are universally consistent. *J. Complexity*, 18(3):768–791, September 2002.
 - [38] Maxwell B. Stinchcombe. Neural network approximation of continuous functionals and continuous functions on compactifications. *Neural Networks*, 12(3):467–477, 1999.

- [39] Vladimir Vapnik. *The Nature of Statistical Learning Theory*. Springer Verlag, New York, 1995.
- [40] Nathalie Villa and Fabrice Rossi. Support vector machine for functional data classification. In *Proceedings of ESANN 2005*, pages 467–472, Bruges, Belgium, April 2005.