

Stability of Multi-Task Kernel Regression Algorithms

Julien Audiffren

Hachem Kadri

*Aix-Marseille Université, CNRS, LIF UMR 7279,
13000, Marseille, France*

JULIEN.AUDIFFREN@LIF.UNIV-MRS.FR

HACHEM.KADRI@LIF.UNIV-MRS.FR

Editor: Cheng Soon Ong and Tu Bao Ho

Abstract

We study the stability properties of nonlinear multi-task regression in reproducing Hilbert spaces with operator-valued kernels. Such kernels, a.k.a. multi-task kernels, are appropriate for learning problems with nonscalar outputs like multi-task learning and structured output prediction. We show that multi-task kernel regression algorithms are uniformly stable in the general case of infinite-dimensional output spaces. We then derive under mild assumption on the kernel generalization bounds of such algorithms, and we show their consistency even with non Hilbert-Schmidt operator-valued kernels¹. We demonstrate how to apply the results to various multi-task kernel regression methods such as vector-valued SVR and functional ridge regression.

Keywords: Algorithmic Stability; Multi-task kernels; Hilbert space valued RKHS; Regression; Generalization bounds.

1. Introduction

A central issue in the field of machine learning is to design and analyze the generalization ability of learning algorithms. Since the seminal work of [Vapnik and Chervonenkis \(1971\)](#), various approaches and techniques have been advocated and a large body of literature has emerged in learning theory providing rigorous generalization and performance bounds ([Herbrich and Williamson, 2002](#)). This literature has mainly focused on scalar-valued function learning algorithms like binary classification ([Boucheron et al., 2005](#)) and real-valued regression ([Györfi et al., 2002](#)). However, interest in learning vector-valued functions is increasing ([Micchelli and Pontil, 2005b](#)). Much of this interest stems from the need for more sophisticated learning methods suitable for complex-output learning problems such as multi-task learning ([Caruana, 1997](#)) and structured output prediction ([Bakir et al., 2007](#)). Developing generalization bounds for vector-valued function learning algorithms then becomes more and more crucial to the theoretical understanding of such complex algorithms. Although relatively recent, the effort in this area has already produced several successful results, including ([Baxter, 2000](#); [Maurer, 2005, 2006](#); [Ando and Zhang, 2005](#)). Yet, these studies have considered only the case of finite-dimensional output spaces, and have focused more on linear machines than nonlinear ones. To the best of our knowledge, the only work investigating the generalization performance of nonlinear multi-task learning methods when output spaces can be infinite-dimensional is that of [Caponnetto and Vito \(2006\)](#). In their

1. See Definition 1 for a precise statement of what we mean by Hilbert-Schmidt operator-valued kernel.

study, the authors have derived from a theoretical (minimax) analysis generalization bounds for regularized least squares regression in reproducing kernel Hilbert spaces (RKHS) with operator-valued kernels. It should be noted that, unlike the scalar-valued function learning setting, the reproducing kernel in this context is a positive-definite operator-valued function². The operator has the advantage of allowing us to take into account dependencies between different tasks and then to model task relatedness. Hence, these kernels are known to extend linear multi-task learning methods to the nonlinear case, and are referred to as *multi-task kernels*³ (Micchelli and Pontil, 2005a; Evgeniou et al., 2005).

Evgeniou et al. (2005) refer to the issue of generalization bounds for multi-task kernel based learning as an open question. The result proposed by Caponnetto and Vito (2006) is restricted only to the least squares framework, and to our knowledge, no other work in the literature has so far addressed this issue. This paper tries to fill this gap by establishing generalization bounds for other loss functions, such as vector-valued support vector regression (SVR) and multi-task logistic regression (LR). Moreover, the convergence rates in (Caponnetto and Vito, 2006), although optimal in the case of finite-dimensional output spaces, require assumptions on the kernel that can be restrictive in the infinite-dimensional case. Indeed, their proof depends upon the fact that the kernel is Hilbert-Schmidt (see Definition 1) and this restricts the applicability of their results when the output space is infinite-dimensional. To illustrate this, let us consider the identity operator-based multi-task kernel $K(\cdot, \cdot) = k(\cdot, \cdot)I$, where k is a scalar-valued kernel and I is the identity operator. This kernel which was already used by Brouard et al. (2011) and Grunewalder et al. (2012) for structured output prediction and conditional mean embedding, respectively, does not satisfy the Hilbert-Schmidt assumption (see Remark 2), and therefore the results of (Caponnetto and Vito, 2006) cannot be applied in this case (for more details see Section 5). It is also important to note that, since the analysis of Caponnetto and Vito (2006) is based on a measure of the complexity of the hypothesis space independently of the algorithm, it does not take into account the properties of learning algorithms.

In this paper, we address these issues by studying the stability of multi-task kernel regression algorithms when the output space is a (possibly infinite-dimensional) Hilbert space. The notion of *algorithmic stability*, which is the behavior of a learning algorithm following a change of the training data, was used successfully by Bousquet and Elisseeff (2002) to derive bounds on the generalization error of deterministic scalar-valued learning algorithms. Subsequent studies extended this result to cover other learning algorithms such as randomized, transductive and ranking algorithms (Elisseeff et al., 2005; Cortes et al., 2008; Agarwal and Niyogi, 2009), both in i.i.d⁴ and non-i.i.d scenarios (Mohri and Rostamizadeh, 2010). But, none of these papers is directly concerned with the stability of nonscalar-valued learning algorithms. It is the aim of the present work to extend the stability results of (Bousquet and Elisseeff, 2002) to cover vector-valued learning schemes associated with multi-task kernels. Specifically, we make the following contributions in this paper: **1)** we show that multi-task kernel regression algorithms are uniformly stable for the general case of infinite-dimensional output spaces, **2)** we derive under mild assumption on the kernel generalization bounds of such algorithms, and we show their consistency even

2. The kernel is a matrix-valued function in the case of finite dimensional output spaces.

3. In the context of this paper, operator-valued kernels and multi-task kernels mean the same thing.

4. The abbreviation ‘‘i.i.d.’’ stands for ‘‘independently and identically distributed’’

with non Hilbert-Schmidt operator-valued kernels (see Definition 1), **3**) we demonstrate how to apply these results to various multi-task regression methods such as vector-valued support vector regression (SVR) and functional ridge regression, **4**) we provide examples of infinite-dimensional multi-task kernels which are not Hilbert-Schmidt, showing that our assumption on the kernel is weaker than the one in (Caponnetto and Vito, 2006).

The rest of this paper is organized as follows. In Section 2 we introduce the necessary notations and briefly recall the main concepts of operator-valued kernels and the corresponding Hilbert-valued RKHS. Moreover, we describe in this section the mathematical assumptions required by the subsequent developments. In Section 3 we state the result establishing the stability and providing the generalization bounds of multi-task kernel based learning algorithms. In Section 4, we show that many existing multi-task kernel regression algorithms such as vector-valued SVR and functional ridge regression do satisfy the stability requirements. In Section 5 we give examples of non Hilbert-Schmidt operator-valued kernels that illustrate the usefulness of our result. Section 6 concludes the paper.

2. Notations, Background and Assumptions

In this section we introduce the notations we will use in this paper. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, \mathcal{X} a Polish space, \mathcal{Y} a (possibly infinite-dimensional) separable Hilbert space, \mathcal{H} a separable Reproducing Kernel Hilbert Space (RKHS) $\subset \mathcal{Y}^{\mathcal{X}}$ with K its reproducing kernel, and $L(\mathcal{Y})$ ⁵ the space of continuous endomorphisms of \mathcal{Y} equipped with the operator norm. Let $\lambda > 0$ and $(X_1, Y_1), \dots, (X_m, Y_m)$ be m i.i.d. copies of the pair of random variables (X, Y) following the unknown distribution P .

We consider a training set $Z = \{(x_1, y_1), \dots, (x_m, y_m)\}$ consisting of a realization of m i.i.d. copies of (X, Y) , and we denote by $Z^i = Z \setminus (x_i, y_i)$ the set Z where the couple (x_i, y_i) is removed. Let $c : \mathcal{Y} \times \mathcal{H} \times \mathcal{X} \rightarrow \mathbb{R}^+$ be a loss function. We will describe stability and consistency results in Section 3 for a general loss function, while in Section 5 we will provide examples to illustrate them with specific forms of c . The goal of multi-task kernel regression is to find a function f , $\mathcal{H} \ni f : \mathcal{X} \rightarrow \mathcal{Y}$, that minimizes a risk functional

$$R(f) = \int c(y, f(x), x) dP(x, y).$$

The empirical risk of f on Z is then

$$R_{emp}(f, Z) = \frac{1}{m} \sum_{k=1}^m c(y_k, f, x_k),$$

and its regularized version is given by

$$R_{reg}(f, Z) = R_{emp}(f, Z) + \lambda \|f\|_{\mathcal{H}}^2.$$

We will denote by

$$f_Z = \arg \min_{f \in \mathcal{H}} R_{reg}(f, Z), \tag{1}$$

the function minimizing the regularized risk over \mathcal{H} .

5. We denote by $My = M(y)$ the application of the operator $M \in L(\mathcal{Y})$ to $y \in \mathcal{Y}$.

Let us now recall the definition of the operator-valued kernel K associated to the RKHS \mathcal{H} when \mathcal{Y} is infinite dimensional. For more details see (Micchelli and Pontil, 2005b).

Definition 1 *The application $K : \mathcal{X} \times \mathcal{X} \rightarrow L(\mathcal{Y})$ is called the Hermitian positive-definite reproducing operator-valued kernel of the RKHS \mathcal{H} if and only if :*

(i) $\forall x \in \mathcal{X}, \forall y \in \mathcal{Y}$, the application

$$\begin{aligned} K(., x)y : \mathcal{X} &\rightarrow \mathcal{Y} \\ x' &\mapsto K(x', x)y \end{aligned}$$

belongs to \mathcal{H} .

(ii) $\forall f \in \mathcal{H}, \forall x \in \mathcal{X}, \forall y \in \mathcal{Y}$,

$$\langle f(x), y \rangle_{\mathcal{Y}} = \langle f, K(., x)y \rangle_{\mathcal{H}},$$

(iii) $\forall x_1, x_2 \in \mathcal{X}$,

$$K(x_1, x_2) = K(x_2, x_1)^* \in L(\mathcal{Y}), \text{ (* denotes the adjoint)}$$

(iv) $\forall n \geq 1, \forall (x_i, i \in \{1..n\}), (x'_i, i \in \{1..n\}) \in \mathcal{X}^n, \forall (y_i, i \in \{1..n\}), (y'_i, i \in \{1..n\}) \in \mathcal{Y}^n$,

$$\sum_{k, \ell=0}^n \langle K(., x_k)y_k, K(., x'_\ell)y'_\ell \rangle_{\mathcal{Y}} \geq 0.$$

(i) and (ii) define a reproducing kernel, (iii) and (iv) correspond to the Hermitian and positive-definiteness properties, respectively.

Moreover, the kernel K will be called *Hilbert-Schmidt* if and only if $\forall x \in \mathcal{X}, \exists (y_i)_{i \in \mathbb{N}}$, a base of \mathcal{Y} , such that $\text{Tr}(K(x, x)) = \sum_{i \in \mathbb{N}} |\langle K(x, x)y_i, y_i \rangle_{\mathcal{H}}| < \infty$. This is equivalent to saying that the operator $K(x, x) \in L(\mathcal{Y})$ is Hilbert-Schmidt.

We now discuss the main assumptions we need to prove our results. We start by the following hypothesis on the kernel K .

Hypothesis 1 $\exists \kappa > 0$ such that $\forall x \in \mathcal{X}$,

$$\|K(x, x)\|_{op} \leq \kappa^2,$$

where $\|K(x, x)\|_{op} = \sup_{y \in \mathcal{Y}} \frac{\|K(x, x)y\|_{\mathcal{Y}}}{\|y\|_{\mathcal{Y}}}$ is the operator norm of $K(x, x)$ on $L(\mathcal{Y})$.

Remark 2 *It is important to note that Hypothesis 1 is weaker than the one used in (Caponnetto and Vito, 2006) which requires that K is Hilbert-Schmidt and $\sup_{x \in \mathcal{X}} \text{Tr}(K(x, x)) < +\infty$. While the two assumptions are equivalent when the output space \mathcal{Y} is finite dimensional, this is no longer the case when, as in this paper, $\dim \mathcal{Y} = +\infty$. Moreover, we can observe that if the hypothesis of (Caponnetto and Vito, 2006) is satisfied, then our Hypothesis 1 holds (see proof below). The converse is not true (see Section 5 for some counterexamples).*

PROOF OF REMARK 1 : Let K be a multi-task kernel satisfying the hypotheses of (Caponetto and Vito, 2006), i.e K is Hilbert-Schmidt and $\sup_{x \in \mathcal{X}} \text{Tr}(K(x, x)) < +\infty$. Then, $\exists \eta > 0, \forall x \in \mathcal{X}, \exists (e_j^x)_{j \in \mathbb{N}}$ an orthonormal basis of $\mathcal{Y}, \exists (h_j^x)_{j \in \mathbb{N}}$ an orthogonal family of \mathcal{H} with $\sum_{j \in \mathbb{N}} \|h_j^x\|_{\mathcal{H}}^2 \leq \eta$ such that $\forall y \in \mathcal{Y}$,

$$K(x, x)y = \sum_{j, \ell} \langle h_j^x, h_\ell^x \rangle_{\mathcal{H}} \langle y, e_j^x \rangle_{\mathcal{Y}} e_\ell^x.$$

Thus, $\forall i \in \mathbb{N}$

$$\begin{aligned} K(x, x)e_i^x &= \sum_{j, \ell} \langle h_j^x, h_\ell^x \rangle_{\mathcal{H}} \langle e_i, e_j^x \rangle_{\mathcal{Y}} e_\ell^x \\ &= \sum_{\ell} \langle h_i^x, h_\ell^x \rangle_{\mathcal{H}} e_\ell^x. \end{aligned}$$

Hence

$$\begin{aligned} \|K(x, x)\|_{op}^2 &= \sup_{i \in \mathbb{N}} \|K(x, x)e_i^x\|_{\mathcal{Y}}^2 \\ &= \sup_{i \in \mathbb{N}} \sum_{j, \ell} \langle h_i^x, h_\ell^x \rangle_{\mathcal{H}} \langle h_i^x, h_j^x \rangle_{\mathcal{H}} \langle e_j^x, e_\ell^x \rangle_{\mathcal{Y}} \\ &= \sup_{i \in \mathbb{N}} \sum_{\ell} (\langle h_i^x, h_\ell^x \rangle_{\mathcal{H}})^2 \\ &\leq \sup_{i \in \mathbb{N}} \|h_i^x\|_{\mathcal{H}}^2 \sum_{\ell} \|h_\ell^x\|_{\mathcal{H}}^2 \leq \eta^2. \end{aligned}$$

□

As a consequence of Hypothesis 1, we immediately obtain the following elementary Lemma which allows us to control $\|f(x)\|_{\mathcal{Y}}$ with $\|f\|_{\mathcal{H}}$. This is crucial to the proof of our main results.

Lemma 3 *Let K be a Hermitian positive kernel satisfying Hypothesis 1. Then $\forall f \in \mathcal{H}, \|f(x)\|_{\mathcal{Y}} \leq \kappa \|f\|_{\mathcal{H}}$.*

PROOF :

$$\begin{aligned} \|f(x)\|_{\mathcal{Y}}^2 &= \langle f(x), f(x) \rangle_{\mathcal{Y}} \\ &= \langle f, K(\cdot, x)f(x) \rangle_{\mathcal{H}} \\ &= \langle K(\cdot, x)^* f, f(x) \rangle_{\mathcal{Y}} = \langle K(x, \cdot)f, f(x) \rangle_{\mathcal{Y}} \\ &= \langle K(x, x)f, f \rangle_{\mathcal{H}} \\ &\leq \|f\|_{\mathcal{H}}^2 \|K(x, x)\|_{op} \leq \kappa^2 \|f\|_{\mathcal{H}}^2 \end{aligned}$$

□

Moreover, in order to avoid measurability problems, we assume that, $\forall y_1, y_2 \in \mathcal{Y}$, the application :

$$\begin{aligned} \mathcal{X} \times \mathcal{X} &\longrightarrow \mathbb{R} \\ (x_1, x_2) &\longrightarrow \langle K(x_1, x_2)y_1, y_1 \rangle_{\mathcal{Y}}, \end{aligned}$$

is measurable. Since \mathcal{H} is separable, this implies that all the functions used in this paper are measurable (for more details see (Caponnetto and Vito, 2006)).

A regularized multi-task kernel based learning algorithm with respect to a loss function c is the function defined by:

$$\bigcup_{n \in \mathbb{N}} (\mathcal{X} \times \mathcal{Y})^n \rightarrow \mathcal{H} \tag{2}$$

$$Z \rightarrow f_Z,$$

where f_Z is determined by equation (1). This leads us to introduce our second hypothesis.

Hypothesis 2 *The minimization problem defined by (1) is well posed. In other words, the function f_Z exists for all Z and is unique.*

Now, let us recall the notion of uniform stability of an algorithm.

Definition 4 *An algorithm $Z \rightarrow f_Z$ is said to be β uniformly stable if and only if: $\forall m \geq 1, \forall 1 \leq i \leq m, \forall Z$ a training set, and $\forall (x, y) \in \mathcal{X} \times \mathcal{Y}$ a realisation of (X, Y) Z -independent,*

$$|c(y, f_Z, x) - c(y, f_{Z^i}, x)| \leq \beta.$$

From now and for the rest of the paper, a β -stable algorithm will refer to the uniform stability. We make the following assumption regarding the loss function.

Hypothesis 3 *The application $(y, f, x) \rightarrow c(y, f, x)$ is C -admissible, i.e. convex with respect to f and Lipschitz continuous with respect to $f(x)$, with C its Lipschitz constant.*

The above three hypotheses are sufficient to prove the β -stability for a family of multi-task kernel regression algorithms. However, to show their consistency we need an additional hypothesis.

Hypothesis 4 $\exists M > 0$ such that $\forall (x, y)$ a realization of the couple (X, Y) , and $\forall Z$ a training set,

$$c(y, f_Z, x) \leq M.$$

Note that the Hypotheses 2, 3 and 4 are the same as the ones used in (Bousquet and Elisseeff, 2002). Hypothesis 1 is a direct extension of the assumption on the scalar-valued kernel to multi-task setting.

3. Stability of Multi-Task Kernel Regression

In this section, we state a result concerning the uniform stability of regularized multi-task kernel regression. This result is a direct extension of Theorem 22 in (Bousquet and Elisseeff, 2002) to the case of infinite-dimensional output spaces. It is worth pointing out that its proof does not differ much from the scalar-valued case, and requires only small modifications of the original to fit the operator-valued kernel approach. For the convenience of the reader, we present in Appendix A the proof taking into account these modifications.

Theorem 5 Under the assumptions 1, 2 and 3, the regularized multi-task kernel based learning algorithm $A : Z \rightarrow f_Z$ defined in (2) is β stable with

$$\beta = \frac{C^2 \kappa^2}{2m\lambda}.$$

PROOF : See Appendix A. □

Note that the β obtained in Theorem 5 is a $O(\frac{1}{m})$. This allows to prove the consistency of the multi-task kernel based estimator from a result of (Bousquet and Elisseeff, 2002).

Theorem 6 Let $Z \rightarrow f_Z$ be a β -stable algorithm, whose cost function c satisfies Hypothesis 4. Then, $\forall m \geq 1, \forall 0 \leq \delta \leq 1$, the following bound holds :

$$\mathbb{P} \left(\mathbb{E}(c(Y, f_Z, X)) \leq R_{emp}(f_Z, Z) + 2\beta + (4m\beta + M) \sqrt{\frac{\ln(1/\delta)}{2m}} \right) \geq 1 - \delta.$$

PROOF : See theorem 12 in (Bousquet and Elisseeff, 2002). □

Since the right term of the previous inequality tends to 0 when $m \rightarrow \infty$, theorem 6 proves the consistency of a class of multi-task kernel regression methods even when the dimensionality of the output space is infinite. We give in the next section several examples to illustrate the above results.

4. Stable Multi-Task Kernel Regression Algorithms

In this section, we show that multi-task extension of a number of existing kernel-based regression methods exhibit good uniform stability properties. In particular, we focus on functional ridge regression (RR) (Kadri et al., 2010), vector-valued support vector regression (SVR) (Brudnak, 2006), and multi-task logistic regression (LR) (Zhu and Hastie, 2002). We assume in this section that all of these algorithms satisfy Hypothesis 2.

Functional response RR. It is an extension of ridge regression (or regularized least squares regression) to functional data analysis domain (Kadri et al., 2010), where the goal is to predict a functional response by considering the output as a single function observation rather than a collection of individual observations. The operator-valued kernel RR algorithm is linked to the square loss function, and is defined as follows:

$$\arg \min_{f \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \|y - f(x)\|_Y^2 + \lambda \|f\|_{\mathcal{H}}^2.$$

We should note that Hypothesis 3 is not satisfied in the least squares context. However, we will show that the following hypothesis is a sufficient condition to prove the stability when Hypothesis 1 is verified (see Lemma 7).

Hypothesis 5 $\exists C_y > 0$ such that $\|Y\|_Y < C_y$ a.s.

Lemma 7 Let $c(y, f, x) = \|y - f(x)\|_{\mathcal{Y}}^2$. If Hypotheses 1 and 5 hold, then

$$|c(y, f_Z, x) - c(y, f_{Z^i}, x)| \leq C \|f_Z(x) - f_{Z^i}(x)\|_{\mathcal{Y}},$$

with $C = 2C_y(1 + \frac{\kappa}{\sqrt{\lambda}})$.

It is important to note that this Lemma can replace the Lipschitz property of c in the proof of Theorem 5.

PROOF OF LEMMA 7 : First, note that c is convex with respect to its second argument. Since \mathcal{H} is a vector space, $0 \in \mathcal{H}$. Thus,

$$\begin{aligned} \lambda \|f_Z\|^2 &\leq R_{reg}(f_Z, Z) \leq R_{reg}(0, Z) \\ &\leq \frac{1}{m} \sum_{k=1}^m \|y_k\|^2 \\ &\leq C_y^2, \end{aligned} \tag{3}$$

where the first line follows from the definition of f_Z (see Equation 1), and in the third line we used the bound on Y . This inequality is uniform over Z , and thus holds for f_{Z^i} .

Moreover, $\forall x \in \mathcal{X}$,

$$\begin{aligned} \|f_Z(x)\|_{\mathcal{Y}}^2 &= \langle f_Z(x), f_Z(x) \rangle_{\mathcal{Y}} \\ &= \langle K(x, x) f_Z, f_Z \rangle_{\mathcal{H}} \\ &\leq \|K(x, x)\|_{op} \|f_Z\|_{\mathcal{H}}^2 \leq \kappa^2 \frac{C_y^2}{\lambda}. \end{aligned}$$

Hence,

$$\begin{aligned} \|Y - f_Z(X)\|_{\mathcal{Y}} &\leq \|Y\|_{\mathcal{Y}} + \|f_Z(X)\|_{\mathcal{Y}} \\ &\leq C_y + \kappa \frac{C_y}{\sqrt{\lambda}}, \end{aligned} \tag{4}$$

where we used Lemma 3 and (3), then

$$\begin{aligned} &| \|y - f_Z(x)\|_{\mathcal{Y}}^2 - \|y - f_{Z^i}(x)\|_{\mathcal{Y}}^2 | \\ &= | \|y - f_Z(x)\|_{\mathcal{Y}} - \|y - f_{Z^i}(x)\|_{\mathcal{Y}} | \times | \|y - f_Z(x)\|_{\mathcal{Y}} + \|y - f_{Z^i}(x)\|_{\mathcal{Y}} | \\ &\leq 2C_y(1 + \frac{\kappa}{\sqrt{\lambda}}) \|f_Z(x) - f_{Z^i}(x)\|_{\mathcal{Y}}. \end{aligned}$$

□

Hypothesis 4 is also satisfied with $M = (C/2)^2$. We can see that from Equation 4. Using Theorem 5, we obtain that the RR algorithm is β -stable with

$$\beta = \frac{2C_y^2 \kappa^2 (1 + \frac{\kappa}{\sqrt{\lambda}})^2}{m\lambda},$$

and one can apply Theorem 6 to obtain the following generalization bound, with probability at least $1 - \delta$:

$$\begin{aligned} \mathbb{E}(c(Y, f_Z, X)) &\leq R_{emp}(f_Z, Z) + \frac{1}{m} \frac{4C_y^2 \kappa^2 (1 + \frac{\kappa}{\sqrt{\lambda}})^2}{\lambda} \\ &\quad + \sqrt{\frac{1}{m}} C_y^2 (1 + \frac{\kappa}{\sqrt{\lambda}})^2 (\frac{8\kappa^2}{\lambda} + 1) \sqrt{\frac{\ln(1/\delta)}{2}}. \end{aligned}$$

Vector-valued SVR. It was introduced in (Brudnak, 2006) to learn a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^d$ which maps inputs $x \in \mathbb{R}^n$ to vector-valued outputs $y \in \mathbb{R}^d$, where d is the number of tasks. In the paper, only the finite-dimensional output case was addressed, but a general class of loss functions associated with the p -norm of the error was studied. In the spirit of the scalar-valued SVR, the ϵ -insensitive loss function which was considered has the following form: $c(y, f, x) = \| \|y - f(x)\|_p |_\epsilon = \max(\|y - f(x)\|_p - \epsilon, 0)$, and from this general form of the p -norm formulation, the special cases of 1-, 2- and ∞ -norms was discussed. Since in our work we mainly focus our attention to the general case of any infinite-dimensional Hilbert space \mathcal{Y} , we consider here the following vector-valued SVR algorithm:

$$\arg \min_{f \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m c(y_i, f, x_i) + \lambda \|f\|_{\mathcal{H}}^2,$$

where the associated loss function is defined by:

$$c(y, f, x) = \| \|y - f(x)\|_{\mathcal{Y}} |_\epsilon = \begin{cases} 0 & \text{if } \|y - f(x)\|_{\mathcal{Y}} \leq \epsilon \\ \|y - f(x)\|_{\mathcal{Y}} - \epsilon & \text{otherwise.} \end{cases}$$

This algorithm satisfies Hypothesis 3. Hypothesis 4 is also verified with $M = C_y(1 + \frac{\kappa}{\sqrt{\lambda}})$ when Hypothesis 5 holds. This can be proved by the same way as the RR case. Theorem 5 gives that the vector-valued SVR algorithm is β -stable with

$$\beta = \frac{\kappa^2}{2m\lambda}.$$

We then obtain the following generalization bound, with probability at least $1 - \delta$:

$$\mathbb{E}(c(Y, f_Z, X)) \leq R_{emp}(f_Z, Z) + \frac{1}{m} \frac{\kappa^2}{\lambda} + \sqrt{\frac{1}{m}} (\frac{2\kappa^2}{\lambda} + C_y(1 + \frac{\kappa}{\sqrt{\lambda}})) \sqrt{\frac{\ln(1/\delta)}{2}}.$$

Multi-task LR. As in the case of SVR, kernel logistic regression (Zhu and Hastie, 2002) can be extended to the multi-task learning setting. The logistic loss can then be expanded in the manner of the ϵ -insensitive loss, that is $c(y, f, x) = \ln \left(1 + e^{-\langle y, f(x) \rangle_{\mathcal{Y}}} \right)$. It is easy to see that the multi-task LR algorithm satisfies Hypothesis 3 with $C = 1$ and Hypothesis 4 since $c(y, f, x) \leq \ln(2)$. Thus the algorithm is β -stable with

$$\beta = \frac{\kappa^2}{2m\lambda}.$$

The associated generalization bound, with probability at least $1 - \delta$, is :

$$\mathbb{E}(c(Y, f_Z, X)) \leq R_{emp}(f_Z, Z) + \frac{1}{m} \frac{\kappa^2}{\lambda} + \sqrt{\frac{1}{m} \left(\frac{2\kappa^2}{\lambda} + \ln(2) \right)} \sqrt{\frac{\ln(1/\delta)}{2}}.$$

Hence, we have obtained generalization bounds for the RR, SVR and LR algorithms even when the kernel does not satisfy the Hilbert Schmidt property (see the following section for examples of such kernels).

5. Discussion and Examples

We provide generalization bounds for multi-task kernel regression when the output space is infinite dimensional Hilbert space using the notion of algorithmic stability. As far as we are aware, the only previous study of this problem was carried out in (Caponnetto and Vito, 2006). However, only learning rates of the regularized least squares algorithm was provided when the operator-valued kernel is assumed to be Hilbert-Schmidt. We have shown in Section 3 that one may use non Hilbert-Schmidt kernels with a large family of loss functions and obtains theoretical guarantees. It should be pointed out that in the finite-dimensional case the Hilbert-Schmit assumption is always satisfied, so it is important to discuss applied machine learning situations where infinite-dimensional output spaces can be encountered. Note that our bound can be recovered from (Caponnetto and Vito, 2006) when both our and their hypotheses are satisfied

Functional regression. From a functional data analysis (FDA) point of view, infinite-dimensional output spaces for operator estimation problems are frequently encountered in functional response regression analysis, where the goal is to predict an entire function. FDA is an extension of multivariate data analysis suitable when the data are curves, see (Ramsay and Silverman, 2005) for more details. A functional response regression problem takes the form $y_i = f(x_i) + \epsilon_i$ where both predictors x_i and responses y_i are functions in some functional Hilbert space, most often the space L^2 of square integrable functions. In this context, the function f is an operator between two infinite-dimensional Hilbert spaces. Most previous work on this model suppose that the relation between functional responses and predictors is linear. The functional regression model is an extension of the multivariate linear model and has the following form:

$$y(t) = \alpha(t) + \beta(t)x(t) + \epsilon(t)$$

for a regression parameter β . In this setting, an extension to nonlinear contexts can be found in (Kadri et al., 2010) where the authors showed how Hilbert spaces of function-valued functions and infinite-dimensional operator-valued reproducing kernels can be used as a theoretical framework to develop nonlinear functional regression methods. A multiplication based operator-valued kernel was proposed, since the linear functional regression model is based on the multiplication operator.

Structured output prediction. One approach to dealing with this problem is kernel dependency estimation (KDE) (Weston et al., 2003). It is based on defining a scalar-valued kernel k_y on the outputs, such that one can transform the problem of learning a mapping

between input data x_i and structured outputs y_i to a problem of learning a Hilbert space valued function f between x_i and $\Phi(y_i)$, where $\Phi(y_i)$ is the projection of y_i by $k_{\mathcal{Y}}$ into a real-valued RKHS $\mathcal{H}_{\mathcal{Y}}$. Depending on the kernel $k_{\mathcal{Y}}$, the RKHS $\mathcal{H}_{\mathcal{Y}}$ can be infinite-dimensional. In this context, extending KDE for RKHS with multi-task kernels was first introduced in (Brouard et al., 2011), where an identity based operator-valued kernel was used to learn the function f .

Conditional mean embedding. As in the case of structured output learning, the output space in the context of conditional mean embedding is a scalar-valued RKHS. In the framework of probability distributions embedding, Grünewälder et al. (Grünewälder et al., 2012) have shown an equivalence between RKHS embeddings of conditional distributions and multi-task kernel regression. On the basis of this link, the authors derived a sparse embedding algorithm using the identity based operator-valued kernel.

Collaborative filtering. The goal of collaborative filtering (CF) is to build a model to predict preferences of clients “users” over a range of products “items” based on information from customer’s past purchases. In (Abernethy et al., 2009), the authors show that several CF methods such as rank-constrained optimization, trace-norm regularization, and those based on Frobenius norm regularization, can all be cast as special cases of spectral regularization on operator spaces. Using operator estimation and spectral regularization as a framework for CF permit to use potentially more information and incorporate additional *user-item* attributes to predict preferences. A generalized CF approach consists in learning a preference function $f(\cdot, \cdot)$ that takes the form of a linear operator from a Hilbert space of users to a Hilbert space of items, $f(\cdot, \cdot) = \langle x, Fy \rangle$ for some compact operator F .

Now we want to emphasize that in the case of infinite dimensions Hypothesis 1 on the kernel is not equivalent to that used in (Caponnetto and Vito, 2006). We have shown in Section 2 that our assumption on the kernel is weaker. To illustrate this, we provide below examples of operator-valued kernels which satisfy Hypothesis 1 but are not Hilbert-Schmidt, as was assumed in (Caponnetto and Vito, 2006).

Example 1 Identity operator. Let $c > 0$, $d \in \mathbb{N}^*$, $\mathcal{X} = \mathbb{R}^d$, I be the identity morphism in $L(\mathcal{Y})$, and $K(x, t) = \exp(-c\|x - t\|_2^2) \times I$. The kernel K is positive, Hermitian, and

$$\begin{aligned} \|K(x, x)\|_{op} &= \|I\|_{op} = 1 \\ \text{Tr}(K(x, x)) &= \text{Tr}(I) = +\infty. \end{aligned} \tag{5}$$

This result is also true for any kernel which can be written as $K(x, t) = k(x, t)I$, where k is a positive-definite scalar-valued kernel satisfying that $\sup_{x \in \mathcal{X}} k(x, x) < \infty$. Identity based operator-valued kernels are already used in structured output learning (Brouard et al., 2011) and conditional mean embedding (Grünewälder et al., 2012).

Example 2 Multiplication operator - Separable case. Let k be a positive-definite scalar-valued such that $\sup_{x \in \mathcal{X}} k(x, x) < \infty$, \mathcal{I} an interval of \mathbb{R} , $C > 0$, and $\mathcal{Y} = L^2(\mathcal{I}, \mathbb{R})$. Let $f \in L^\infty(\mathcal{I}, \mathbb{R})$ be such that $\|f\|_\infty < C$.

We now define the multiplication based operator-valued kernel K as follows

$$K(x, z)y(\cdot) = k(x, z)f^2(\cdot)y(\cdot) \in \mathcal{Y}.$$

Such kernels are suited to extend linear functional regression to nonlinear context (Kadri et al., 2010). K is a positive Hermitian kernel but, even if K always satisfy Hypothesis 1, the Hilbert-Schmidt property of K depends on the choice of f and may be difficult to verify. For instance, let $f(t) = \frac{C}{2}(\exp(-t^2) + 1)$. Then

$$\begin{aligned} \|K(x, x)\|_{op} &\leq C^2 k(x, x) \\ \text{Tr}(K(x, x)) &= \sum_{i \in \mathbb{N}} \langle K(x, x)y_i, y_i \rangle \geq k(x, x) \frac{C}{2} \sum_{i \in \mathbb{N}} \|y_i\|_2^2 = \infty, \end{aligned} \quad (6)$$

where $(y_i)_{i \in \mathbb{N}}$ is an orthonormal basis of Y (which exists, since Y is separable).

Example 3 *Multiplication Operator - Non separable case*⁶. Let \mathcal{I} an interval of \mathbb{R} , $C > 0$, $\mathcal{Y} = L^2(\mathcal{I}, \mathbb{R})$, $\mathcal{X} = \{f \in L^\infty(\mathcal{I}, \mathbb{R}) \text{ such that } \|f\|_\infty \leq C\}$.

Let K be the following operator-valued function:

$$K(x, z)y(\cdot) = x(\cdot)z(\cdot)y(\cdot) \in \mathcal{Y}.$$

K is a positive Hermitian kernel satisfying Hypothesis 1. Indeed,

$$\|K(x, x)\|_{op} = \max_{y \in \mathcal{Y}} \frac{\sqrt{\int_{\mathcal{I}} x^4(t)y^2(t)dt}}{\|y\|_2} \leq C^2 \max_{y \in \mathcal{Y}} \frac{\sqrt{\int_{\mathcal{I}} y^2(t)dt}}{\|y\|_2} \leq C^2.$$

On the other hand, $K(x, x)$ is not Hilbert-Schmidt for all choice of x (in fact it is not Hilbert Schmidt as long as $\exists \varepsilon > 0$ such that $\forall t \in \mathcal{I}$, $x(t) > \varepsilon > 0$). To illustrate this, let choose $x = \frac{C}{2}(\exp(-t^2) + 1)$ as defined in the previous example. Then, for $(y_i)_{i \in \mathbb{N}}$ an orthonormal basis of Y , we have

$$\begin{aligned} \text{Tr}(K(x, x)) &= \sum_{i \in \mathbb{N}} \langle K(x, x)y_i, y_i \rangle = \sum_{i \in \mathbb{N}} \int_{\mathcal{I}} x^2(t)y_i^2(t)dt \\ &\geq \frac{C}{2} \sum_{i \in \mathbb{N}} \int_{\mathcal{I}} y_i^2(t)dt = \frac{C}{2} \sum_{i \in \mathbb{N}} \|y_i\|_2^2 = +\infty. \end{aligned}$$

Example 4 *Sum of kernels*. This example is provided to show that in the case of multiple kernels the sum of a non Hilbert-Schmidt kernel with an Hilbert-Schmidt one gives a non-Hilbert-Schmidt kernel. This makes the assumption on the kernel of (Caponnetto and Vito, 2006) inconvenient for multiple kernel learning (MKL) (Kadri et al., 2012), since one would like to learn a combination of different kernels which can be non Hilbert-Schmidt (like the basic identity based operator-valued kernel).

Let k be a positive-definite scalar-valued kernel satisfying $\sup_{x \in \mathcal{X}} k(x, x) < \infty$, \mathcal{Y} a Hilbert space, and $y_0 \in \mathcal{Y}$. Let K be the following kernel:

$$K(x, z)y = k(x, z)(y + \langle y, y_0 \rangle y_0).$$

K is a positive and Hermitian kernel. Note that a similar kernel was proposed for multi-task learning (Kadri et al., 2012), where the identity operator is used to encode the relation

6. A kernel $K(x, t)$ is called non separable, as opposed to separable, if it cannot be written as the product of a scalar valued kernel $k(x, t)$ and a $L(\mathcal{Y})$ operator independent of the choice of $x, t \in \mathcal{X}$.

between a task and itself, and a second kernel is added for sharing the information between tasks. K satisfies Hypothesis 1, since

$$\|K(x, x)\|_{op} = \max_{y \in \mathcal{Y}, \|y\|_{\mathcal{Y}}=1} \|k(x, x) (y + \langle y, y_0 \rangle y_0)\|_{\mathcal{Y}} \leq |k(x, x)|(1 + \|y_0\|_{\mathcal{Y}}^2).$$

However, K is not Hilbert Schmidt. Indeed, it is the sum of a Hilbert Schmidt kernel (resp. $K_1(x, z)y = k(x, z) \langle y, y_0 \rangle y_0$) and a Hilbert-Schmidt one (resp. $K_2(x, z)y = k(x, z)y$), which is not Hilbert Schmidt. To see this, note that since the trace of K_1 is the sum over an absolutely summable family, and the trace is linear, so the trace of K is the sum of an convergent series and a divergent one, hence it diverges, so K is not Hilbert Schmidt.

6. Conclusion

We have shown that a large family of multi-task kernel regression algorithms, including functional ridge regression and vector-valued SVR, are β -stable even when the output space is infinite-dimensional. This result allows us to provide generalization bounds and to prove under mild assumptions on the kernel the consistency of these algorithms. In the future, it will be interesting to derive learning bounds with optimal and fast rates for infinite-dimensional multi-task kernel based algorithms.

Acknowledgments

This research was funded by the Ministry of Higher Education and Research, the ANR project LAMPADA (ANR-09-EMER-007), and the CNRS PEPS project FLAME.

References

- Jacob Abernethy, Francis Bach, Theodoros Engeniou, and Jean-Philippe Vert. A New Approach to Collaborative Filtering : Operator Estimation with Spectral Regularization. *Journal of Machine Learning Research*, 10:803–826, 2009.
- Shivani Agarwal and Partha Niyogi. Generalization bounds for ranking algorithms via algorithmic stability. *Journal of Machine Learning Research*, 10:441–474, 2009.
- Rie Kubota Ando and Tong Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6:1817–1853, 2005.
- Gökhan Bakir, Thomas Hofmann, Bernhard Schölkopf, Alexander J. Smola, Ben Taskar, and S.V.N. Vishwanathan. *Predicting Structured Data*. The MIT Press, 2007.
- Jonathan Baxter. A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 12:149–198, 2000.
- Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: a survey of some recent advances. *ESAIM: Probability and Statistics*, 9:323 – 375, 2005.

- Olivier Bousquet and Andre Elisseeff. Stability and generalisation. *Journal of Machine Learning Research*, 2:499–526, 2002.
- Céline Brouard, Florence D’Alche-Buc, and Marie Szafranski. Semi-supervised penalized output kernel regression for link prediction. ICML ’11, June 2011.
- Mark Brudnak. Vector-valued support vector regression. In *IJCNN*, pages 1562–1569. IEEE, 2006.
- Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least square algorithm. *Foundations of the computational mathematics*, 7:361–368, 2006.
- Rich Caruana. *Multitask Learning*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1997.
- C. Cortes, M. Mohri, D. Pechyony, and A. Rastogi. Stability of transductive regression algorithms. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, pages 176–183, 2008.
- André Elisseeff, Theodoros Evgeniou, and Massimiliano Pontil. Stability of randomized learning algorithms. *Journal of Machine Learning Research*, 6:55–79, 2005.
- Theodoros Evgeniou, Charles A. Micchelli, and Massimiliano Pontil. Learning multiple tasks with kernel methods. *Journal of Machine Learning Research*, 6:615–637, 2005.
- Steffen Grunewalder, Guy Lever, Luca Baldassarre, Sam Patterson, Arthur Gretton, and Massi Pontil. Conditional mean embeddings as regressors. ICML ’12, July 2012.
- László Györfi, Michael Kohler, Adam Krzyżak, and Harro Walk. *A distribution-free theory of nonparametric regression*. Springer Series in Statistics. Springer-Verlag, New York, 2002.
- Ralf Herbrich and Robert C Williamson. Learning and generalization: Theoretical bounds. In *Handbook of Brain Theory and Neural Networks*, pages 619–623. 2nd edition edition, 2002.
- Hachem Kadri, Emmanuel Duflos, Philippe Preux, Stéphane Canu, and Manuel Davy. Non-linear functional regression: a functional rkhs approach. AISTATS ’10, May 2010.
- Hachem Kadri, Alain Rakotomamonjy, Francis Bach, and Philippe Preux. Multiple Operator-valued Kernel Learning. In *Neural Information Processing Systems (NIPS)*, 2012.
- Andreas Maurer. Algorithmic stability and meta-learning. *Journal of Machine Learning Research*, 6:967–994, 2005.
- Andreas Maurer. Bounds for linear multi-task learning. *Journal of Machine Learning Research*, 7:117–139, 2006.
- C. Micchelli and M. Pontil. Kernels for multi-task learning. In *Advances in Neural Information Processing Systems 17*, pages 921–928, 2005a.

Charles A. Micchelli and Massimiliano Pontil. On learning vector-valued functions. *Neural Computation*, 17:177–204, 2005b.

Mehryar Mohri and Afshin Rostamizadeh. Stability bounds for stationary phi-mixing and beta-mixing processes. *Journal of Machine Learning Research*, 11:789–814, 2010.

Jim Ramsay and Bernard Silverman. *Functional Data Analysis*. Springer Series in Statistics. Springer Verlag, 2005.

Vladimir N. Vapnik and Alexey Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16(2):264–280, 1971.

Jason Weston, Olivier Chapelle, Andre Elisseeff, Bernhard Schölkopf, and Vladimir Vapnik. Kernel dependency estimation. In *Advances in Neural Information Processing Systems (NIPS)*, volume 15, pages 873–880. MIT Press, 2003.

Ji Zhu and Trevor Hastie. Kernel logistic regression and the import vector machine. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14 (NIPS)*, pages 1081–1088. MIT Press, 2002.

Appendix A. Proof of Theorem 5

First, since c is convex with respect to f , we have $\forall 0 \leq t \leq 1$

$$c(y, f_Z + t(f_{Z^i} - f_Z), x) - c(y, f_Z, x) \leq t(c(y, f_{Z^i}, x) - c(y, f_Z, x)).$$

Then, by summing over all couples (x_k, y_k) in Z^i ,

$$R_{emp}(f_Z + t(f_{Z^i} - f_Z), Z^i) - R_{emp}(f_Z, Z^i) \leq t(R_{emp}(f_{Z^i}, Z^i) - R_{emp}(f_Z, Z^i)). \quad (7)$$

Symmetrically, we also have

$$R_{emp}(f_{Z^i} + t(f_Z - f_{Z^i}), Z^i) - R_{emp}(f_{Z^i}, Z^i) \leq t(R_{emp}(f_Z, Z^i) - R_{emp}(f_{Z^i}, Z^i)). \quad (8)$$

Thus, by summing (7) and (8), we obtain

$$\begin{aligned} R_{emp}(f_Z + t(f_{Z^i} - f_Z), Z^i) - R_{emp}(f_Z, Z^i) \\ + R_{emp}(f_{Z^i} + t(f_Z - f_{Z^i}), Z^i) - R_{emp}(f_{Z^i}, Z^i) \leq 0. \end{aligned} \quad (9)$$

Now, by definition of f_Z and f_{Z^i} ,

$$\begin{aligned} R_{reg}(f_Z, Z) - R_{reg}(f_Z + t(f_{Z^i} - f_Z), Z) \\ + R_{reg}(f_{Z^i}, Z^i) - R_{reg}(f_{Z^i} + t(f_Z - f_{Z^i}), Z^i) \leq 0. \end{aligned} \quad (10)$$

By using (9) and (10) we get

$$\begin{aligned} c(y_i, f_Z, x_i) - c(y_i, f_Z + t(f_{Z^i} - f_Z), x_i) \\ + m\lambda (\|f_Z\|_{\mathcal{H}}^2 - \|f_Z + t(f_{Z^i} - f_Z)\|_{\mathcal{H}}^2 + \|f_{Z^i}\|_{\mathcal{H}}^2 - \|f_{Z^i} + t(f_Z - f_{Z^i})\|_{\mathcal{H}}^2) \leq 0, \end{aligned}$$

hence, since c is C -Lipschitz continuous with respect to $f(x)$, and the inequality is true $\forall t \in [0, 1]$,

$$\begin{aligned}
 \|f_Z - f_{Z^i}\|_{\mathcal{H}}^2 &\leq \frac{1}{2t} (\|f_Z\|_{\mathcal{H}}^2 - \|f_Z + t(f_{Z^i} - f_Z)\|_{\mathcal{H}}^2 + \|f_{Z^i}\|_{\mathcal{H}}^2 - \|f_{Z^i} + t(f_Z - f_{Z^i})\|_{\mathcal{H}}^2) \\
 &\leq \frac{1}{2tm\lambda} (c(y_i, f_Z + t(f_{Z^i} - f_Z), x_i) - c(y_i, f_Z, x_i)) \\
 &\leq \frac{C}{2m\lambda} \|f_{Z^i}(x_i) - f_Z(x_i)\|_{\mathcal{Y}} \\
 &\leq \frac{C\kappa}{2m\lambda} \|f_{Z^i} - f_Z\|_{\mathcal{H}},
 \end{aligned}$$

which gives that

$$\|f_Z - f_{Z^i}\|_{\mathcal{H}} \leq \frac{C\kappa}{2m\lambda}.$$

This implies that, $\forall (x, y)$ a realization of (X, Y) ,

$$\begin{aligned}
 |c(y, f_Z, x) - c(y, f_{Z^i}, x)| &\leq C \|f_Z(x) - f_{Z^i}(x)\|_{\mathcal{Y}} \\
 &\leq C\kappa \|f_Z - f_{Z^i}\|_{\mathcal{H}} \\
 &\leq \frac{C^2\kappa^2}{2m\lambda}
 \end{aligned}$$

□