Multiple Operator-valued Kernel Learning

Hachem Kadri
LIF - CNRS / INRIA Lille - Sequel Project
Université Aix-Marseille
Marseille, France
hachem.kadri@lif.univ-mrs.fr

Alain Rakotomamonjy
LITIS EA 4108
Université de Rouen
St Etienne du Rouvray, France
alain.rakotomamony@insa-rouen.fr

Francis Bach
INRIA - Sierra Project
Ecole Normale Supérieure
Paris, France
francis.bach@inria.fr

Philippe Preux
INRIA Lille - Sequel Project
LIFL - CNRS, Université de Lille
Villeneuve d'Ascq, France
philippe.preux@inria.fr

Abstract

Positive definite operator-valued kernels generalize the well-known notion of reproducing kernels, and are naturally adapted to multi-output learning situations. This paper addresses the problem of learning a finite linear combination of infinite-dimensional operator-valued kernels which are suitable for extending functional data analysis methods to nonlinear contexts. We study this problem in the case of kernel ridge regression for functional responses with an ℓ_r-norm constraint on the combination coefficients ($r \geq 1$). The resulting optimization problem is more involved than those of multiple scalar-valued kernel learning since operator-valued kernels pose more technical and theoretical issues. We propose a multiple operator-valued kernel learning algorithm based on solving a system of linear operator equations by using a block coordinate-descent procedure. We experimentally validate our approach on a functional regression task in the context of finger movement prediction in brain-computer interfaces.

1 Introduction

During the past decades, a large number of algorithms have been proposed to deal with learning problems in the case of single-valued functions (*e.g.*, binary-output function for classification or real output for regression). Recently, there has been considerable interest in estimating vector-valued functions [21, 5, 7]. Much of this interest has arisen from the need to learn tasks where the target is a complex entity, not a scalar variable. Typical learning situations include multi-task learning [11], functional regression [12], and structured output prediction [4].

In this paper, we are interested in the problem of functional regression with functional responses in the context of brain-computer interface (BCI) design. More precisely, we are interested in finger movement prediction from electrocorticographic signals [23]. Indeed, from a set of signals measuring brain surface electrical activity on d channels during a given period of time, we want to predict, for any instant of that period whether a finger is moving or not and the amplitude of the finger flexion. Formally, the problem consists in learning a functional dependency between a set of d signals and a sequence of labels (a step function indicating whether a finger is moving or not) and between the same set of signals and vector of real values (the amplitude function). While, it is clear that this problem can be formalized as functional regression problem, from our point of view, such problem can benefit from the multiple operator-valued kernel learning framework. Indeed, for these problems, one of the difficulties arises from the unknown latency between the signal related to the finger

1

movement and the actual movement [23]. Hence, instead of fixing in advance some value for this latency in the regression model, our framework allows to learn it from the data by means of several operator-valued kernels.

If we wish to address functional regression problem in the principled framework of reproducing kernel Hilbert spaces (RKHS), we have to consider RKHSs whose elements are operators that map a function to another function space, possibly source and target function spaces being different. Working in such RKHSs, we are able to draw on the important core of work that has been performed on scalar-valued and vector-valued RKHSs [28, 21]. Such a functional RKHS framework and associated operator-valued kernels have been introduced recently [12, 13]. A basic question with reproducing kernels is how to build these kernels and what is the optimal kernel choice for a given application. In order to overcome the need for choosing a kernel before the learning process, several works have tried to address the problem of learning the scalar-valued kernel jointly with the decision function [18, 29]. Since these seminal works, many efforts have been carried out in order to theoretically analyze the kernel learning framework [9, 3] or in order to provide efficient algorithms [24, 1, 15]. While many works have been devoted to multiple *scalar-valued* kernel learning, this problem of kernel learning have been barely investigated for *operator-valued* kernels. One motivation of this work is to bridge the gap between multiple kernel learning (MKL) and operator-valued kernels by proposing a framework and an algorithm for learning a finite linear combination of operator-valued kernels. While each step of the scalar-valued MKL framework can be extended without major difficulties to operator-valued kernels, technical challenges arise at all stages because we deal with infinite dimensional spaces. It should be pointed out that in a recent work [10], the problem of learning the output kernel was formulated as an optimization problem over the cone of positive semidefinite matrices, and a block-coordinate descent method was proposed to solve it. However, they did not focus on learning the input kernel. In contrast, our multiple operator-valued kernel learning formulation can be seen as a way of learning simultaneously input and output kernels, although we consider a linear combination of kernels that are fixed in advance.

In this paper, we make the following contributions: **1)** we introduce a novel approach to infinite-dimensional multiple operator-valued kernel learning (MovKL) suitable for learning the functional dependencies and interactions between continuous data; **2)** we extend the original formulation of ridge regression in dual variables to the functional data analysis domain, showing how to perform nonlinear functional regression with functional responses by constructing a linear regression operator in an operator-valued kernel feature space (Section 2); **3)** we derive a dual form of the MovKL problem with functional ridge regression, and show that a solution of the related optimization problem exists (Section 2); **4)** we propose a block-coordinate descent algorithm to solve the MovKL optimization problem which involves solving a challenging linear system with a sum of block operator matrices (Section 3); **5)** we provide an empirical evaluation of MovKL performance which demonstrates its effectiveness on a BCI dataset (Section 4).

2 Problem Setting

Before describing the multiple operator-valued kernel learning algorithm that we will study and experiment with in this paper, we first review notions and properties of reproducing kernel Hilbert spaces with operator-valued kernels, show their connection to learning from multiple response data (multiple outputs; see [21] for discrete data and [12] for continuous data), and describe the optimization problem for learning kernels with functional response ridge regression.

2.1 Notations and Preliminaries

We start by some standard notations and definitions used all along the paper. Given a Hilbert space $\mathcal{H}$, $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and $\| \cdot \|_{\mathcal{H}}$ refer to its inner product and norm, respectively. We denote by $\mathcal{G}_x$ and $\mathcal{G}_y$ the separable real Hilbert spaces of input and output functional data, respectively. In functional data analysis domain, continuous data are generally assumed to belong to the space of square integrable functions L^2. In this work, we consider that $\mathcal{G}_x$ and $\mathcal{G}_y$ are the Hilbert space $L^2(\Omega)$ which consists of all equivalence classes of square integrable functions on a finite set Ω. Ω being potentially different for $\mathcal{G}_x$ and $\mathcal{G}_y$. We denote by $\mathcal{F}(\mathcal{G}_x, \mathcal{G}_y)$ the vector space of functions $f : \mathcal{G}_x \longrightarrow \mathcal{G}_y$, and by $\mathcal{L}(\mathcal{G}_y)$ the set of bounded linear operators from $\mathcal{G}_y$ to $\mathcal{G}_y$.

2

We consider the problem of estimating a function f such that $f(x_i) = y_i$ when observed functional data $(x_i, y_i)_{i=1,\ldots,n} \in (\mathcal{G}_x, \mathcal{G}_y)$. Since $\mathcal{G}_x$ and $\mathcal{G}_y$ are spaces of functions, the problem can be thought of as an operator estimation problem, where the desired operator maps a Hilbert space of factors to a Hilbert space of targets. We can define the regularized operator estimate of $f \in \mathcal{F}$ as:

$$f_\lambda \triangleq \arg\min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} \|y_i - f(x_i)\|_{\mathcal{G}_y}^2 + \lambda \|f\|_{\mathcal{F}}^2. \tag{1}$$

In this work, we are looking for a solution to this minimization problem in a function-valued reproducing kernel Hilbert space $\mathcal{F}$. More precisely, we mainly focus on the RKHS $\mathcal{F}$ whose elements are continuous linear operators on $\mathcal{G}_x$ with values in $\mathcal{G}_y$. The continuity property is obtained by considering a special class of reproducing kernels called Mercer kernels [7, Proposition 2.2]. Note that in this case, $\mathcal{F}$ is separable since $\mathcal{G}_x$ and $\mathcal{G}_y$ are separable [6, Corollary 5.2].

We now introduce (function) $\mathcal{G}_y$-valued reproducing kernel Hilbert spaces and show the correspondence between such spaces and positive definite (operator) $\mathcal{L}(\mathcal{G}_y)$-valued kernels. These extend the traditional properties of scalar-valued kernels.

Definition 1 *(function-valued RKHS)*

A Hilbert space $\mathcal{F}$ of functions from $\mathcal{G}_x$ to $\mathcal{G}_y$ is called a reproducing kernel Hilbert space if there is a positive definite $\mathcal{L}(\mathcal{G}_y)$-valued kernel $K_\mathcal{F}(w, z)$ on $\mathcal{G}_x \times \mathcal{G}_x$ such that:

   i. the function $z \longmapsto K_\mathcal{F}(w, z)g$ belongs to $\mathcal{F}$, $\forall z \in \mathcal{G}_x$, $w \in \mathcal{G}_x$, $g \in \mathcal{G}_y$,

   ii. $\forall f \in \mathcal{F}, w \in \mathcal{G}_x, g \in \mathcal{G}_y$, $\langle f, K_\mathcal{F}(w, \cdot)g \rangle_\mathcal{F} = \langle f(w), g \rangle_{\mathcal{G}_y}$ (reproducing property).

Definition 2 *(operator-valued kernel)*

An $\mathcal{L}(\mathcal{G}_y)$-valued kernel $K_\mathcal{F}(w, z)$ on $\mathcal{G}_x$ is a function $K_\mathcal{F}(\cdot, \cdot) : \mathcal{G}_x \times \mathcal{G}_x \longrightarrow \mathcal{L}(\mathcal{G}_y)$; furthermore:

   i. $K_\mathcal{F}$ is Hermitian if $K_\mathcal{F}(w, z) = K_\mathcal{F}(z, w)^*$, where $*$ denotes the adjoint operator,

   ii. $K_\mathcal{F}$ is positive definite on $\mathcal{G}_x$ if it is Hermitian and for every natural number r and all $\{(w_i, u_i)_{i=1,\ldots,r}\} \in \mathcal{G}_x \times \mathcal{G}_y$, $\sum_{i,j} \langle K_\mathcal{F}(w_i, w_j)u_j, u_i \rangle_{\mathcal{G}_y} \geq 0$.

Theorem 1 *(bijection between function-valued RKHS and operator-valued kernel)*
An $\mathcal{L}(\mathcal{G}_y)$-valued kernel $K_\mathcal{F}(w, z)$ on $\mathcal{G}_x$ is the reproducing kernel of some Hilbert space $\mathcal{F}$, if and only if it is positive definite.

The proof of Theorem 1 can be found in [21]. For further reading on operator-valued kernels and their associated RKHSs, see, e.g., [5, 6, 7].

2.2 Functional Response Ridge Regression in Dual Variables

We can write the ridge regression with functional responses optimization problem (1) as follows:

$$\min_{f \in \mathcal{F}} \frac{1}{2} \|f\|_{\mathcal{F}}^2 + \frac{1}{2n\lambda} \sum_{i=1}^{n} \|\xi_i\|_{\mathcal{G}_y}^2 \tag{2}$$
$$\text{with } \xi_i = y_i - f(x_i).$$

Now, we introduce the Lagrange multipliers $\alpha_i, i = 1, \ldots, n$ which are functional variables since the output space is the space of functions $\mathcal{G}_y$. For the optimization problem (2), the Lagrangian multipliers exist and the Lagrangian function is well defined. The method of Lagrange multipliers on Banach spaces, which is a generalization of the classical (finite-dimensional) Lagrange multipliers method suitable to solve certain infinite-dimensional constrained optimization problems, is applied here. For more details, see [16]. Let $\alpha = (\alpha_i)_{i=1,\ldots,n} \in \mathcal{G}_y^n$ the vector of functions containing the Lagrange multipliers, the Lagrangian function is defined as

$$L(f, \alpha, \xi) = \frac{1}{2} \|f\|_{\mathcal{F}}^2 + \frac{1}{2n\lambda} \|\xi\|_{\mathcal{G}_y^n}^2 + \langle \alpha, y - f(x) - \xi \rangle_{\mathcal{G}_y^n}, \tag{3}$$

where $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathcal{G}_y^n$, $y = (y_1, \ldots, y_n) \in \mathcal{G}_y^n$, $f(x) = (f(x_1), \ldots, f(x_n)) \in \mathcal{G}_y^n$, $\xi = (\xi_1, \ldots, \xi_n) \in \mathcal{G}_y^n$, and $\forall a, b \in \mathcal{G}_y^n, \langle a, b \rangle_{\mathcal{G}_y^n} = \sum_{i=1}^{n} \langle a_i, b_i \rangle_{\mathcal{G}_y}$.

3

Differentiating (3) with respect to $f \in \mathcal{F}$ and setting to zero, we obtain

$$f(.) = \sum_{i=1}^{n} K(x_i, .)\alpha_i, \tag{4}$$

where $K : \mathcal{G}_x \times \mathcal{G}_x \longrightarrow \mathcal{L}(\mathcal{G}_y)$ is the operator-valued kernel of $\mathcal{F}$.

Substituting this into (3) and minimizing with respect to ξ, we obtain the dual of the functional response kernel ridge regression (KRR) problem

$$\max_{\alpha} -\frac{n\lambda}{2}\|\alpha\|_{\mathcal{G}_y^n}^2 - \frac{1}{2}\langle \mathbf{K}\alpha, \alpha \rangle_{\mathcal{G}_y^n} + \langle \alpha, y \rangle_{\mathcal{G}_y^n}, \tag{5}$$

where $\mathbf{K} = [K(x_i, x_j)]_{i,j=1}^n$ is the block operator kernel matrix. The computational details regrading the dual formulation of functional KRR are derived in Appendix B of [14].

2.3 MovKL in Dual Variables

Let us now consider that the function $f(\cdot)$ is sum of M functions $\{f_k(\cdot)\}_{k=1}^M$ where each f_k belongs to a $\mathcal{G}_y$-valued RKHS with kernel $K_k(\cdot, \cdot)$. Similarly to scalar-valued multiple kernel learning, we can cast the problem of learning these functions f_k as

$$\min_{d \in \mathcal{D}} \min_{f_k \in \mathcal{F}_k} \sum_{k=1}^M \frac{\|f_k\|_{\mathcal{F}_k}^2}{2d_k} + \frac{1}{2n\lambda}\sum_{i=1}^n \|\xi_i\|_{\mathcal{G}_y}^2 \tag{6}$$
$$\text{with } \xi_i = y_i - \sum_{k=1}^M f_k(x_i),$$

where $d = [d_1, \cdots, d_M]$, $\mathcal{D} = \{d : \forall k, d_k \geq 0 \text{ and } \sum_k d_k^r \leq 1\}$ and $1 \leq r \leq \infty$. Note that this problem can equivalently be rewritten as an unconstrained optimization problem. Before deriving the dual of this problem, it can be shown by means of the generalized Weierstrass theorem [17] that this problem admits a solution. We report the proof in Appendix A of [14].

Now, following the lines of [24], a dualization of this problem leads to the following equivalent one

$$\min_{\mathbf{d} \in \mathcal{D}} \max_{\alpha \in \mathcal{G}_y^n} -\frac{n\lambda}{2}\|\alpha\|_{\mathcal{G}_y^n}^2 - \frac{1}{2}\langle \mathbf{K}\alpha, \alpha \rangle_{\mathcal{G}_y^n} + \langle \alpha, y \rangle_{\mathcal{G}_y^n}, \tag{7}$$

where $\mathbf{K} = \sum_{k=1}^M d_k \mathbf{K}_k$ and $\mathbf{K}_k$ is the block operator kernel matrix associated to the operator-valued kernel K_k. The KKT conditions also state that at optimality we have $f_k(\cdot) = \sum_{i=1}^n d_k K_k(x_i, \cdot)\alpha_i$.

3 Solving the MovKL Problem

After having presented the framework, we now devise an algorithm for solving this MovKL problem.

3.1 Block-coordinate descent algorithm

Since the optimization problem (6) has the same structure as a multiple scalar-valued kernel learning problem, we can build our MovKL algorithm upon the MKL literature. Hence, we propose to borrow from [15], and consider a block-coordinate descent method. The convergence of a block coordinate descent algorithm which is related closely to the Gauss-Seidel method was studied in works of [30] and others. The difference here is that we have operators and block operator matrices rather than matrices and block matrices, but this doesn't increase the complexity if the inverse of the operators are computable (typically analytically or by spectral decomposition). Our algorithm iteratively solves the problem with respect to α with d being fixed and then with respect to d with α being fixed (see Algorithm 1). After having initialized $\{d_k\}$ to non-zero values, this boils down to the following steps :

    1. with $\{d_k\}$ fixed, the resulting optimization problem with respect to α has the following form:

$$(\mathbf{K} + \lambda I)\alpha = y, \tag{8}$$

4

where $\mathbf{K} = \sum_{k=1}^{M} d_k \mathbf{K}_k$. While the form of solution is rather simple, solving this linear system is still challenging in the operator setting and we propose below an algorithm for its resolution.

2. with $\{f_k\}$ fixed, according to problem (6), we can rewrite the problem as

$$\min_{d \in \mathcal{D}} \sum_{k=1}^{M} \frac{\|f_k\|_{\mathcal{F}_k}^2}{d_k} \tag{9}$$

which has a closed-form solution and for which optimality occurs at [20]:

$$d_k = \frac{\|f_k\|^{\frac{2}{r+1}}}{(\sum_k \|f_k\|^{\frac{2r}{r+1}})^{1/r}}. \tag{10}$$

This algorithm is similar to that of [8] and [15] both being based on alternating optimization. The difference here is that we have to solve a linear system involving a block-operator kernel matrix which is a combination of basic kernel matrices associated to M operator-valued kernels. This makes the system very challenging, and we present an algorithm for solving it in the next paragraph. We also report in Appendix C of [14] a convergence proof of a modified version of the MovKL algorithm that minimizes a perturbation of the objective function (6) with a small positive parameter required to guarantee convergence [2].

3.2 Solving a linear system involving multiple operator-valued kernel matrices

One common way to construct operator-valued kernels is to build scalar-valued ones which are carried over to the vector-valued (resp. function-valued) setting by a positive definite matrix (resp. operator). In this setting an operator-valued kernel has the following form:

$$K(w, z) = G(w, z)T,$$

where G is a scalar-valued kernel and T is a positive operator in $\mathcal{L}(\mathcal{G}_y)$. In multi-task learning, T is a finite dimensional matrix that is expected to share information between tasks [11, 5]. More recently and for supervised functional output learning problems, T is chosen to be a multiplication or an integral operator [12, 13]. This choice is motivated by the fact that functional linear models for functional responses [25] are based on these operators and then such kernels provide an interesting alternative to extend these models to nonlinear contexts. In addition, some works on functional regression and structured-output learning consider operator-valued kernels constructed from the identity operator as in [19] and [4]. In this work we adopt a functional data analysis point of view and then we are interested in a finite combination of operator-valued kernels constructed from the identity, multiplication and integral operators. A problem encountered when working with operator-valued kernels in infinite-dimensional spaces is that of solving the system of linear operator equations (8). In the following we show how to solve this problem for two cases of operator-valued kernel combinations.

Case 1: multiple scalar-valued kernels and one operator. This is the simpler case where the combination of operator-valued kernels has the following form

$$K(w, z) = \sum_{k=1}^{M} d_k G_k(w, z)T, \tag{11}$$

where G_k is a scalar-valued kernel, T is a positive operator in $\mathcal{L}(\mathcal{G}_y)$, and d_k are the combination coefficients. In this setting, the block operator kernel matrix $\mathbf{K}$ can be expressed as a Kronecker product between the multiple scalar-valued kernel matrix $\mathbf{G} = \sum_{k=1}^{M} d_k \mathbf{G}_k$, where $\mathbf{G}_k = [G_k(x_i, x_j)]_{i,j=1}^n$, and the operator T. Thus we can compute an analytic solution of the system of equations (8) by inverting $\mathbf{K} + \lambda I$ using the eigendecompositions of $\mathbf{G}$ and T as in [13].

Case 2: multiple scalar-valued kernels and multiple operators. This is the general case where multiple operator-valued kernels are combined as follows

$$K(w, z) = \sum_{k=1}^{M} d_k G_k(w, z)T_k, \tag{12}$$

5

Algorithm 1 ℓ_r-norm MovKL

Input $\mathbf{K}_k$ for $k = 1, \ldots, M$

$d_k^1 \longleftarrow \dfrac{1}{M}$ for $k = 1, \ldots, M$

$\alpha \longleftarrow 0$

for $t = 1, 2, \ldots$ **do**

    $\alpha' \longleftarrow \alpha$

    $\mathbf{K} \longleftarrow \sum_k d_k^t \mathbf{K}_k$

    $\alpha \longleftarrow$ solution of $(\mathbf{K} + \lambda I)\alpha = y$

    if $\|\alpha - \alpha'\| < \epsilon$ **then**

        break

    end if

    $d_k^{t+1} \longleftarrow \dfrac{\|f_k\|^{\frac{2}{r+1}}}{(\sum_k \|f_k\|^{\frac{2r}{r+1}})^{1/r}}$ for $k = 1, \ldots, M$

end for

Algorithm 2 Gauss-Seidel Method

choose an initial vector of functions $\alpha^{(0)}$

repeat

    for $i = 1, 2, \ldots, n$

        $\alpha_i^{(t)} \longleftarrow$ sol. of (13):

            $[K(x_i, x_i) + \lambda I]\alpha_i^{(t)} = s_i$

    end for

until convergence

---

where G_k is a scalar-valued kernel, T_k is a positive operator in $\mathcal{L}(\mathcal{G}_y)$, and d_k are the combination coefficients. Inverting the associated block operator kernel matrix $\mathbf{K}$ is not feasible in this case, that is why we propose a Gauss-Seidel iterative procedure (see Algorithm 2) to solve the system of linear operator equations (8). Starting from an initial vector of functions $\alpha^{(0)}$, the idea is to iteratively compute, until a convergence condition is satisfied, the functions α_i according to the following expression

$$[K(x_i, x_i) + \lambda I]\alpha_i^{(t)} = y_i - \sum_{j=1}^{i-1} K(x_i, x_j)\alpha_j^{(t)} - \sum_{j=i+1}^{n} K(x_i, x_j)\alpha_j^{(t-1)}, \qquad (13)$$

where t is the iteration index. This problem is still challenging because the kernel $K(\cdot, \cdot)$ still involves a positive combination of operator-valued kernels. Our algorithm is based on the idea that instead of inverting the finite combination of operator-valued kernels $[K(x_i, x_i) + \lambda I]$, we can consider the following variational formulation of this system

$$\min_{\alpha_i^{(t)}} \frac{1}{2} \langle \sum_{k=1}^{M+1} K_k(x_i, x_i)\alpha_i^{(t)}, \alpha_i^{(t)} \rangle_{\mathcal{G}_y} - \langle s_i, \alpha_i^{(t)} \rangle_{\mathcal{G}_y},$$

where $s_i = y_i - \sum_{j=1}^{i-1} K(x_i, x_j)\alpha_j^{(t)} - \sum_{j=i+1}^{n} K(x_i, x_j)\alpha_j^{(t-1)}$, $K_k = d_k G_k T_k$, $\forall k \in \{1, \ldots, M\}$, and $K_{M+1} = \lambda I$.

Now, by means of a variable-splitting approach, we are able to decouple the role of the different kernels. Indeed, the above problem is equivalent to the following one :

$$\min_{\boldsymbol{\alpha}_i^{(t)}} \frac{1}{2} \langle \hat{\mathbf{K}}(x_i, x_i)\boldsymbol{\alpha}_i^{(t)}, \boldsymbol{\alpha}_i^{(t)} \rangle_{\mathcal{G}_y^M} - \langle \mathbf{s_i}, \boldsymbol{\alpha}_i^{(t)} \rangle_{\mathcal{G}_y^M} \quad \text{with} \quad \alpha_{i,1}^{(t)} = \alpha_{i,k}^{(t)} \quad \text{for} \quad k = 2, \ldots, M+1,$$

where $\hat{\mathbf{K}}(x_i, x_i)$ is the $(M+1) \times (M+1)$ diagonal matrix $[K_k(x_i, x_i)]_{k=1}^{M+1}$. $\boldsymbol{\alpha}_i^{(t)}$ is the vector $(\alpha_{i,1}^{(t)}, \ldots, \alpha_{i,M+1}^{(t)})$ and the $(M+1)$-dimensional vector $\mathbf{s_i} = (s_i, 0, \ldots, 0)$. We now have to deal with a quadratic optimization problem with equality constraints. Writing down the Lagrangian of this optimization problem and then deriving its first-order optimality conditions leads us to the following set of linear equations

$$\begin{cases} K_1(x_i, x_i)\alpha_{i,1} - s_i + \sum_{k=1}^{M} \gamma_k &= 0 \\ K_k(x_i, x_i)\alpha_{i,k} - \gamma_k &= 0 \\ \alpha_{i,1} - \alpha_{i,k} &= 0 \end{cases} \qquad (14)$$

where $k = 2, \ldots, M+1$ and $\{\gamma_k\}$ are the Lagrange multipliers related to the M equality constraints. Finally, in this set of equations, the operator-valued kernels have been decoupled and thus, if their inversion can be easily computed (which is the case in our experiments), one can solve the problem (14) with respect to $\{\alpha_{i,k}\}$ and γ_k by means of another Gauss-Seidel algorithm after simple reorganization of the linear system.

6
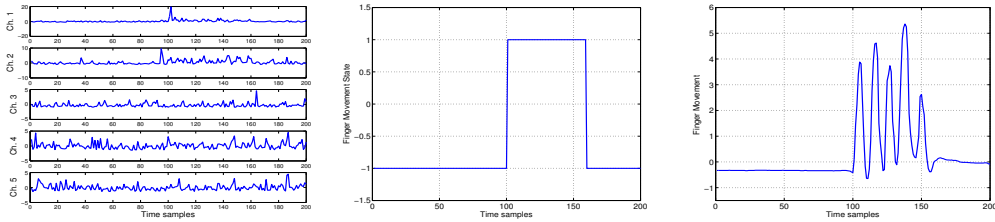
Figure 1: Example of a couple of input-output signals in our BCI task. (left) Amplitude modulation features extracted from ECoG signals over 5 pre-defined channels. (middle) Signal of labels denoting whether the finger is moving or not. (right) Real amplitude movement of the finger.

4   Experiments

In order to highlight the benefit of our multiple operator-valued kernel learning approach, we have considered a series of experiments on a real dataset, involving functional output prediction in a brain-computer interface framework. The problem we addressed is a sub-problem related to finger movement decoding from Electrocorticographic (ECoG) signals. We focus on the problem of estimating if a finger is moving or not and also on the direct estimation of the finger movement amplitude from the ECoG signals. The development of the full BCI application is beyond the scope of this paper and our objective here is to prove that this problem of predicting finger movement can benefit from multiple kernel learning.

To this aim, the fourth dataset from the BCI Competition IV [22] was used. The subjects were 3 epileptic patients who had platinium electrode grids placed on the surface of their brains. The number of electrodes varies between 48 to 64 depending on the subject, and their position on the cortex was unknown. ECoG signals of the subject were recorded at a 1KHz sampling using BCI2000 [27]. A band-pass filter from 0.15 to 200Hz was applied to the ECoG signals. The finger flexion of the subject was recorded at 25Hz and up-sampled to 1KHz by means of a data glove which measures the finger movement amplitude. Due to the acquisition process, a delay appears between the finger movement and the measured ECoG signal [22]. One of our hopes is that this time-lag can be properly learnt by means of multiple operator-valued kernels. Features from the ECoG signals are built by computing some band-specific amplitude modulation features, which is defined as the sum of the square of the band-specific filtered ECoG signal samples during a fixed time window.

For our finger movement prediction task, we have kept 5 channels that have been manually selected and split ECoG signals in portions of 200 samples. For each of these time segments, we have the label of whether at each time sample, the finger is moving or not as well as the real movement amplitudes. The dataset is composed of 487 couples of input-output signals, the output signals being either the binary movement labels or the real amplitude movement. An example of input-output signals are depicted in Figure 1. In a nutshell, the problem boils down to be a functional regression task with functional responses.

To evaluate the performance of the multiple operator-valued kernel learning approach, we use both: (1) the percentage of labels correctly recognized (LCR) defined by $(W_r/T_n) \times 100\%$, where W_r is the number of well-recognized labels and T_n the total number of labels to be recognized; (2) the residual sum of squares error (RSSE) as evaluation criterion for curve prediction

$$RSSE = \int \sum_i \{y_i(t) - \widehat{y}_i(t)\}^2 dt, \tag{15}$$

where $\widehat{y}_i(t)$ is the prediction of the function $y_i(t)$ corresponding to real finger movement or the finger movement state.

For the multiple operator-valued kernels having the form (12), we have used a Gaussian kernel with 5 different bandwidths and a polynomial kernel of degree 1 to 3 combined with three operators T: identity $Ty(t) = y(t)$, multiplication operator associated with the function e^{-t^2} defined by $Ty(t) = e^{-t^2} y(t)$, and the integral Hilbert-Schmidt operator with the kernel $e^{-|t-s|}$ proposed in [13], $Ty(t) = \int e^{-|t-s|}y(s)ds$. The inverses of these operators can be computed analytically.

7

Table 1: (**Left**) Results for the movement state prediction. Residual Sum of Squares Error (RSSE) and the percentage number of Labels Correctly Recognized (LCR) of : (1) baseline KRR with the Gaussian kernel, (2) functional response KRR with the integral operator-valued kernel, (3) MovKL with ℓ_∞, ℓ_1 and ℓ_2-norm constraint. (**Right**) Residual Sum of Squares Error (RSSE) results for finger movement prediction.

Algorithm	RSSE	LCR(%)	Algorithm	RSSE
KRR - scalar-valued -	68.32	72.91	KRR - scalar-valued -	88.21
KRR - functional response -	49.40	80.20	KRR - functional response -	79.86
MovKL - ℓ_∞ norm -	45.44	81.34	MovKL - ℓ_∞ norm -	76.52
MovKL - ℓ_1 norm -	48.12	80.66	MovKL - ℓ_1 norm -	78.24
MovKL - ℓ_2 norm -	**39.36**	**84.72**	MovKL - ℓ_2 norm -	**75.15**

While the inverses of the identity and the multiplication operators are easily and directly computable from the analytic expressions of the operators, the inverse of the integral operator is computed from its spectral decomposition as in [13]. The number of eigenfunctions as well as the regularization parameter λ are fixed using "one-curve-leave-out cross-validation" [26] with the aim of minimizing the residual sum of squares error.

Empirical results on the BCI dataset are summarized in Table 1. The dataset was randomly partitioned into 65% training and 35% test sets. We compare our approach in the case of ℓ_1 and ℓ_2-norm constraint on the combination coefficients with: (1) the baseline scalar-valued kernel ridge regression algorithm by considering each output independently of the others, (2) functional response ridge regression using an integral operator-valued kernel [13], (3) kernel ridge regression with an evenly-weighted sum of operator-valued kernels, which we denote by ℓ_∞-norm MovKL.

As in the scalar case, using multiple operator-valued kernels leads to better results. By directly combining kernels constructed from identity, multiplication and integral operators we could reduce the residual sum of squares error and enhance the label classification accuracy. Best results are obtained using the MovKL algorithm with ℓ_2-norm constraint on the combination coefficients. RSSE and LCR of the baseline kernel ridge regression are significantly outperformed by the operator-valued kernel based functional response regression. These results confirm that by taking into account the relationship between outputs we can improve performance. This is due to the fact that an operator-valued kernel induces a similarity measure between two pairs of input/output.

5 Conclusion

In this paper we have presented a new method for learning simultaneously an operator and a finite linear combination of operator-valued kernels. We have extended the MKL framework to deal with functional response kernel ridge regression and we have proposed a block coordinate descent algorithm to solve the resulting optimization problem. The method is applied on a BCI dataset to predict finger movement in a functional regression setting. Experimental results show that our algorithm achieves good performance outperforming existing methods. It would be interesting for future work to thoroughly compare the proposed MKL method for operator estimation with previous related methods for multi-class and multi-label MKL in the contexts of structured-output learning and collaborative filtering.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments. This research was funded by the Ministry of Higher Education and Research, Nord-Pas-de-Calais Regional Council and FEDER (Contrat de Projets Etat Region CPER 2007-2013), ANR projects LAMPADA (ANR-09-EMER-007) and ASAP (ANR-09-EMER-001), and by the IST Program of the European Community under the PASCAL2 Network of Excellence (IST-216886). This publication only reflects the authors' views. Francis Bach was partially supported by the European Research Council (SIERRA Project).

8

References

[1] J. Aflalo, A. Ben-Tal, C. Bhattacharyya, J. Saketha Nath, and S. Raman. Variable sparsity kernel learning. *JMLR*, 12:565–592, 2011.

[2] A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.

[3] F. Bach. Consistency of the group Lasso and multiple kernel learning. *JMLR*, 9:1179–1225, 2008.

[4] C. Brouard, F. d'Alché-Buc, and M. Szafranski. Semi-supervised penalized output kernel regression for link prediction. In *Proc. ICML*, 2011.

[5] A. Caponnetto, C. A. Micchelli, M. Pontil, and Y. Ying. Universal multi-task kernels. *JMLR*, 68:1615–1646, 2008.

[6] C. Carmeli, E. De Vito, and A. Toigo. Vector valued reproducing kernel Hilbert spaces of integrable functions and mercer theorem. *Analysis and Applications*, 4:377–408, 2006.

[7] C. Carmeli, E. De Vito, and A. Toigo. Vector valued reproducing kernel Hilbert spaces and universality. *Analysis and Applications*, 8:19–61, 2010.

[8] C. Cortes, M. Mohri, and A. Rostamizadeh. L_2 regularization for learning kernels. In *Proc. UAI*, 2009.

[9] C. Cortes, M. Mohri, and A. Rostamizadeh. Generalization bounds for learning kernels. In *ICML*, 2010.

[10] F. Dinuzzo, C. S. Ong, P. Gehler, and G. Pillonetto. Learning output kernels with block coordinate descent. In *Proc. ICML*, 2011.

[11] T. Evgeniou, C. A. Micchelli, and M. Pontil. Learning multiple tasks with kernel methods. *JMLR*, 6:615–637, 2005.

[12] H. Kadri, E. Duflos, P. Preux, S. Canu, and M. Davy. Nonlinear functional regression: a functional RKHS approach. In *Proc. AISTATS*, pages 111–125, 2010.

[13] H. Kadri, A. Rabaoui, P. Preux, E. Duflos, and A. Rakotomamonjy. Functional regularized least squares classification with operator-valued kernels. In *Proc. ICML*, 2011.

[14] H. Kadri, A. Rakotomamonjy, F. Bach, and P. Preux. Multiple operator-valued kernel learning. Technical Report 00677012, INRIA, 2012.

[15] M. Kloft, U. Brefeld, S. Sonnenburg, and A. Zien. ℓ_p-norm multiple kernel learning. *JMLR*, 12:953–997, 2011.

[16] S. Kurcyusz. On the existence and nonexistence of lagrange multipliers in Banach spaces. *Journal of Optimization Theory and Applications*, 20:81–110, 1976.

[17] A. Kurdila and M. Zabarankin. *Convex Functional Analysis*. Birkhauser Verlag, 2005.

[18] G. Lanckriet, N. Cristianini, L. El Ghaoui, P. Bartlett, and M. Jordan. Learning the kernel matrix with semi-definite programming. *JMLR*, 5:27–72, 2004.

[19] H. Lian. Nonlinear functional models for functional responses in reproducing kernel Hilbert spaces. *The Canadian Journal of Statistics*, 35:597–606, 2007.

[20] C. Micchelli and M. Pontil. Learning the kernel function via regularization. *JMLR*, 6:1099–1125, 2005.

[21] C. A. Micchelli and M. Pontil. On learning vector-valued functions. *Neural Comput.*, 17:177–204, 2005.

[22] K. J. Miller and G. Schalk. Prediction of finger flexion: 4th brain-computer interface data competition. BCI Competition IV, 2008.

[23] T. Pistohl, T. Ball, A. Schulze-Bonhage, A. Aertsen, and C. Mehring. Prediction of arm movement trajectories from ECoG-recordings in humans. *Journal of Neuroscience Methods*, 167(1):105–114, 2008.

[24] A. Rakotomamonjy, F. Bach, Y. Grandvalet, and S. Canu. SimpleMKL. *JMLR*, 9:2491–2521, 2008.

[25] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis, 2nd ed*. Springer Verlag, New York, 2005.

[26] John A. Rice and B. W. Silverman. Estimating the mean and covariance structure nonparametrically when the data are curves. *Journal of the Royal Statistical Society. Series B*, 53(1):233–243, 1991.

[27] G. Schalk, D. J. McFarland, T. Hinterberger, N. Birbaumer, and J. R. Wolpaw. BCI2000: a general-purpose brain-computer interface system. *Biomedical Engineering, IEEE Trans. on*, 51:1034–1043, 2004.

[28] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2002.

[29] S. Sonnenburg, G. Rätsch, C. Schäfer, and B. Schölkopf. Large scale multiple kernel learning. *JMLR*, 7:1531–1565, 2006.

[30] P. Tseng. Convergence of block coordinate descent method for nondifferentiable minimization. *J. Optim. Theory Appl.*, 109:475–494, 2001.

9