

Feature space perspectives for learning the kernel

Charles A. Micchelli · Massimiliano Pontil

Received: 15 August 2006 / Revised: 17 August 2006 / Accepted: 11 October 2006 / Published online:
9 January 2007
Springer Science + Business Media, LLC 2007

Abstract In this paper, we continue our study of learning an optimal kernel in a prescribed convex set of kernels (Micchelli & Pontil, 2005). We present a reformulation of this problem within a feature space environment. This leads us to study regularization in the dual space of all continuous functions on a compact domain with values in a Hilbert space with a mix norm. We also relate this problem in a special case to \mathcal{L}^p regularization.

Keywords Banach space regularization · Convex optimization · Learning the kernels · Kernel methods · Sparsity

1 Introduction

A central theme of this paper is the problem of learning a kernel in a prescribed convex set of kernels \mathcal{K} . Our previous work on this problem which was motivated by its potential application in machine learning focused on finding a suitable optimal kernel. Here, we study an equivalent feature space formulation of this problem. This leads us to explore the relationship between the problem of finding an optimal kernel and regularization in the dual space

Editors: Olivier Bousquet and Andre Elisseeff

This work was supported by NSF Grant ITR-0312113, EPSRC Grant GR/T18707/01 and by the IST Programme of the European Community, under the PASCAL Network of Excellence IST-2002-506778.

C. A. Micchelli
Department of Mathematics and Statistics, State University of New York, The University at Albany,
1400 Washington Avenue, Albany, NY 12222, USA
e-mail: charles_micchelli@hotmail.com

M. Pontil (✉)
Department of Computer Science, University College London, Gower Street, London WC1E, England,
UK
e-mail: m.pontil@cs.ucl.ac.uk

of the space of continuous functions on a compact domain with values in a Hilbert space. We also describe related regularization techniques in \mathcal{L}^p spaces which naturally arise in our investigation.

Argyriou, Micchelli, and Pontil (2005) and Micchelli and Pontil (2005) proposed to find a good kernel K by solving the variational problem

$$\min \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, f(x_j)) + \mu \|f\|_K^2, : f \in \mathcal{H}_K, K \in \mathcal{K} \right\} \quad (1.1)$$

where $L : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ is a prescribed loss function, $\|\cdot\|_K$ the norm in a reproducing kernel Hilbert space of real-valued functions on some input space \mathcal{X} with kernel K , μ a positive parameter, \mathcal{K} is a prescribed set of kernels and $\mathbb{N}_m := \{1, \dots, m\}$. This problem has been studied from different perspectives in a substantial number of papers. Specifically, in statistics, it has been motivated by Lee et al. (2004) and Lin and Zhang (2003) as a generalization of the *lasso*, a technique introduced in Tibshirani (1996) which also relates to basis pursuit denoising (Chen, Donoho, & Saunders, 1998) and to a linear programming approach for feature selection (Fung & Mangasarian, 2004); in machine learning, problem (1.1) has been studied in the context of support vector machines as a mean to optimize the margin or soft-margin error used therein (Bach, Lanckriet & Jordan, 2004; Lanckriet et al., 2004); in learning theory, it has been investigated with the intention of improving the approximation error (Wu, Ying & Zhou, to appear; Micchelli et al., 2005). For additional interesting observations related to the theme of this paper see (Bousquet & Herrmann, 2003; Cristianini et al., 2002; Gunn & Kandola, 2002; Herbster, 2004; Ong, Smola, & Williamson, 2003).

In Section 2, we describe the main result of the paper which relates problem (1.1) to our feature space extremal problem. Indeed, the problem described above concerns the choice of an optimal kernel for kernel based learning algorithms while the second problem we study is the reformulation of it within a feature space environment. We demonstrate in great generality that these problems are equivalent and characterize the form of the solutions for both problems. We also provide a description for an optimal feature map solution analogous to the one we derived in our earlier work on learning the kernel (Argyriou, Micchelli, & Pontil, 2005; Micchelli & Pontil, 2005). A detailed description of this result appear in Section 2. However, the proof is postponed until Section 6. In Section 3, we present specific motivating examples when \mathcal{K} is the convex hull of a *finite* set of prescribed kernels. Moreover, for these examples we provide an alternate derivation of the main result in Section 2 by using a decomposition theorem from Aronszajn (1950). In Section 4, we discuss the connection between learning the kernel and \mathcal{L}^1 regularization. Section 5 contains related results for \mathcal{L}^p regularization and provide a representer theorem in the spirit of Micchelli and Pontil (2004). We end the paper with a discussion of future research directions and commentaries on our results.

We remark that an interesting aspect of the feature space regularization we present here is not only does it involve linear functionals, but also that it is a Banach space regularization method. Indeed, as we shall show, the appropriate norm for the functionals is induced by a mix norm on a space of functions with values in the Hilbert space associated with the feature map. Finally, we also explore similar issues for an \mathcal{L}^p analog of the convex hull of a fix set of kernels.

2 Main result

In this section, we present our main result. First, we recall the notion of reproducing kernel Hilbert spaces and continuously parameterized convex set of kernels.

2.1 Integrals of kernels

Let \mathcal{X} be an *input set*. By a *kernel* we mean a function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that for every finite set of inputs $\mathbf{x} = \{x_j : j \in \mathbb{N}_m\} \subseteq \mathcal{X}$ and every $m \in \mathbb{N}$, the $m \times m$ matrix $K_{\mathbf{x}} := (K(x_i, x_j) : i, j \in \mathbb{N}_m)$ is *positive semi-definite*. According to Aronszajn and Moore, every kernel has associated to it an (essentially) *unique* Hilbert space \mathcal{H}_K of real-valued functions on \mathcal{X} with inner product $\langle \cdot, \cdot \rangle_K$ such that K is its reproducing kernel (Aronszajn, 1950). This means, for every $f \in \mathcal{H}_K$ and $x \in \mathcal{X}$, that $\langle f, K(x, \cdot) \rangle_K = f(x)$.

We use the notation $\mathcal{A}(\mathcal{X})$ for the set of all kernels on the set \mathcal{X} and $\mathcal{A}_+(\mathcal{X})$ for the subset of kernels K such that, for *each* input \mathbf{x} , the matrix $K_{\mathbf{x}}$ is positive definite.

Let Θ be a compact Hausdorff space, $C(\Theta)$ the space of continuous real-valued functions on Θ and $\mathcal{M}(\Theta)$ the set of all *probability measures* on Θ . Let $G : \Theta \rightarrow \mathcal{A}_+(\mathcal{X})$ be a continuous map. By this we mean that, for each $x, t \in \mathcal{X}$, the function of $\theta \mapsto G(\theta)(x, t)$ is continuous on Θ . The set of kernels $\mathcal{G} := \{G(\theta) : \theta \in \Theta\}$ induces the convex set of kernels

$$\mathcal{K}(\mathcal{G}) := \left\{ \int_{\Theta} G(\theta) dp(\theta) : p \in \mathcal{M}(\Theta) \right\} \tag{2.1}$$

which we shall consider below. We note that when $\Theta = \mathbb{N}_n$ then $\mathcal{K}(\mathcal{G})$ equal the convex hull of \mathcal{G} .

2.2 Regularization error functional

Let $D := \{(x_j, y_j) : j \in \mathbb{N}_m\} \subset \mathcal{X} \times \mathbb{R}$ be prescribed data and \mathbf{y} the vector $(y_j : j \in \mathbb{N}_m)$. Each kernel $K \in \mathcal{K}(\mathcal{G})$ gives rise to a RKHS \mathcal{H}_K . For each $f \in \mathcal{H}_K$, we introduce the *information operator* $I_{\mathbf{x}}(f) := (f(x_j) : j \in \mathbb{N}_m) \in \mathbb{R}^m$ of values of f on the set of inputs in \mathbf{x} . We consider $I_{\mathbf{x}} : \mathcal{H}_K \rightarrow \mathbb{R}^m$ a linear map and on \mathbb{R}^m we put the usual inner product. Thus, for any two vectors $c = (c_j : j \in \mathbb{N}_m)$ and $d = (d_j : j \in \mathbb{N}_m)$ we write $(c, d) := \sum_{j \in \mathbb{N}_m} c_j d_j$. A straightforward computation identifies the adjoint $I_{\mathbf{x}}^* : \mathbb{R}^m \rightarrow \mathcal{H}_K$, for every $c = (c_j : j \in \mathbb{N}_m) \in \mathbb{R}^m$, as

$$I_{\mathbf{x}}^* c = \sum_{j \in \mathbb{N}_m} c_j K_j,$$

where $K_j = K(x_j, \cdot)$. A *regularization error function* is any function $q : \mathbb{R}^m \times \mathbb{R}_+ \rightarrow \mathbb{R}$. We write any vector $v \in \mathbb{R}^m \times \mathbb{R}_+$ in the form (c, t) for some $c \in \mathbb{R}^m$ and $t \in \mathbb{R}_+$. In other words, the vector v is the concatenation of the vector c and the scalar t . There should be no confusion with this “double duty” notation since in this case one argument is a vector and the other is a scalar. Likewise, we shall denote $q(v)$ as $q(c, t)$ and $q_{\inf} := \inf\{q(v) : v \in \mathbb{R}^m \times \mathbb{R}_+\}$. The regularization error function will allow us to balance the data $I_{\mathbf{x}} f$ with the norm $\|f\|_K := \sqrt{\langle f, f \rangle_K}$ and leads us to the following definition.

Definition 2.1. We say that $q : \mathbb{R}^m \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is an acceptable regularization error function provided that

1. q is lower semi-continuous, that is, for each $\lambda \in \mathbb{R}$ the set $\mathcal{U}_\lambda := \{v : v \in \mathbb{R}^m \times \mathbb{R}_+, q(v) \leq \lambda\}$ is a closed subset of $\mathbb{R}^m \times \mathbb{R}_+$;
2. given any $\lambda > 0$ there exists $\rho > 0$ such that whenever $(c, t) \in \mathcal{U}_\lambda$ then $t \in [0, \rho]$;
3. for each $c \in \mathbb{R}^m$ the function $q(c, \cdot) : \mathbb{R}_+ \rightarrow \mathbb{R}$ is non-decreasing on \mathbb{R}_+ .

An important example for an acceptable regularization error function in machine learning has the form

$$q(c, t) = \sum_{j \in \mathbb{N}_m} L(y_j, c_j) + \mu J(t), \tag{2.2}$$

where $L : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ is a continuous loss function (typically convex), μ a positive constant and $J : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a *strictly increasing* function. The standard choice for the function J is $J(t) = t, t \in \mathbb{R}_+$. This leads to the usual kernel-based regularization algorithm studied extensively in the literature. However, as we shall see later, from the feature space point of view, the choice $J(t) = \sqrt{t}, t \in \mathbb{R}_+$, is widely studied. Whenever we consider the special case (2.2) we always assume that the above properties of L and J are satisfied.

An acceptable regularization error function q gives rise to a functional $Q(\cdot, K) : \mathcal{H}_K \rightarrow \mathbb{R}$ defined, for all $f \in \mathcal{H}_K$, as

$$Q(f, K) = q(I_{\mathbf{x}}f, \|f\|_K^2). \tag{2.3}$$

Properties 1–3 above guarantee, for every $K \in \mathcal{A}(\mathcal{X})$, that $Q(\cdot, K)$ has a minimum f_K over $f \in \mathcal{H}_K$. Since we do not assume here that q is strictly convex this minimum may not be unique, hence our notation f_K is to be interpreted to mean that f_K is *any* minimum for $Q(\cdot, K)$. We let

$$E(K) := \min\{Q(f, K) : f \in \mathcal{H}_K\} = Q(f_K, K)$$

and note that using the weak compactness of the unit ball in \mathcal{H}_K , the weak lower semi-continuity of the norm on \mathcal{H}_K and the first hypothesis on an acceptable regularization functional q ensure the existence of the minimum. In our previous work (Argyriou, Micchelli, & Pontil, 2005) we studied the problem of choosing a kernel $K \in \mathcal{K}(\mathcal{G})$ which solves the variational problem

$$E_G := \inf\{E(K) : K \in \mathcal{K}(\mathcal{G})\}. \tag{2.4}$$

Any kernel $\hat{K} \in \mathcal{K}(\mathcal{G})$ for which $E_G = E(\hat{K})$ is called an optimal kernel and $f_{\hat{K}}$ is called an optimal function. A main goal of this paper is to give the variational problem (2.4) a feature space interpretation. To this end, we assume $G : \Theta \rightarrow \mathcal{A}_+(\mathcal{X})$ is expressed in terms of a *feature map* $\Phi : \Theta \times \mathcal{X} \rightarrow \mathcal{W}$, where \mathcal{W} is a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{W}}$ and corresponding norm $\|\cdot\|_{\mathcal{W}}$. That is, for each $x, y \in \mathcal{X}$ and $\theta \in \Theta$, we have that $G(\theta)(x, y) = \langle \Phi(\theta, x), \Phi(\theta, y) \rangle_{\mathcal{W}}$. For example, when $\Theta = \mathbb{N}_n$, we construct $\mathcal{K}(G)$ as the *convex hull* of the *finite set* of kernels $G(\theta_1), \dots, G(\theta_N)$. In addition, when \mathcal{W} is finite dimensional, say $\mathcal{W} = \mathbb{R}^M$, then each of the kernels is constructed from a *finite number of features*.

We require that the feature map has the property that, for each $x \in \mathcal{X}$, the function $\Phi(\cdot, x) : \Theta \rightarrow \mathcal{W}$ is continuous. In particular, we conclude that $\|\Phi(\cdot, x)\|$ is a continuous

real-valued function on Ω . From this it easily follows that, for each $x, y \in \mathcal{X}$, $G(\cdot)(x, y)$ is a continuous function on Ω . Let us use the notation $C(\Theta, \mathcal{W})$ for the set of all continuous functions on Θ with values in \mathcal{W} , where we give any $\Psi \in C(\Theta, \mathcal{W})$ the norm

$$\|\Psi\|_{\infty, \mathcal{W}} := \max \{ \|\Psi(\theta)\|_{\mathcal{W}} : \theta \in \Theta \}.$$

Therefore, by our assumptions above for every $x \in \mathcal{X}$ the function $\Phi(\cdot, x) \in C(\Theta, \mathcal{W})$.

We introduce *feature functions* $\Phi_j : \Theta \rightarrow \mathcal{W}$ induced by the inputs \mathbf{x} and defined, for every $\theta \in \Theta$, as $\Phi_j(\theta) = \Phi(\theta, x_j)$ and assume that they are linearly independent on Θ . This is equivalent to the assumption that G maps into $\mathcal{A}_+(\mathcal{X})$. We use the standard notation $C^*(\Theta, \mathcal{W})$ for the space of all continuous linear functionals on $C(\Theta, \mathcal{W})$. A precise description of $C^*(\Theta, \mathcal{W})$ can be given by using the representation theorem for $C(\Theta, \mathbb{R})$ in term of regular Borel measures on Θ . However, this information is not needed in our presentation. Instead, corresponding to any feature map and any regularization error function as described above we introduce a feature space regularization functional $V : C^*(\Theta, \mathcal{W}) \rightarrow \mathbb{R}$ given, for any $B \in C^*(\Theta, \mathcal{W})$, by

$$V(B) = q(D_{\mathbf{x}}(B), \|B\|^2), \tag{2.5}$$

where $D_{\mathbf{x}} : C^*(\Theta, \mathcal{W}) \rightarrow \mathbb{R}^m$ is the linear operator defined as $D_{\mathbf{x}}(B) = (B(\Phi_j) : j \in \mathbb{N}_m)$. Recall that the norm of the linear functional B is defined as

$$\|B\| := \sup\{B(\Psi) : \|\Psi\|_{\infty, \mathcal{W}} \leq 1\}.$$

We introduce the variational problem

$$V_{\Phi} := \inf\{V(B) : B \in C^*(\Theta, \mathcal{W})\} \tag{2.6}$$

which we will henceforth refer to as the feature space variational problem. Any linear functional $\hat{B} \in C^*(\Theta, \mathcal{W})$ for which $V_{\Phi} = V(\hat{B})$ is called an optimal linear functional. Our main result is the following fact.

Theorem 2.1. *Under the above hypotheses, we have that $E_G = V_{\Phi}$. Moreover, there exist a finitely supported measure $\hat{p} \in \mathcal{M}(\Theta)$ with at most $m + 1$ atoms and a vector $\hat{c} \in \mathbb{R}^m$ such that the kernel*

$$\hat{K} = \int_{\Theta} G(\theta) d\hat{p}(\theta)$$

is an optimal kernel, the function

$$f_{\hat{K}} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \hat{K}_j$$

is an optimal function and the linear functional $\hat{B} \in C^*(\Theta, \mathcal{W})$ defined, for any $\Psi \in C(\Theta, \mathcal{W})$, as

$$\hat{B}(\Psi) = \int_{\Theta} (\hat{\Gamma}(\theta), \Psi(\theta))_{\mathcal{W}} d\hat{p}(\theta)$$

is an optimal linear functional, where the function $\hat{\Gamma} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \Phi_j$ has the property, for every atom θ of \hat{p} , that $\|\hat{\Gamma}(\theta)\|_{\mathcal{W}} = \|\hat{\Gamma}\|_{\infty, \mathcal{W}} = \|\hat{B}\|$.

This theorem not only shows that the values of the extremal problems (2.4) and (2.6) are the same, that is $E_G = V_\Phi$, but also connects their solutions. In particular, the kernel \hat{K} defined in the theorem is an *optimal kernel* and the *optimal linear functional* for the feature space variational problem is *connected* to \hat{K} by means of the vector \hat{c} and function $\hat{\Gamma}$ as described in the theorem.

We also wish to point out another connection between the optimal function, optimal linear functional and the feature map, namely, for every $x \in \mathcal{X}$, we have that

$$f_{\hat{K}}(x) = \hat{B}(\Phi(\cdot, x)).$$

We mention here an immediate corollary of the above theorem.

Corollary 2.1. *Under the hypothesis above, we have that $E_G = V_\Phi$ where*

$$E_G = \inf \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, f(x_j)) + \mu J(\|f\|_K^2) : f \in \mathcal{H}_K, K \in \mathcal{K}(\mathcal{G}) \right\}$$

and

$$V_\Phi = \min \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, B(\Phi_j)) + \mu J(\|B\|^2) : B \in C^*(\Theta, \mathcal{W}) \right\}.$$

Moreover, there exist a finitely supported measure $\hat{p} \in \mathcal{M}(\Theta)$ with at most $m + 1$ atoms and a vector $\hat{c} \in \mathbb{R}^m$ such that the kernel

$$\hat{K} = \int_{\Theta} G(\theta) d\hat{p}(\theta)$$

is an optimal kernel, the function

$$f_{\hat{K}} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \hat{K}_j$$

is an optimal function and the linear functional $\hat{B} \in C^*(\Theta, \mathcal{W})$ defined, for any $\Psi \in C(\Theta, \mathcal{W})$, as

$$\hat{B}(\Psi) = \int_{\Theta} \langle \hat{\Gamma}(\theta), \Psi(\theta) \rangle_{\mathcal{W}} d\hat{p}(\theta)$$

is an optimal linear functional, where the function $\hat{\Gamma} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \Phi_j$ has the property, for every atom θ of \hat{p} , that $\|\hat{\Gamma}(\theta)\|_{\mathcal{W}} = \|\hat{\Gamma}\|_{\infty, \mathcal{W}} = \|\hat{B}\|$.

Before we prove Theorem 2.1 we describe several examples of some potential practical importance in the next sections. We postpone the proof of Theorem 2.1 to Section 6.

3 A practical example: Finitely many kernels

In this section, we specialize our analysis to the practically important case that $\Theta = \mathbb{N}_n$. Thus, we express \mathcal{G} as a finite collection of kernels, that is $\mathcal{G} = \{G_r : r \in \mathbb{N}_n\} \subseteq \mathcal{A}_+(\mathcal{X})$, $\mathcal{M}(\Theta) = \mathbb{S}^n$, the n -dimensional simplex given by $\mathbb{S}^n = \{(\lambda_r : r \in \mathbb{N}_n) \in \mathbb{R}^n : \lambda_r \geq 0, \sum_{s \in \mathbb{N}_n} \lambda_s = 1\}$, and

$$\mathcal{K}(\mathcal{G}) = \left\{ \sum_{r \in \mathbb{N}_n} \lambda_r G_r : (\lambda_r : r \in \mathbb{N}_n) \in \mathbb{S}^n \right\}.$$

We let \mathcal{W}^n be the n -fold cross product of \mathcal{W} , that is,

$$\mathcal{W}^n = \{w = (w_r : r \in \mathbb{N}_n) : w_r \in \mathcal{W}, r \in \mathbb{N}_n\}$$

equipped it with the ℓ_n^p norm, where $1 \leq p \leq \infty$, given, for any $w \in \mathcal{W}^n$, as

$$\|w\|_{p, \mathcal{W}} := \left(\sum_{r \in \mathbb{N}_n} \|w_r\|_{\mathcal{W}}^p \right)^{\frac{1}{p}}$$

and denote the resulting Banach space by $\mathcal{W}^{p,n}$. In the case that $p = 2$, this is a Hilbert space with inner product defined, for every $u, w \in \mathcal{W}^n$ as $\langle w, u \rangle_{\mathcal{W}^n} := \sum_{r \in \mathbb{N}_n} \langle w_r, u_r \rangle_{\mathcal{W}}$. Clearly, the space $C(\Theta, \mathcal{W})$ is identified with $\mathcal{W}^{\infty,n}$ and its dual space, $C^*(\Theta, \mathcal{W})$, with $\mathcal{W}^{1,n}$. Thus, a linear functional $B \in C^*(\Theta, \mathcal{W})$ corresponds uniquely to a vector $w \in \mathcal{W}^n$ by means of the equation

$$B(u) = \langle w, u \rangle_{\mathcal{W}^n}, \quad u \in \mathcal{W}^n$$

and its norm is given by

$$\|B\| = \|w\|_{1, \mathcal{W}}.$$

We shall now specialize Corollary 2.1 to this case.

Corollary 3.1. *Under the hypotheses above, we have that*

$$E_G = V_\Phi, \tag{3.1}$$

where

$$E_G = \inf \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, f(x_j)) + \mu J(\|f\|_K^2) : f \in \mathcal{H}_K, K \in \mathcal{K}(\mathcal{G}) \right\}$$

and

$$V_\Phi = \min \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, \langle w, \Phi(x_j) \rangle_{\mathcal{W}^n}) + \mu J(\|w\|_{1, \mathcal{W}}^2) : w \in \mathcal{W}^n \right\}. \tag{3.2}$$

Moreover, there exist $\hat{\lambda} \in \mathbb{S}^n$ and a vector $\hat{c} \in \mathbb{R}^m$ such that the set $\mathcal{R} := \{r : \hat{\lambda}_r > 0\}$ has cardinality at most $\min(m + 1, n)$, the kernel

$$\hat{K} = \sum_{r \in \mathcal{R}} \hat{\lambda}_r G_r \tag{3.3}$$

is an optimal kernel, the function

$$f_{\hat{K}} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \hat{K}_j$$

is an optimal function and the vector

$$\hat{w} = (\hat{\lambda}_r \hat{\Gamma}_r : r \in \mathbb{N}_n)$$

is an optimal solution to problem (3.2), where the $\hat{\Gamma} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \Phi(x_j) \in \mathcal{W}^n$ has the property, for every $r \in \mathcal{R}$, that $\|\hat{\Gamma}_r\|_{\mathcal{W}} = \|\hat{\Gamma}\|_{\infty, \mathcal{W}} = \|\hat{w}\|_{1, \mathcal{W}}$.

Note that the last statement in the corollary concerning the optimal vector $\hat{w} = (\hat{w}_r : r \in \mathbb{N}_n)$ says, for every $r \in \mathbb{N}_n$, that

$$\|\hat{w}_r\|_{\mathcal{W}} = \hat{\lambda}_r \|\hat{\Gamma}\|_{\infty, \mathcal{W}}.$$

Thus, summing both sides of this equation over $r \in \mathbb{N}_n$, we conclude that

$$\|\hat{w}\|_{1, \mathcal{W}} = \|\hat{\Gamma}\|_{\infty, \mathcal{W}}$$

and, so

$$\hat{\lambda}_r = \frac{\|\hat{w}_r\|_{\mathcal{W}}}{\|\hat{w}\|_{1, \mathcal{W}}}. \tag{3.4}$$

This formula demonstrate that a solution to the feature space variational problem provides a choice for the optimal kernel in Eq. (3.3).

When $\mathcal{W} = \mathbb{R}$, problem (3.2) becomes

$$\min \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, \langle w, \Phi(x_j) \rangle_{\mathbb{R}^n}) + \mu J \left(\left(\sum_{r \in \mathbb{N}_n} |w_r| \right)^2 \right) : w \in \mathbb{R}^n \right\}. \tag{3.5}$$

This problem is closely related to some well-known function estimation techniques. In particular, when the loss function L is the square loss and J is chosen to be $J(t) = \sqrt{t}$, $t \in \mathbb{R}_+$, the variational problem (3.5) has been studied in statistics under the name of *lasso* (Tibshirani, 1996), in signal processing as *basis pursuit denoising* (Chen, Donoho, & Saunders, 1998), and in linear programming (Fung & Mangasarian, 2004) as a *feature selection* method. The common theme of these methods is that the solution \hat{w} of (3.5) is sparse, that is, most of its components are zero. The nonzero components of \hat{w} identify informative features for representing the given data. Indeed, Corollary 3.1 establishes that there exists an optimal vector \hat{w} with at most $\min(n, m + 1)$ non zero components.

We note that the equivalence between the functionals E_G and V_Φ described in Corollary 3.1 when the loss function L is the hinge loss used in support vector machines is also described in Bach, Lanckriet and Jordan (2004). Moreover, a method similar to problem (3.2) has been recently proposed in statistics under the name of *cosso*, where it has been proposed as a generalization of the *lasso* (see Lin & Zhang, (2003) and Lee et al. (2004) for related results). A specific instance of the above environment is provided by ANOVA kernels (see Wahba, 1990; Lin & Zhang, 2003 for a detailed discussion).

We now present an alternate derivation of Eq. (3.1). To this end, we require the following result of Aronszajn concerning the norm induced by a sum of reproducing kernels (Aronszajn, 1950, Section 7).

Theorem 3.1. *If $\{G_r, r \in \mathbb{N}_n\} \subseteq \mathcal{A}(\mathcal{X})$ and $K = \sum_{r \in \mathbb{N}_n} G_r$ then, for every $f \in \mathcal{H}_K$, we have that*

$$\|f\|_K^2 = \min \left\{ \sum_{r \in \mathbb{N}_n} \|f_r\|_{G_r}^2 : f = \sum_{r \in \mathbb{N}_n} f_r, f_s \in \mathcal{H}_{G_s}, s \in \mathbb{N}_n \right\}.$$

Without elaborating on the technical details, we note, by Theorem 3.1 that

$$\begin{aligned} E_G &= \min \left\{ \sum_{j \in \mathbb{N}_m} L \left(y_j, \sum_{r \in \mathbb{N}_n} f_r(x_j) \right) + \mu J \left(\sum_{r \in \mathbb{N}_n} \frac{\|f_r\|_{G_r}^2}{\lambda_r} \right) : f_r \in \mathcal{H}_{G_r}, \lambda \in \mathbb{S}^n \right\} \\ &= \min \left\{ \sum_{j \in \mathbb{N}_m} L \left(y_j, \sum_{r \in \mathbb{N}_n} f_r(x_j) \right) + \mu J \left(\left(\sum_{r \in \mathbb{N}_n} \|f_r\|_{G_r} \right)^2 \right) : f_r \in \mathcal{H}_{G_r}, r \in \mathbb{N}_n \right\} \\ &= \min \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, \langle w, \Phi(x_j) \rangle_{\mathcal{W}}) + \mu J(\|w\|_{1, \mathcal{W}}^2) : w \in \mathcal{W}^n \right\} = V_\Phi, \end{aligned}$$

where the first equality follows by Theorem 3.1, the second follows by taking the minimum over $\lambda \in \mathbb{S}^n$ and the third equality uses the feature map representation of the function f_r , that is, $f_r = \langle w_r, \Phi_r \rangle_{\mathcal{W}}$, $r \in \mathbb{N}_n$, where $\Phi = (\Phi_r : r \in \mathbb{N}_n)$, and $\|f_r\|_{G_r} = \|w_r\|_{\mathcal{W}}$. Moreover, the optimal value of $\lambda = (\lambda_r : r \in \mathbb{N}_n)$ is given by

$$\lambda_r = \frac{\|w_r\|_{\mathcal{W}}}{\|w\|_{1, \mathcal{W}}}. \tag{3.6}$$

As mentioned above, this derivation requires the result of Aronszajn in Theorem 3.1 whose proof is elaborate and only applies easily to the case of finite number of kernels. The approach we use to prove the general result in Theorem 2.1 is self-contained and gives more information. Indeed, note that the alternate derivation after Theorem 3.1 only reveals Eq. (3.1) and does not provide information about the structure for the extremal solutions for the associated variational problems. Nevertheless, it suggests an interesting extension of that

equation. To this end, we choose $h \in \mathbb{R}_+$, set $\mathbb{S}^{n,h} := \{(\lambda_r : r \in \mathbb{N}_n) : \lambda_r \geq 0, \sum_{r \in \mathbb{N}_n} \lambda_r = 1\}$,

$$\mathcal{K}_h(\mathcal{G}) = \left\{ \sum_{r \in \mathbb{N}_n} \lambda_r G_r : \lambda \in \mathbb{S}^{n,h} \right\},$$

$$E_{G,h} = \inf \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, f(x_j)) + \mu J(\|f\|_K^2) : f \in \mathcal{H}_K, K \in \mathcal{K}_h(\mathcal{G}) \right\}$$

and

$$E_{\Phi,p} = \min \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, \langle w, \Phi(x_j) \rangle_{\mathcal{W}}) + \mu J(\|w\|_{p,\mathcal{W}}^2) : w \in \mathcal{W}^n \right\}. \tag{3.7}$$

Proposition 3.1. *If $h \in \mathbb{R}_+$ and $p = \frac{2h}{h+1}$ then $E_{G,h} = E_{\Phi,p}$.*

The proof of this proposition follows from Theorem 3.1, the following lemma, whose proof can be found in the appendix of Micchelli and Pontil (2005) and the same technique used above to give an alternate proof of Eq. (3.1).

Lemma 3.1. *If $h \geq 0$, $p := \frac{2h}{h+1}$, and $a = (a_r : r \in \mathbb{N}_n) \in \mathbb{R}^n$ then*

$$\min \left\{ \left(\sum_{r \in \mathbb{N}_n} \frac{|a_r|^2}{\lambda_r} \right)^{\frac{1}{2}} : \lambda = (\lambda_r : r \in \mathbb{N}_n) \in \mathbb{S}^{n,h} \right\} = \|a\|_p^2$$

and the equality occurs for

$$\lambda_r := \left(\frac{|a_r|}{\|a\|_p} \right)^{2-p}, \quad r \in \mathbb{N}_n. \tag{3.8}$$

We hope on a future occasion, to use this alternate approach to discover the structure of the optimal solutions for $E_{G,h}$ and $V_{\Phi,p}$.

4 Single feature kernels and \mathcal{L}^1 regularization

In this section, we consider another case of our main result in Section 2 corresponding to the choice $\mathcal{W} = \mathbb{R}$. Equivalently, the kernels in \mathcal{G} are all expressed as a *single feature*. We have already observed in the previous section that in this case, under the additional assumption that Θ is a finite set, problem (3.2) reduces to problem (3.5) which is a type of \mathcal{L}^1 regularization problem. An analogous observation is summarized in the corollary below in the general case that Θ is any compact set. We note that in this case $C(\Theta, \mathbb{R}) = C(\Theta)$.

Corollary 4.1. *Under the hypothesis above, we have that $E_G = V_\Phi$ where*

$$E_G = \inf \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, f(x_j)) + \mu J(\|f\|_K^2) : f \in \mathcal{H}_K, K \in \mathcal{K}(\mathcal{G}) \right\}$$

and

$$V_\Phi = \min \left\{ \sum_{j \in \mathbb{N}_m} L(y_j, B(\Phi_j)) + \mu J(\|B\|^2) : B \in C^*(\Theta) \right\}.$$

Moreover, there exist a finitely supported measure $\hat{p} \in \mathcal{M}(\Theta)$ with at most $m + 1$ atoms and a vector $\hat{c} \in \mathbb{R}^m$ such that the kernel

$$\hat{K} = \int_{\Theta} G(\theta) d\hat{p}(\theta)$$

is an optimal kernel, the function

$$f_{\hat{K}} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \hat{K}_j$$

is an optimal function and the linear functional $\hat{B} \in C^*(\Theta)$ defined, for any $\Psi \in C(\Theta)$, as

$$\hat{B}(\Psi) = \int_{\Theta} \hat{\Gamma}(\theta) \Psi(\theta) d\hat{p}(\theta)$$

is an optimal linear functional, where the function $\hat{\Gamma} = \sum_{j \in \mathbb{N}_m} \hat{c}_j \Phi_j$ has the property, for every atom θ of \hat{p} , that $|\hat{\Gamma}(\theta)| = \|\hat{\Gamma}\|_\infty = \|\hat{B}\|$.

Note that we can rewrite the linear functional \hat{B} as a finitely supported signed measure, namely,

$$\hat{B} = \sum_{j \in \mathbb{N}_k} \hat{\gamma}_j \delta(\theta_j - \cdot),$$

where $\{\theta_i : i \in \mathbb{N}_k\} \subseteq \Theta$ are the atoms of \hat{p} , $k \leq m + 1$, $\hat{\gamma}_i = \hat{\lambda}_i \|\hat{\Gamma}\|_\infty \text{sgn} \hat{\Gamma}(\theta_i)$ and for each $\theta \in \Theta$, we interpret $\delta(\theta - \cdot)$ as the delta function concentrated at θ . Moreover, we can alternatively express the optimal function $f_{\hat{K}}$ as a linear combination of the features evaluated at the atoms of \hat{p} , that is

$$f_{\hat{K}} = \sum_{j \in \mathbb{N}_k} \hat{\gamma}_j \Phi(\theta_j, \cdot).$$

We view the feature space variational problem appearing in the above corollary as the analog of \mathcal{L}^1 regularization. We shall now change our perspective and explain in detail what we have in mind. Our intention is also to have this discussion encompass an \mathcal{L}^p extension of the variational problem above for $p \in (1, \infty)$. To this end, we describe the necessary terminology and notation to cover this case too.

The appropriate context for this discussion is a measurable space Θ with finite measure ν and $\mathcal{L}^p(\Theta, \nu)$ the space of functions $\omega : \Theta \rightarrow \mathbb{R}$ with norm $\|\omega\|_p$ defined, for $p \in [1, \infty)$, as

$$\|\omega\|_p := \left(\int_{\Theta} |\omega(\theta)|^p d\nu(\theta) \right)^{\frac{1}{p}}.$$

We wish to learn a function in $\mathcal{L}^p(\Theta, \nu)$ based on a finite set of linear functionals of it, that is, we have a set of examples of the form $\{(M_j, y_j) : j \in \mathbb{N}_m\}$ where M_j are linear functionals in $\mathcal{L}^r(\Theta, \nu)$, where $\frac{1}{p} + \frac{1}{r} = 1$, and $y_j \in \mathbb{R}$ is a noisy measurement of $M_j(\omega)$ from the unknown target function ω . Furthermore, we assume that the linear functionals $\{M_j : j \in \mathbb{N}_m\}$ are linearly independent.

To estimate ω we consider the problem of minimizing the functional $E_p : \mathcal{L}^p(\Theta, \nu) \rightarrow \mathbb{R}$ defined, for $\omega \in \mathcal{L}^p(\Theta, \nu)$, as

$$E_p(\omega) := q(M(\omega), \|\omega\|_p) \quad (4.1)$$

over its domain, where q is an admissible regularization function and $M : \mathcal{L}^p(\Theta, \nu) \rightarrow \mathbb{R}^m$ is the linear operator defined, for $\omega \in \mathcal{L}^p(\Theta, \nu)$, as $M(\omega) = (M_j(\omega) : j \in \mathbb{N}_m)$. Recall, for $p \in [1, \infty)$, that the linear functionals M_j can be expressed as

$$M_j(\omega) = \int_{\Theta} \varphi_j(\theta)\omega(\theta)d\nu(\theta),$$

where the function $\varphi_j \in \mathcal{L}^r(\Theta, \nu)$ and $\frac{1}{p} + \frac{1}{r} = 1$ (see, for example, Royden, 1964, p. 103). A special case of this setup is covered by the regularization error functional

$$E_p(\omega) := \sum_{j \in \mathbb{N}_m} L(y_j, M_j(\omega)) + \mu \|\omega\|_p^p, \quad \omega \in \mathcal{L}^p(\Theta, \nu), \quad (4.2)$$

where $L : \mathbb{R} \rightarrow \mathbb{R}_+$ is some prescribed loss function and μ is a positive parameter.

As an example of the above we let $N : \Theta \times \Theta \rightarrow \mathbb{R}$ be a prescribed continuous function and $\mathcal{N} : \mathcal{L}^p(\Theta, \nu) \rightarrow C(\Theta)$ the associated integral operator, that is, for $\omega \in \mathcal{L}^p(\Theta, \nu)$, we define

$$\mathcal{N}\omega(\cdot) = \int_{\Theta} N(\cdot, \theta)\omega(\theta)d\nu(\theta).$$

We introduce a linear space of functions $T := \text{range}\mathcal{N}$. We assume \mathcal{N} is one-to-one and observe that the norm of $h \in T$ defined as $\|h\| := \|\omega\|_p$ where $h = \mathcal{N}\omega$ makes T a Banach space.

We choose $\varphi_j := N_j$, $j \in \mathbb{N}_m$, where $N_j(\cdot) := N(\theta_j, \cdot)$, and express the regularization function (4.1) in the form

$$E_p(\omega) = q(I_{\theta}(h), \|h\|), \quad (4.3)$$

where $h = \mathcal{N}\omega$ and $\theta = \{\theta_j : j \in \mathbb{N}_m\}$ is a prescribed set of inputs. Clearly, minimizing the left hand side of equation above over $\omega \in \mathcal{L}^p(\Theta, \nu)$ is equivalent to minimize the right hand side of this equation over $h \in T$.

In general, T is not a Hilbert space. However, for the special case that $p = 2$, T is a reproducing kernel Hilbert space with inner product defined as $\langle h, h' \rangle = \int_{\Theta} \omega(\theta)\omega'(\theta)d\nu(\theta)$, where $h = \mathcal{N}\omega, h' = \mathcal{N}\omega'$ and the reproducing kernel K is given by

$$K(\theta, \theta') = \int_{\Theta} N(\theta, s)N(\theta', s)d\nu(s), \quad \theta, \theta' \in \Theta.$$

Indeed, to see that K is the reproducing kernel of T we observe that the above formula means, for $\theta \in \Theta$, that $K(\theta, \cdot) = \mathcal{N}(N(\theta, \cdot))$ and, so, if $h = \mathcal{N}\omega$, by definition we have that

$$\langle K(\theta, \cdot), h \rangle = \int_{\Theta} N(\theta, s)\omega(s)d\nu(s) = h(\theta).$$

In this special case, it is well-known that the unique minimizer \hat{h} of the functional in the right hand side of Eq. (4.3) has the form

$$\hat{h} = \sum_{j \in \mathbb{N}_m} c_j K_j, \tag{4.4}$$

where $K_j = K(\theta_j, \cdot)$ and the vector of coefficients $c = (c_j : j \in \mathbb{N}_m) \in \mathbb{R}^m$ is obtained by substituting formula (4.4) for \hat{h} into the right hand side of Eq. (4.3) and minimizing over c . Note that, if $\hat{h} = \mathcal{N}\hat{\omega}$ then $\hat{\omega}$ is given by

$$\hat{\omega} = \sum_{j \in \mathbb{N}_m} c_j N_j$$

or, equivalently, recalling the definition of φ_j in this example, we have that $\hat{\omega} = \sum_{j \in \mathbb{N}_m} c_j \varphi_j$.

Let us turn our attention to the case that $p = 1$. We begin our discussion by observing that, in general, E_1 does not have a minimum on $\mathcal{L}^1(\Theta, \nu)$. We illustrate this point with the following example.

Example 1. Let $d\theta$ be the Lebesgue measure on $[0, 1]$ and consider the problem of minimizing the functional

$$W(\omega; \alpha) = \alpha \left| 1 - \int_0^1 \omega(\theta)\theta d\theta \right| + \int_0^1 |\omega(\theta)|d\theta$$

over $\omega \in \mathcal{L}^1([0, 1], d\theta)$, where α is a nonnegative number. Call the value of this minimum $W_{\infty}(\alpha) := \inf\{W(\omega; \alpha) : \omega \in \mathcal{L}^1([0, 1], d\theta)\}$. With minimal effort it follows that

$$W_{\infty}(\alpha) = \begin{cases} 1, & \alpha > 1 \\ \alpha, & \alpha \in (0, 1]. \end{cases}$$

Moreover, the minimum does not exist in $L([0, 1], d\theta)$ but does exist as a distribution $\hat{\omega}$ given, for $\theta \in [0, 1]$ by

$$\hat{\omega}(\theta) = \begin{cases} \delta(\theta - 1), & \alpha > 1 \\ 0, \text{ a.e.}, & \alpha \in (0, 1]. \end{cases}$$

If we embed $\mathcal{L}^1([0, 1], d\theta)$ into $C^*([0, 1])$ and minimize $W(\cdot; \alpha)$ over $C^*([0, 1])$ the minimum exists and is given as above. Likewise, in general, the minimum of functional E_1 defined in (4.2) exists in $C^*(\Theta)$. Indeed, this corresponds to the feature space variational problem treated in Corollary 4.1. Keep in mind that our description of the \mathcal{L}^p regularization above takes the point of view of learning ω from the data $M(\omega)$. However, in the case $p = 1$, from our remarks above we *also* view it from the feature space perspective presented in Corollary 4.1.

We complete this digression by describing a representer theorem for the $\mathcal{L}^p(\Theta, \nu)$ regularization when $1 < p < \infty$. Before doing so, we think it is advantageous to present another proof for the \mathcal{L}^1 regularization case independent of the general theorem presented in Section 2.

Proposition 4.1. *There exist an integer $k \leq m + 1$, a set $\{\theta_j : j \in \mathbb{N}_k\} \subseteq \Theta$ and $\hat{\lambda} = (\hat{\lambda}_j : j \in \mathbb{N}_k) \in \mathbb{S}^k$ such that*

$$\hat{B} = \sum_{j \in \mathbb{N}_k} \hat{\lambda}_j \delta(\theta_j - \cdot)$$

is a minimizer of the regularization functional E_1 above. Moreover, there is a vector $\hat{c} = (\hat{c}_j : j \in \mathbb{N}_m)$ satisfying the constraint that $(\hat{c}, \hat{y}) = 1$, where $\hat{y} := (\hat{B}(\varphi_j) : j \in \mathbb{N}_m)$, such that the function $\hat{\Gamma} := \sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j$ has the property that, for every $j \in \mathbb{N}_k$,

$$|\hat{\Gamma}(\theta_j)| = \|\hat{\Gamma}\|_\infty = \min \left\{ \left\| \sum_{j \in \mathbb{N}_m} d_j \varphi_j \right\|_\infty : d \in \mathbb{R}^m, (d, \hat{y}) = 1 \right\}.$$

Proof: The existence of a minimum of E_1 over $C^*(\Theta)$ follows from weak-*compactness in the unit ball in $C^*(\Theta)$ (see Royden, 1964, p. 173). If \hat{B} is a minimizer of E_1 , we set $\hat{y}_j = \hat{B}(\varphi_j)$, $j \in \mathbb{N}_m$, choose any $d \in \mathbb{R}^m$ and note that, for every $B \in C^*(\Theta)$ such that $B(\varphi_j) = \hat{y}_j$, $j \in \mathbb{N}_m$, that

$$(d, \hat{y}) = \sum_{j \in \mathbb{N}_m} d_j B(\varphi_j) = B \left(\sum_{j \in \mathbb{N}_m} d_j \varphi_j \right) \leq \|B\| \left\| \sum_{j \in \mathbb{N}_m} d_j \varphi_j \right\|_\infty.$$

Consequently, we have that

$$\|B\| \geq \frac{\sum_{j \in \mathbb{N}_m} d_j \hat{y}_j}{\left\| \sum_{j \in \mathbb{N}_m} d_j \varphi_j \right\|_\infty} \geq \sigma^{-1},$$

where we have defined

$$\sigma := \min \left\{ \left\| \sum_{j \in \mathbb{N}_m} d_j \varphi_j \right\|_\infty : \sum_{j \in \mathbb{N}_m} d_j \hat{y}_j = 1 \right\}. \tag{4.5}$$

We observe that the variational problem (4.5) has a minimum because the function $U : \mathbb{R}^m \rightarrow \mathbb{R}_+$ defined for each $d = (d_j : j \in \mathbb{N}_m)$ by $U(d) := \|\sum_{j \in \mathbb{N}_m} d_j \varphi_j\|_\infty$ is continuous, homogeneous and nonzero for $d \neq 0$. Hence, it tends to infinity as $d \rightarrow \infty$.

Let $\hat{c} \in \mathbb{R}^m$ be a minimizer of U . Therefore, this vector is characterized by the fact that the right directional derivative of U at \hat{c} in all directions $a \in \mathbb{R}^m$ such that $(a, \hat{y}) = 1$ is nonnegative. We denote this derivative by $U'_+(\hat{c}; a)$ which is given by

$$U'_+(\hat{c}; a) = \max \left\{ \left(\sum_{j \in \mathbb{N}_m} a_j \varphi_j(\theta) \right) \operatorname{sgn} \left(\sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j(\theta) \right) : \theta \in \Theta(\hat{c}) \right\},$$

where the set $\Theta(\hat{c})$ is defined as

$$\Theta(\hat{c}) := \left\{ \theta : \theta \in \Theta, \left| \sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j(\theta) \right| = \left\| \sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j \right\|_\infty \right\}.$$

For each $\theta \in \Theta$, we define the vector $z(\theta) := (\varphi_j(\theta) \operatorname{sgn}(\sum_i \hat{c}_i \varphi_i(\theta)) : j \in \mathbb{N}_m) \in \mathbb{R}^m$ and the set of vectors $Z(\hat{c}) := \{z(\theta) : \theta \in \Theta(\hat{c})\} \subseteq \mathbb{R}^m$.

The condition that \hat{c} is a solution to problem (4.5) means, for all $a \in \mathbb{R}^m$ satisfying $(a, \hat{y}) = 0$, that

$$\max\{(z, a) : z \in Z(\hat{c})\} \geq 0.$$

Clearly, $Z(\hat{c})$ is a bounded subset of \mathbb{R}^m . Therefore, its closed convex hull $A := \overline{co(Z(\hat{c}))}$ is compact. We claim that A intersects the line spanned by the vector \hat{y} . Indeed, if this is not true then there exists a hyperplane $H := \{d : d \in \mathbb{R}^m, (\beta, d) + \alpha = 0\}$, $\alpha \in \mathbb{R}$, $\beta \in \mathbb{R}^m$, which strictly separates A from the set $\{\rho \hat{y} : \rho \in \mathbb{R}\}$ (see, for example, Royden, 1964, p. 176). In other words, we must have that

$$(\beta, \rho \hat{y}) + \alpha > 0, \quad \rho \in \mathbb{R}$$

and

$$(\beta, z) + \alpha \leq 0, \quad z \in Z(\hat{c}).$$

The first condition implies that $\alpha > 0$ and, so we conclude that

$$\max\{(\beta, z) : z \in Z(\hat{c})\} < 0$$

which contradicts our hypothesis that \hat{c} is a minimum of U . Hence, we have that $\rho_0 \hat{y} \in A$ for some $\rho_0 \in \mathbb{R}$. By the Caratheodory theorem (see, for example, Borwein & Lewis, 2000, Ch. 2), every vector in A can be expressed as a convex combination of at most $m + 1$ vectors in $Z(\hat{c})$. In particular, there exists ρ_0 such that

$$\rho_0 \hat{y} = \int_{\Theta} z(\theta) d\hat{p}(\theta), \tag{4.6}$$

where \hat{p} is a probability measure with k atoms, that is,

$$\hat{p} = \sum_{j \in \mathbb{N}_k} \gamma_j \delta(\cdot - \theta_j)$$

k is at most $m + 1$, $\{\theta_j : j \in \mathbb{N}_k\} \subseteq \Theta(\hat{c})$, $\gamma_j \geq 0$, $j \in \mathbb{N}_m$ and $\sum_{j \in \mathbb{N}_m} \gamma_j = 1$. Taking the inner product of both sides of Eq. (4.6) with c we conclude that

$$\rho_0 = \left\| \sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j \right\|_{\infty} = \|g\|_{\infty} > 0,$$

where $g = \sum_{j \in \mathbb{N}_k} \hat{c}_j \varphi_j$. Next, we introduce the linear functional $\hat{B} : C(\Theta) \rightarrow \mathbb{R}$ defined for $\omega \in C(\Theta)$ as $\hat{B}(\omega) = \int_{\Theta} \hat{\Gamma}(\theta) \omega(\theta) d\hat{p}(\theta)$ where

$$\hat{\Gamma} = \frac{1}{\rho_0} \text{sgn}(g) \hat{p} = \sum_{j \in \mathbb{N}_k} \hat{\lambda}_j \delta(\theta_j - \cdot),$$

where we have defined $\hat{\lambda}_j = \frac{\gamma_j}{\rho_0}$. The result follows by noting that $\hat{B}(\varphi_j) = \hat{y}_j$ and $\|\hat{B}\| = \|\hat{\Gamma}\|_{\infty}^{-1}$. □

5 Regularization in \mathcal{L}^p spaces

In this section, we provide a representation result for the minimizer of functional (4.1) when $p \in (1, \infty)$.

Proposition 5.1. *If the function $q : \mathbb{R}^m \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is admissible and $p \in (1, \infty)$ then there exists a minimizer $\hat{\omega} \in \mathcal{L}^p(\Theta, \nu)$ of functional (4.1), given by the form*

$$\hat{\omega} = \left(\left| \sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j \right| \right)^{r-1} \text{sign} \left(\sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j \right), \tag{5.1}$$

where $\hat{c} = (\hat{c}_j : j \in \mathbb{N}_m) \in \mathbb{R}^m$.

Proof: E_p has a minimum in $\mathcal{L}^p(\Theta, \nu)$ since q is an admissible regularization function and the unit ball in $\mathcal{L}^p(\Theta, \nu)$ is weakly compact, as it is well known (Royden, 1964, p. 173). Let $\hat{\omega}$ a minimizer of E_p , define data $\hat{y}_j = M_j(\hat{\omega})$, $j \in \mathbb{N}_m$ and choose $\tilde{\omega}$ as the solution to the minimal norm interpolation problem

$$\min\{\|\omega\|_p : \omega \in \mathcal{L}^p(\Theta, \nu), M_j(\omega) = \hat{y}_j, j \in \mathbb{N}_m\}. \tag{5.2}$$

Clearly $\tilde{\omega}$ is also a minimizer of E_p and, for any $d \in \mathbb{R}^m$, we have that

$$(d, \hat{y}) = \sum_{j \in \mathbb{N}_m} d_j M_j(\tilde{\omega}) = \int_{\Theta} \tilde{\omega}(\theta) \sum_{j \in \mathbb{N}_m} d_j \varphi_j(\theta) d\nu(\theta) \leq \|\tilde{\omega}\|_p \left\| \sum_{j \in \mathbb{N}_m} d_j \varphi_j \right\|_r,$$

where the last inequality follows by Hölder’s inequality. Consequently, we have that

$$\|\tilde{\omega}\|_p \geq \frac{\sum_{j \in \mathbb{N}_m} d_j \hat{y}_j}{\left\| \sum_{j \in \mathbb{N}_m} d_j \varphi_j \right\|_r} \geq \sigma^{-1},$$

where we have defined

$$\sigma = \min \left\{ \left\| \sum_{j \in \mathbb{N}_m} d_j \varphi_j \right\|_r : \sum_{j \in \mathbb{N}_m} d_j \hat{y}_j = 1 \right\}. \tag{5.3}$$

Let $\hat{c} \in \mathbb{R}^m$ be a minimizer of (5.3), define $\varphi := \sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j$, set $\tilde{\omega} = \|\varphi\|^{-r} |\varphi|^{r-1} \text{sgn}(\varphi)$ and note that

$$\|\tilde{\omega}\|_p = \|\varphi\|^{-r} \left(\int_{\Theta} |\varphi(\theta)|^{p(r-1)} d\nu(\theta) \right)^{\frac{1}{p}} = \|\varphi\|^{-r} \|\varphi\|^{\frac{r}{p}} = \|\varphi\|_r^{-1}.$$

This proves the claimed result. □

In order to compute the coefficient vector $\hat{c} = (\hat{c}_j : j \in \mathbb{N}_m)$ in Eq. (5.1) we substitute this equation for $\hat{\omega}$ in the right hand side of Eq. (4.1) obtaining the function

$$\hat{E}_p(\hat{c}) := E_p \left(\left| \sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j \right|^{r-1} \text{sgn} \left(\sum_{j \in \mathbb{N}_m} \hat{c}_j \varphi_j \right) \right), \quad \hat{c} \in \mathbb{R}^m$$

and then optimize \hat{E}_p over the vector $\hat{c} \in \mathbb{R}^m$. Unfortunately, the function \hat{E}_p is not, in general, convex. In fact, for the square loss $A(y, t) = \frac{1}{2}(y - t)^2$, $t \in \mathbb{R}$ with $\Theta = [0, 1]$, $m = 1$, $y_1 = 1$ and $\varphi = 1$ we have, for every $p \neq 2$, that

$$\hat{E}_p(\hat{c}) = (1 - |\hat{c}|^{r-1} \text{sgn}(\hat{c}))^2 + \gamma |\hat{c}|^{r-1}.$$

This function is not convex when $\gamma < 1$.

6 Proof of the main result

In this section, we present the proof of our main result in Theorem 2.1. We divide the proof in two parts. In the first part we establish that

$$E_G \geq V_\Phi. \tag{6.1}$$

Recall Eq. (2.3) and that, for every $K \in \mathcal{K}(\mathcal{G})$ we defined f_K to be any function in \mathcal{H}_K such that

$$E(K) := \min\{Q(f, K) : f \in \mathcal{H}_K\} = Q(f_K, K).$$

For every $y \in \mathbb{R}^m$ we use the notation $\rho(K, y)$ to denote the minimum norm squared of all functions $f \in \mathcal{H}_K$ which interpolate y at \mathbf{x} , that is, $I_{\mathbf{x}}f = y$ and set

$$\rho(K, y) := \min \{ \|f\|_K^2 : f \in \mathcal{H}_K, I_{\mathbf{x}}f = y \}. \tag{6.2}$$

As we remarked earlier f_K may not be unique and, hence, we are not certain of its structure. To overcome this difficulty we introduce the vector $y_K := I_{\mathbf{x}}f_K$ and let $g_K \in \mathcal{H}_K$ solve the minimal norm interpolation problem

$$\|g_K\|_K^2 = \rho(K, y_K).$$

Hence, $g_K = I_{\mathbf{x}}^*c_K$ for a unique $c_K \in \mathbb{R}^m$ identified by the linear equation $I_{\mathbf{x}}g_K = y_K$. This follows from the so-called representer theorem see, for example, Micchelli & Pontil, 2005 and references therein. Consequently, we have that

$$E(K) = q(I_{\mathbf{x}}f_K, \|f_K\|_K^2) = q(y_K, \|f_K\|_K^2) \geq q(y_K, \|g_K\|_K^2).$$

Note that in the last step we used property 3 of an acceptable regularization error function (see Definition 2.1).

Since $K \in \mathcal{K}(\mathcal{G})$, there exists $p \in \mathcal{M}(\Theta)$ such that

$$K = \int_{\Theta} G(\theta)dp(\theta).$$

We observe, for any $x \in \mathcal{X}$, that

$$g_K(x) = \int_{\Theta} \langle \Gamma_K(\theta), \Phi(\theta, x) \rangle dp(\theta),$$

where $\Gamma_K := \sum_{j \in \mathbb{N}_m} c_{K,j} \Phi_j$ and $c_K := (c_{K,j} : j \in \mathbb{N}_m)$. This computation suggests that we introduce the linear functional B_K defined, for every $\Psi \in C(\Theta, \mathcal{W})$, as

$$B_K(\Psi) = \int_{\Theta} \langle \Gamma_K(\theta), \Psi(\theta) \rangle dp(\theta). \tag{6.3}$$

Therefore, with these observations we obtain that

$$y_K = I_{\mathbf{x}}g_K = D_{\mathbf{x}}(B_K), \tag{6.4}$$

where the linear operator $D : C^*(\Theta, \mathcal{W}) \rightarrow \mathbb{R}^m$ was defined earlier, below Eq. (2.5). Also, observe, for any $\Psi \in C(\Theta, \mathcal{W})$ by the Cauchy-Schwarz’s inequality used twice (once in \mathcal{W} and then another time in $\mathcal{L}^2(\Theta, dp)$), that

$$\begin{aligned} |B_K(\Psi)| &\leq \int_{\Theta} \|\Gamma_K(\theta)\|_{\mathcal{W}} \|\Psi(\theta)\|_{\mathcal{W}} dp(\theta) \\ &\leq \left(\int_{\Theta} \|\Gamma_K(\theta)\|_{\mathcal{W}}^2 dp(\theta) \right)^{\frac{1}{2}} \left(\int_{\Theta} \|\Psi(\theta)\|_{\mathcal{W}}^2 dp(\theta) \right)^{\frac{1}{2}} \end{aligned}$$

and, so, we get the following inequality for the norm of B_K ,

$$\|B_K\|^2 \leq \int_{\Theta} \|\Gamma_K(\theta)\|_{\mathcal{V}}^2 dp(\theta). \tag{6.5}$$

A straightforward computation shows that

$$\int_{\Theta} \|\Gamma_K(\theta)\|_{\mathcal{V}}^2 dp(\theta) = \|g_K\|_K^2. \tag{6.6}$$

Consequently, combining Eqs. (6.5) and (6.6) we have demonstrated that

$$\|B_K\| \leq \|g_K\|_K. \tag{6.7}$$

Now, we observe from Eqs. (6.4) and (6.7), and the definition of the functional V_{Φ} , for any $K \in \mathcal{K}(\mathcal{G})$, that

$$\begin{aligned} E(K) &= E(y_K, f_K) \geq q(y_K, \|g_K\|_K^2) \\ &= q(D_{\mathbf{x}}(B_K), \|g_K\|_K^2) \geq q(D_{\mathbf{x}}(B_K), \|B_K\|^2) \geq V_{\Phi}. \end{aligned}$$

Since this lower bound for $E(K)$ holds for any $K \in \mathcal{K}(\mathcal{G})$, we have proved that

$$E_G \geq V_{\Phi}.$$

To show the reverse inequality we use a result from Micchelli and Pontil (2005). For the convenience of the reader we describe it in detail. To this end, we recall some notation used there. Before, we used $\rho(K, y)$ for the minimum norm squared of all functions $f \in \mathcal{H}_K$ which interpolate the data at y , see Eq. (6.2). The infimum of this quantity over all $K \in \mathcal{K}(\mathcal{G})$ will be denoted by $\rho(\mathcal{K}(\mathcal{G}), y)$. Moreover, since $\mathcal{K}(\mathcal{G}) \subseteq \mathcal{A}_+$, for each $y \in \mathbb{R}^m$ there is a unique $c_K \in \mathbb{R}^m$ such that

$$(c_K, I_{\mathbf{x}} I_{\mathbf{x}}^* c_K) = \|I^* c_K\|_K^2 = \rho(K, y)$$

and $I_{\mathbf{x}} I_{\mathbf{x}}^* c_K = y$. Since $I_{\mathbf{x}} I_{\mathbf{x}}^* = K_{\mathbf{x}}$ we also have that $c_K = K_{\mathbf{x}}^{-1} y$. For the statement of the theorem below we introduce the vector $\hat{c}_K := \rho^{-1}(K, y) c_K$ which has the property that $(\hat{c}_K, y) = 1$.

Theorem 6.1. *If Θ is a compact Hausdorff topological space and $G : \Theta \rightarrow \mathcal{A}_+(\mathcal{X})$ is continuous then there exists a kernel $\hat{K} = \int_{\Theta} G(\theta) d\hat{p}(\theta) \in \mathcal{K}(\mathcal{G})$ such that \hat{p} is a discrete probability measure in $\mathcal{M}(\Theta)$ with at most $m + 1$ atoms. Moreover, for $\hat{c} := \hat{c}_{\hat{K}}$ and any atom $\theta \in \Theta$ of \hat{p} , we have that*

$$(\hat{c}, G_{\mathbf{x}}(\theta)\hat{c}) = \max\{(\hat{c}, G_{\mathbf{x}}(\theta)\hat{c}) : \theta \in \Theta\}, \tag{6.8}$$

$$\rho(\mathcal{K}(\mathcal{G}), y) = \rho(\hat{K}, y) = (y, \hat{K}_{\mathbf{x}}^{-1} y) \tag{6.9}$$

and for every $c \in \mathbb{R}^m$ with $(c, y) = 1$ and every $K \in \mathcal{K}(\mathcal{G})$ there holds

$$(\hat{c}, K_{\mathbf{x}}\hat{c}) \leq (\hat{c}, \hat{K}_{\mathbf{x}}\hat{c}) \leq (c, \hat{K}_{\mathbf{x}}c). \tag{6.10}$$

Returning to the proof of Theorem 2.1, we observe by weak*-compactness in $C^*(\Theta, \mathcal{W})$ that there is a solution to the problem of minimizing $V_\Phi : C^*(\Theta, \mathcal{W}) \rightarrow \mathbb{R}_+$ over its domain (see, for example Royden, 1964, p. 173). We call the minimum \hat{B} and set $\hat{y} := D_x(\hat{B})$. Hence, by definition we have that

$$V_\Phi = q(\hat{y}, \|\hat{B}\|^2).$$

To estimate this quantity from below, we consider the problem

$$\gamma := \min \left\{ \left\| \sum_{j \in \mathbb{N}_m} c_j \Phi_j \right\|_{\infty, \mathcal{W}}^2 : c \in \mathbb{R}^m, (c, \hat{y}) = 1 \right\}.$$

Vector-valued problems of this type were considered in Micchelli (1992). Note that the minimum exists because the functions $\Phi_j, j \in \mathbb{N}_m$ were assumed to be linearly independent over Θ . For the problem at hand we note that

$$\begin{aligned} \gamma &= \min \left\{ \max \left\{ \left\| \sum_{j \in \mathbb{N}_m} c_j \Phi_j(\theta) \right\|^2 : \theta \in \Theta \right\} : c \in \mathbb{R}^m, (c, \hat{y}) = 1 \right\} \\ &= \min \left\{ \max \left\{ \int_{\Theta} \left\| \sum_{j \in \mathbb{N}_m} c_j \Phi_j(\theta) \right\|^2 dp(\theta) : p \in \mathcal{M}(\Theta) \right\} : c \in \mathbb{R}^m, (c, \hat{y}) = 1 \right\} \\ &= \min\{\max\{(c, K_x c) : K \in \mathcal{K}(\mathcal{G})\} : c \in \mathbb{R}^m, (c, \hat{y}) = 1\} = \rho^{-1}(\mathcal{K}(\mathcal{G}), y), \end{aligned}$$

where we recall that $K_x = (K(x_i, x_j) : i, j \in \mathbb{N}_m)$. Thus, by Theorem 6.1 there is a discrete probability measure $\hat{p} \in \mathcal{M}(\Theta)$ of support $\leq m + 1$ and a vector $\hat{c} \in \mathbb{R}^m$ such that for all $c \in \mathbb{R}^m, K \in \mathcal{K}(\mathcal{G})$, there holds the inequality

$$(\hat{c}, K_x, \hat{c}) \leq (\hat{c}, \hat{K}_x \hat{c}) \leq (c, \hat{K}_x c),$$

where

$$\hat{K} = \int_{\Theta} G(\theta) d\hat{p}(\theta).$$

Moreover, for each atom θ of \hat{p} we have that

$$\rho^{-1}(\mathcal{K}(\mathcal{G}), y) = (\hat{c}, G_x(\theta)\hat{c}) = \max\{(\hat{c}, G_x(\theta)\hat{c}) : \theta \in \Theta\} \tag{6.11}$$

and the kernel \hat{K} , has the property that $\hat{K}_x \hat{c} = \gamma \hat{y}$. As before, we let $\hat{\Gamma} := \sum_{j \in \mathbb{N}_m} \hat{c}_j \Phi_j$ and observe that $\|\hat{\Gamma}(\theta)\|_{\mathcal{W}}^2 = (\hat{c}, G_x(\theta)\hat{c})$. Consequently, for each atom θ of \hat{p} , we have that $\|\hat{\Gamma}(\theta)\|_{\mathcal{W}} = \|\hat{\Gamma}\|_{\infty, \mathcal{W}} = \sqrt{\gamma}$.

Now, let us consider the linear functional $F \in C^*(\Theta, \mathcal{W})$ defined for each $\Psi \in C^*(\Theta, \mathcal{W})$ as

$$F(\Psi) = \gamma^{-1} \int_{\Theta} \langle \hat{\Gamma}(\theta), \Psi(\theta) \rangle_{\mathcal{W}} d\hat{p}(\theta).$$

As before in the proof of (6.1), we conclude that $\|F\| \leq \gamma^{-1} \|\hat{\Gamma}\|_{\infty, \mathcal{W}}$. However, in the present circumstance, since $F(\hat{\Gamma}) = \gamma^{-1} \|\hat{\Gamma}\|_{\infty, \mathcal{W}}^2$ we additionally obtain, by the definition of the norm of F , that $\gamma^{-1} \|\hat{\Gamma}\|_{\infty, \mathcal{W}}^2 \leq \|F\| \|\hat{\Gamma}\|_{\infty, \mathcal{W}}$. In other words, we obtain that $\gamma^{\frac{1}{2}} = \gamma^{-1} \|\hat{\Gamma}\|_{\infty, \mathcal{W}} = \|F\|$.

To proceed further, we shall demonstrate that F is the minimal norm interpolant to the data \hat{y} . First, we observe that the functional F interpolate the data at \hat{y} . Indeed, since $K_x \hat{c} = \gamma y$ we have that

$$D_x(F) = \gamma^{-1} \hat{K}_x^{-1} \hat{c} = \hat{y}.$$

Now, let $B \in C^*(\Theta, \mathcal{W})$ be any linear functional which interpolates the data, that is, $D_x(B) = \hat{y}$. Therefore, we get that

$$1 = (\hat{c}, \hat{y}) = B \left(\sum_{j \in \mathbb{N}_m} \hat{c}_j \Phi_j \right) = B(\hat{\Gamma}) \leq \|B\| \|\hat{\Gamma}\|_{\infty, \mathcal{W}} = \sqrt{\gamma} \|B\|$$

from which we obtain that $\|F\| \leq \|B\|$. In other words, we have established, as anticipated above, that

$$\|F\| = \min\{\|B\| : B \in C^*(\Theta, \mathcal{W}), D_x(B) = \hat{y}, j \in \mathbb{N}_m\}.$$

Next, we introduce the function $\hat{g} := \gamma^{-1} I_x^* \hat{c}$ and observe as before that $\|\hat{g}\|_{\hat{K}} = \|F\|$. Finally, since $I_x \hat{g} = \hat{y}$, we conclude that

$$V_\Phi = q(\hat{y}, \|\hat{B}\|^2) \geq q(\hat{y}, \|F\|^2) = q(I_x \hat{g}, \|\hat{g}\|_{\hat{K}}^2) \geq E_G.$$

Thus, we have that $V_\Phi = E_G$, g_K is a minimizer for $E(\hat{K})$, \hat{K} is optimal for E and F is optimal for V_Φ .

7 Summary

We have presented an equivalence between the problem of learning a kernel within a prescribed set of continuously parameterized kernels studied by Argyriou, Micchelli, and Pontil (2005) and Micchelli and Pontil (2005) and a feature space reformulation of it. This leads us to study a variational regularization problem in the dual space of all continuous functions with values in the Hilbert space associated with the features maps. This equivalence requires only weak conditions on the form of the regularization error function. Not only does it establish that these variational problems are the same but, also, it provides a choice of the optimal solutions to both problems. In particular, it generalizes some results from Argyriou, Micchelli, and Pontil (2005) which required the loss function to be differentiable.

Furthermore, we demonstrated that the problem of learning the kernel, which has been investigated extensively in the literature (see Argyriou, Micchelli, & Pontil, 2005; Bach, Lanckriet, & Jordan, 2004; Bousquet & Herrmann, 2003; Cristianini et al., 2002; Lanckriet et al., 2004; Lee et al., 2004; Lin & Zhang, 2003; Ong, Smola, & Williamson, 2003; Wu, Ying, & Zhou, to appear), in special cases reduces to \mathcal{L}^p regularization (Micchelli and Pontil, 2004). This connection highlights the importance of studying regularization in a non-Hilbert space framework in machine learning. Indeed, special cases of the feature space problem

have been widely studied under the names *lasso* (Tibshirani, 1996), *basis pursuit denoising* (Chen, Donoho, & Saunders, 1998) and, recently, as the *cosso* method (Lin & Zhang, 2003).

There are a number of issues related to the work presented in this paper which would be valuable to explore. For example, how does the form of the optimal solutions to the variational problems evolve with μ ? Our results in Section 2 and the subsequent comments in Section 3 say that there always exists a solution which uses at most $m + 1$ nonzero kernels or features. Do the number of non-zero components diminish with μ , as was seen in Micchelli and Pinkus (1994) for special cases? Finally, the study of generalization error bounds for the methods presented in this paper would definitely be of interest (see Micchelli et al., 2005 for recent progress on this issue). Of course, the central challenge not addressed here is the practical implementation and numerical validation of the methods presented here.

References

- Argyriou, A., Micchelli, C. A., & Pontil, M. (2005). Learning convex combinations of continuously parameterized basic kernels. In *Proc. 18-th Annual Conference on Learning Theory (COLT'05)*, Bertinoro, Italy.
- Aronszajn, N. (1950). Theory of reproducing kernels. *Trans. Amer. Math. Soc.*, 686, 337–404.
- Bach, F. R., Lanckriet, G. R. G., & Jordan, M. I. (2004). Multiple kernels learning, conic duality, and the SMO algorithm. In *Proc. of the Int. Conf. on Machine Learning (ICML'04)*.
- Borwein, J. M., & Lewis, A. S. (2000). *Convex analysis and nonlinear optimization. Theory and examples*. CMS (Canadian Mathematical Society) Springer-Verlag, New York.
- Bousquet, O., & Herrmann, D. J. L. (2003). On the complexity of learning the kernel matrix. *Advances in Neural Information Processing Systems*, 15.
- Chen, S. S., Donoho, D. L., Saunders, M. A. (1998). Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20(1), 33–61.
- Cristianini, N., Shawe-Taylor, J., Elisseeff, A., Kandola, J. S. (2002). On kernel-target alignment. In T. G. Dietterich, S. Becker, & Z. Ghahramani (Eds.), *Advances in Neural Information Processing Systems*, vol. 14.
- Fung, G. M., & Mangasarian, O. L. (2004). A feature selection Newton method for support vector machine classification. *Comput. Optim. Appl.*, 28(2), 185–202.
- Gunn, S. R., & Kandola, J. S. (2002) Structural modelling with sparse kernels *Machine Learning*, 48(1), 137–163.
- Herbster, M. (2004). Relative loss bounds and polynomial-time predictions for the K-LMS-NET algorithm. In *Proc. of the 15th Int. Conference on Algorithmic Learning Theory*.
- Lanckriet, G. R. G., Cristianini, N., Bartlett, P., El Ghaoui, L., Jordan, M. I. (2004). Learning the kernel matrix with semi-definite programming. *J. of Machine Learning Research*, 5, 27–72.
- Lee, Y., Kim, Y., Lee, S., & Koo, J.-Y. (2004). Structured multicategory support vector machine with ANOVA decomposition. Technical Report No. 743, Department of Statistics, The Ohio State University.
- Lin, Y., & Zhang, H. H. (2003). Component selection and smoothing in smoothing spline analysis of variance models—cosso. Institute of Statistics Mimeo Series 2556, NCSU.
- Micchelli, C. A. (1992). Curves from variational principles. *Mathematical Modeling and Numerical Analysis*, 26, 77–93.
- Micchelli, C. A., & Pinkus, A. (1994). Variational problems arising from balancing different error criteria. *Rendiconti di Matematica, Serie VII*, 14, 37–86.
- Micchelli, C. A., & Pontil, M. (2004). A function representation for learning in Banach spaces. In *Proc. of the 17th Annual Conference on Learning Theory (COLT'04)*, Banff, Alberta.
- Micchelli, C. A. & Pontil, M. (2005). On learning vector-valued functions. *Neural Computation*, 17, 177–204.
- Micchelli, C. A., & Pontil, M. (2005). Learning the kernel function via regularization. *J. of Machine Learning Research*, 6, 1099–1125.
- Micchelli, C. A., Pontil, M., Wu, Q., & Zhou, D. X. (2005). Error bounds for learning the kernel. Research Note 05/09, Dept. of Computer Science, University College London.
- Ong, C. S., Smola, A. J., & Williamson, R. C. (2003). Hyperkernels. In S. Becker, S. Thrun, K. Obermayer (Eds.), *Advances in Neural Information Processing Systems*, vol. 15, MIT Press, Cambridge, MA.

- Royden, H. L. (1964). *Real analysis*, 2nd edition. Macmillan Publishing Company, New York.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Royal Statist. Soc. B*, 58, 267–288.
- Wahba, G. (1990). *Splines models for observational data*. Series in Applied Mathematics, vol. 59, SIAM, Philadelphia.
- Wu, Q., Ying, Y., & Zhou, D. X. Multi-kernel regularization classifiers. *J. of Complexity* (to appear).