



Technical Report No. 156

A Review of Kernel Methods in Machine Learning

Thomas Hofmann¹, Bernhard Schölkopf²,
Alexander J. Smola³

December 2006

¹Darmstadt University of Technology, Department of Computer Science and Fraunhofer Institute for Integrated Publication and Information Systems, Darmstadt, Germany, email: thomas.hofmann@ipsi.fraunhofer.de

²Max Planck Institute für Biologische Cybernetics, 72076 Tübingen, Germany, email: bs@tuebingen.mpg.de

³Statistical Machine Learning Program, National ICT Australia, and Computer Science Laboratory, Australian National University, Canberra, Australia, email: Alex.Smola@nicta.com.au

Note: instead of this TR, please cite the updated journal version:

T. Hofmann, B. Schölkopf, A. Smola: Kernel methods in machine learning.
Annals of Statistics (in press).

A Review of Kernel Methods in Machine Learning

Thomas Hofmann, Bernhard Schölkopf, Alexander J. Smola

Abstract. We review recent methods for learning with positive definite kernels. All these methods formulate learning and estimation problems as linear tasks in a reproducing kernel Hilbert space (RKHS) associated with a kernel. We cover a wide range of methods, ranging from simple classifiers to sophisticated methods for estimation with structured data.

(AMS 2000 subject classifications: primary - 30C40 Kernel functions and applications; secondary - 68T05 Learning and adaptive systems. — Key words: machine learning, reproducing kernels, support vector machines, graphical models)

1 Introduction

Over the last ten years, estimation and learning methods utilizing positive definite kernels have become rather popular, particularly in machine learning. Since these methods have a stronger mathematical slant than earlier machine learning methods (e.g., neural networks), there is also significant interest in the statistics and mathematics community for these methods. The present review aims to summarize the state of the art on a conceptual level. In doing so, we build on various sources (including [Vapnik, 1998, Burges, 1998, Cristianini and Shawe-Taylor, 2000, Herbrich, 2002] and in particular [Schölkopf and Smola, 2002]), but we also add a fair amount of more recent material which helps unifying the exposition. We have not had space to include proofs; they can be found either in the long version of the present paper (see Hofmann et al. [2006]), in the references given, or in the above books.

The main idea of all the described methods can be summarized in one paragraph. Traditionally, theory and algorithms of machine learning and statistics has been very well developed for the linear case. Real world data analysis problems, on the other hand, often requires nonlinear methods to detect the kind of dependencies that allow successful prediction of properties of interest. By using a positive definite kernel, one can sometimes have the best of both worlds. The kernel corresponds to a dot product in a (usually high dimensional) feature space. In this space, our estimation methods are linear, but as long as we can formulate everything in terms of kernel evaluations, we never explicitly have to compute in the high dimensional feature space.

The paper has three main sections: Section 2 deals with fundamental properties of *kernels*, with special emphasis on (conditionally) positive definite kernels and their characterization. We give concrete examples for such kernels and discuss kernels and reproducing kernel Hilbert spaces in the context of regularization. Section 3 presents various approaches for estimating dependencies and analyzing data that make use of kernels. We provide an overview of the problem formulations as well as their solution using convex programming techniques. Finally, Section 4 examines the use of reproducing kernel Hilbert spaces as a means to define statistical models, the focus being on structured, multidimensional responses. We also show how such techniques can be combined with Markov networks as a suitable framework to model dependencies between response variables.




Figure 1: A simple geometric classification algorithm: given two classes of points (depicted by ‘o’ and ‘+’), compute their means c_+, c_- and assign a test input x to the one whose mean is closer. This can be done by looking at the dot product between $x - c$ (where $c = (c_+ + c_-)/2$) and $w := c_+ - c_-$, which changes sign as the enclosed angle passes through $\pi/2$. Note that the corresponding decision boundary is a hyperplane (the dotted line) orthogonal to w (from [Schölkopf and Smola, 2002]).

2 Kernels

2.1 An Introductory Example

Suppose we are given empirical data

$$(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathcal{Y}. \quad (1)$$

Here, the domain \mathcal{X} is some nonempty set that the *inputs* (the predictor variables) x_i are taken from; the $y_i \in \mathcal{Y}$ are called *targets* (the response variable). Here and below, $i, j \in [n]$, where we use the notation $[n] := \{1, \dots, n\}$.

Note that we have not made any assumptions on the domain \mathcal{X} other than it being a set. In order to study the problem of learning, we need additional structure. In learning, we want to be able to *generalize* to unseen data points. In the case of binary pattern recognition, given some new input $x \in \mathcal{X}$, we want to predict the corresponding $y \in \{\pm 1\}$ (more complex output domains \mathcal{Y} will be treated below). Loosely speaking, we want to choose y such that (x, y) is in some sense *similar* to the training examples. To this end, we need similarity measures in \mathcal{X} and in $\{\pm 1\}$. The latter is easier, as two target values can only be identical or different. For the former, we require a function

$$k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}, \quad (x, x') \mapsto k(x, x') \quad (2)$$

satisfying, for all $x, x' \in \mathcal{X}$,

$$k(x, x') = \langle \Phi(x), \Phi(x') \rangle, \quad (3)$$

where Φ maps into some dot product space \mathcal{H} , sometimes called the *feature space*. The similarity measure k is usually called a *kernel*, and Φ is called its *feature map*.

The advantage of using such a kernel as a similarity measure is that it allows us to construct algorithms in dot product spaces. For instance, consider the following simple classification algorithm, where $\mathcal{Y} = \{\pm 1\}$. The idea is to compute the means of the two classes in the feature space, $c_+ = \frac{1}{n_+} \sum_{\{i: y_i = +1\}} \Phi(x_i)$, and $c_- = \frac{1}{n_-} \sum_{\{i: y_i = -1\}} \Phi(x_i)$, where n_+ and n_- are the number of examples with positive and negative target values, respectively. We then assign a new point $\Phi(x)$ to the class whose mean is closer to it. This leads to the prediction rule

$$y = \text{sgn}(\langle \Phi(x), c_+ \rangle - \langle \Phi(x), c_- \rangle + b) \quad (4)$$

with $b = \frac{1}{2} (\|c_-\|^2 - \|c_+\|^2)$. Substituting the expressions for c_\pm yields

$$y = \operatorname{sgn} \left(\frac{1}{n_+} \sum_{\{i:y_i=+1\}} \underbrace{\langle \Phi(x), \Phi(x_i) \rangle}_{k(x,x_i)} - \frac{1}{n_-} \sum_{\{i:y_i=-1\}} \underbrace{\langle \Phi(x), \Phi(x_i) \rangle}_{k(x,x_i)} + b \right), \quad (5)$$

where $b = \frac{1}{2} \left(\frac{1}{n_-^2} \sum_{\{(i,j):y_i=y_j=-1\}} k(x_i, x_j) - \frac{1}{n_+^2} \sum_{\{(i,j):y_i=y_j=+1\}} k(x_i, x_j) \right)$.

Let us consider one well-known special case of this type of classifier. Assume that the class means have the same distance to the origin (hence $b = 0$), and that $k(\cdot, x)$ is a density for all $x \in \mathcal{X}$. If the two classes are equally likely and were generated from two probability distributions that are estimated

$$p_+(x) := \frac{1}{n_+} \sum_{\{i:y_i=+1\}} k(x, x_i), \quad p_-(x) := \frac{1}{n_-} \sum_{\{i:y_i=-1\}} k(x, x_i), \quad (6)$$

then (5) is the estimated Bayes decision rule, plugging in the estimates p_+ and p_- for the true densities.

The classifier (5) is closely related to the *Support Vector Machine (SVM)* that we will discuss below. It is linear in the feature space (4), while in the input domain, it is represented by a kernel expansion (5). In both cases, the decision boundary is a hyperplane in the feature space; however, the normal vectors (for (4), $w = c_+ - c_-$) are usually rather different.¹ The normal vector not only characterizes the alignment of the hyperplane, its length can also be used to construct tests for the equality of the two class-generating distributions [Borwardt et al., 2006].

2.2 Positive Definite Kernels

We have required that a kernel satisfy (3), i.e., correspond to a dot product in some dot product space. In the present section, we show that the class of kernels that can be written in the form (3) coincides with the class of positive definite kernels. This has far-reaching consequences. There are examples of positive definite kernels which can be evaluated efficiently even though they correspond to dot products in infinite dimensional dot product spaces. In such cases, substituting $k(x, x')$ for $\langle \Phi(x), \Phi(x') \rangle$, as we have done in (5), is crucial. In the machine learning community, this substitution is called the *kernel trick*.

Definition 1 (Gram matrix) Given a kernel k and inputs $x_1, \dots, x_n \in \mathcal{X}$, the $n \times n$ matrix

$$K := (k(x_i, x_j))_{ij} \quad (7)$$

is called the Gram matrix (or kernel matrix) of k with respect to x_1, \dots, x_n .

Definition 2 (Positive definite matrix) A real $n \times n$ symmetric matrix K_{ij} satisfying

$$\sum_{i,j} c_i c_j K_{ij} \geq 0 \quad (8)$$

for all $c_i \in \mathbb{R}$ is called positive definite. If equality in (8) only occurs for $c_1 = \dots = c_n = 0$ then we shall call the matrix strictly positive definite.

Definition 3 (Positive definite kernel) Let \mathcal{X} be a nonempty set. A function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ which for all $n \in \mathbb{N}$, $x_i \in \mathcal{X}$, $i \in [n]$ gives rise to a positive definite Gram matrix is called a positive definite kernel. A function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ which for all $n \in \mathbb{N}$ and distinct $x_i \in \mathcal{X}$ gives rise to a strictly positive definite Gram matrix is called a strictly positive definite kernel.

¹As an aside, note that if we normalize the targets such that $\hat{y}_i = y_i / |\{j : y_j = y_i\}|$, in which case the \hat{y}_i sum to zero, then $\|w\|^2 = \langle K, \hat{y} \hat{y}^\top \rangle_F$, where $\langle \cdot, \cdot \rangle_F$ is the Frobenius dot product. If the two classes have equal size, then up to a scaling factor involving $\|K\|_2$ and n , this equals the *kernel-target alignment* defined in [Cristianini et al., 2002].

Occasionally, we shall refer to positive definite kernels simply as a *kernels*. Note that for simplicity we have restricted ourselves to the case of real valued kernels. However, with small changes, the below will also hold for the complex valued case.

Since $\sum_{i,j} c_i c_j \langle \Phi(x_i), \Phi(x_j) \rangle = \left\langle \sum_i c_i \Phi(x_i), \sum_j c_j \Phi(x_j) \right\rangle \geq 0$, kernels of the form (3) are positive definite for any choice of Φ . In particular, if \mathcal{X} is already a dot product space, we may choose Φ to be the identity. Kernels can thus be regarded as generalized dot products. Whilst they are not generally bilinear, they share important properties with dot products, such as the Cauchy-Schwartz inequality:

Proposition 4 *If k is a positive definite kernel, and $x_1, x_2 \in \mathcal{X}$, then*

$$k(x_1, x_2)^2 \leq k(x_1, x_1) \cdot k(x_2, x_2). \quad (9)$$

Proof The 2×2 Gram matrix with entries $K_{ij} = k(x_i, x_j)$ is positive definite. Hence both its eigenvalues are nonnegative, and so is their product, K 's determinant, i.e.,

$$0 \leq K_{11}K_{22} - K_{12}K_{21} = K_{11}K_{22} - K_{12}^2. \quad (10)$$

Substituting $k(x_i, x_j)$ for K_{ij} , we get the desired inequality. \blacksquare

2.2.1 Construction of the Reproducing Kernel Hilbert Space

We now define a map from \mathcal{X} into the space of functions mapping \mathcal{X} into \mathbb{R} , denoted as $\mathbb{R}^{\mathcal{X}}$, via

$$\Phi : \mathcal{X} \rightarrow \mathbb{R}^{\mathcal{X}} \text{ where } x \mapsto k(\cdot, x). \quad (11)$$

Here, $\Phi(x) = k(\cdot, x)$ denotes the function that assigns the value $k(x', x)$ to $x' \in \mathcal{X}$.

We next construct a dot product space containing the images of the inputs under Φ . To this end, we first turn it into a vector space by forming linear combinations

$$f(\cdot) = \sum_{i=1}^n \alpha_i k(\cdot, x_i). \quad (12)$$

Here, $n \in \mathbb{N}$, $\alpha_i \in \mathbb{R}$ and $x_i \in \mathcal{X}$ are arbitrary.

Next, we define a dot product between f and another function $g(\cdot) = \sum_{j=1}^{n'} \beta_j k(\cdot, x'_j)$ (with $n' \in \mathbb{N}$, $\beta_j \in \mathbb{R}$ and $x'_j \in \mathcal{X}$) as

$$\langle f, g \rangle := \sum_{i=1}^n \sum_{j=1}^{n'} \alpha_i \beta_j k(x_i, x'_j). \quad (13)$$

To see that this is well-defined although it contains the expansion coefficients and points, note that $\langle f, g \rangle = \sum_{j=1}^{n'} \beta_j f(x'_j)$. The latter, however, does not depend on the particular expansion of f . Similarly, for g , note that $\langle f, g \rangle = \sum_{i=1}^n \alpha_i g(x_i)$. This also shows that $\langle \cdot, \cdot \rangle$ is bilinear. It is symmetric, as $\langle f, g \rangle = \langle g, f \rangle$. Moreover, it is positive definite, since positive definiteness of k implies that for any function f , written as (12), we have

$$\langle f, f \rangle = \sum_{i,j=1}^n \alpha_i \alpha_j k(x_i, x_j) \geq 0. \quad (14)$$

Next, note that given functions f_1, \dots, f_p , and coefficients $\gamma_1, \dots, \gamma_p \in \mathbb{R}$, we have

$$\sum_{i,j=1}^p \gamma_i \gamma_j \langle f_i, f_j \rangle = \left\langle \sum_{i=1}^p \gamma_i f_i, \sum_{j=1}^p \gamma_j f_j \right\rangle \geq 0. \quad (15)$$

Here, the equality follows from the bi-linearity of $\langle \cdot, \cdot \rangle$, and the right hand inequality from (14).

By (15), $\langle \cdot, \cdot \rangle$ is a positive definite kernel, defined on our vector space of functions. For the last step in proving that it even is a dot product, we note that by (13), for all functions (12),

$$\langle k(\cdot, x), f \rangle = f(x), \text{ and in particular } \langle k(\cdot, x), k(\cdot, x') \rangle = k(x, x'). \quad (16)$$

By virtue of these properties, k is called a *reproducing kernel* [Aronszajn, 1950].

Due to (16) and (9), we have

$$|f(x)|^2 = |\langle k(\cdot, x), f \rangle|^2 \leq k(x, x) \cdot \langle f, f \rangle. \quad (17)$$

By this inequality, $\langle f, f \rangle = 0$ implies $f = 0$, which is the last property that was left to prove in order to establish that $\langle \cdot, \cdot \rangle$ is a dot product.

Skipping some details, we add that one can complete the space of functions (12) in the norm corresponding to the dot product, and thus gets a Hilbert space \mathcal{H} , called a *reproducing kernel Hilbert space (RKHS)*.

One can define an RKHS as a Hilbert space \mathcal{H} of functions on a set \mathcal{X} with the property that for all $x \in \mathcal{X}$ and $f \in \mathcal{H}$, the point evaluations $f \mapsto f(x)$ are continuous linear functionals (in particular, all point values $f(x)$ are well defined, which already distinguishes RKHS's from many L_2 Hilbert spaces). From the point evaluation functional, one can then construct the reproducing kernel using the Riesz representation theorem. The Moore-Aronszajn-Theorem [Aronszajn, 1950] states that for every positive definite kernel on $\mathcal{X} \times \mathcal{X}$, there exists a unique RKHS and vice versa.

There is an analogue of the kernel trick for distances rather than dot products, i.e., dissimilarities rather than similarities. This leads to the larger class of *conditionally positive definite kernels*. Those kernels are defined just like positive definite ones, with the one difference being that their Gram matrices need to satisfy (8) only subject to

$$\sum_{i=1}^n c_i = 0. \quad (18)$$

Interestingly, it turns out that many kernel algorithms, including SVMs and kernel PCA (see Section 3), can be applied also with this larger class of kernels, due to their being translation invariant in feature space [Schölkopf and Smola, 2002, Hein et al., 2005].

We conclude this section with a note on terminology. In the early years of kernel machine learning research, it was not the notion of positive definite kernels that was being used. Instead, researchers considered kernels satisfying the conditions of Mercer's theorem [Mercer, 1909], see e.g. Vapnik [1998] and Cristianini and Shawe-Taylor [2000]. However, whilst all such kernels do satisfy (3), the converse is not true. Since (3) is what we are interested in, positive definite kernels are thus the right class of kernels to consider.

2.2.2 Properties of Positive Definite Kernels

We begin with some closure properties of the set of positive definite kernels.

Proposition 5 *Below, k_1, k_2, \dots are arbitrary positive definite kernels on $\mathcal{X} \times \mathcal{X}$, where \mathcal{X} is a nonempty set.*

(i) *The set of positive definite kernels is a closed convex cone, i.e., (a) if $\alpha_1, \alpha_2 \geq 0$, then $\alpha_1 k_1 + \alpha_2 k_2$ is positive definite; and (b) If $k(x, x') := \lim_{n \rightarrow \infty} k_n(x, x')$ exists for all x, x' , then k is positive definite.*

(ii) *The pointwise product $k_1 k_2$ is positive definite.*

(iii) *Assume that for $i = 1, 2$, k_i is a positive definite kernel on $\mathcal{X}_i \times \mathcal{X}_i$, where \mathcal{X}_i is a nonempty set. Then the tensor product $k_1 \otimes k_2$ and the direct sum $k_1 \oplus k_2$ are positive definite kernels on $(\mathcal{X}_1 \times \mathcal{X}_2) \times (\mathcal{X}_1 \times \mathcal{X}_2)$.*

The proofs can be found in Berg et al. [1984]. We only give a short proof of (ii), a result due to Schur (see e.g., Berg and Forst [1975]). Denote by K, L the positive definite Gram matrices of k_1, k_2 with respect to some data set x_1, \dots, x_n . Being positive definite, we can express L as SS^\top in terms of an $n \times n$ matrix S (e.g., the positive definite square root of L). We thus have

$$L_{ij} = \sum_{m=1}^n S_{im} S_{jm}, \quad (19)$$

and therefore, for any $c_1, \dots, c_n \in \mathbb{R}$,

$$\sum_{i,j=1}^n c_i c_j K_{ij} L_{ij} = \sum_{i,j=1}^n c_i c_j K_{ij} \sum_{m=1}^n S_{im} S_{jm} = \sum_{m=1}^n \sum_{i,j=1}^n (c_i S_{im})(c_j S_{jm}) K_{ij} \geq 0, \quad (20)$$

where the inequality follows from (8).

It is reassuring that sums and products of positive definite kernels are positive definite. We will now explain that, loosely speaking, there are no other operations that preserve positive definiteness. To this end, let C denote the set of all functions $\psi: \mathbb{R} \rightarrow \mathbb{R}$ that map positive definite kernels to (conditionally) positive definite kernels (readers who are not interested in the case of conditionally positive definite kernels may ignore the term in parentheses). We define

$$C := \{\psi \mid k \text{ is a positive definite kernel} \Rightarrow \psi(k) \text{ is a (conditionally) positive definite kernel}\}.$$

Moreover, define

$$C' = \{\psi \mid \text{for any Hilbert space } \mathcal{F}, \psi(\langle x, x' \rangle_{\mathcal{F}}) \text{ is (conditionally) positive definite}\}$$

and

$$C'' = \{\psi \mid \text{for all } n \in \mathbb{N}: K \text{ is a positive definite } n \times n \text{ matrix} \Rightarrow \psi(K) \text{ is (conditionally) positive definite}\},$$

where $\psi(K)$ is the $n \times n$ matrix with elements $\psi(K_{ij})$.

Proposition 6 $C = C' = C''$

Proof Let $\psi \in C$. For any Hilbert space \mathcal{F} , $\langle x, x' \rangle_{\mathcal{F}}$ is positive definite, thus $\psi(\langle x, x' \rangle_{\mathcal{F}})$ is (conditionally) positive definite. Hence $\psi \in C'$.

Next, assume $\psi \in C'$. Consider, for $n \in \mathbb{N}$, an arbitrary positive definite $n \times n$ -matrix K . We can express K as the Gram matrix of some $x_1, \dots, x_n \in \mathbb{R}^n$, i.e., $K_{ij} = \langle x_i, x_j \rangle_{\mathbb{R}^n}$. Since $\psi \in C'$, we know that $\psi(\langle x, x' \rangle_{\mathbb{R}^n})$ is a (conditionally) positive definite kernel, hence in particular the matrix $\psi(K)$ is (conditionally) positive definite, thus $\psi \in C''$.

Finally, assume $\psi \in C''$. Let k be a positive definite kernel, $n \in \mathbb{N}$, $c_1, \dots, c_n \in \mathbb{R}$, and $x_1, \dots, x_n \in \mathcal{X}$. Then

$$\sum_{i,j} c_i c_j \psi(k(x_i, x_j)) = \sum_{i,j} c_i c_j (\psi(K))_{ij} \geq 0 \quad (21)$$

(for $\sum_i c_i = 0$), where the inequality follows from $\psi \in C''$. Therefore, $\psi(k)$ is (conditionally) positive definite, and hence $\psi \in C$. \blacksquare

The following proposition follows from a result of FitzGerald et al. [1995] for (conditionally) positive definite matrices; by Proposition 6, it also applies for (conditionally) positive definite kernels, and for functions of dot products. We state the latter case.

Proposition 7 Let $\psi : \mathbb{R} \rightarrow \mathbb{R}$. Then $\psi(\langle x, x' \rangle_{\mathcal{F}})$ is positive definite for any Hilbert space \mathcal{F} if and only if ψ is real entire of the form

$$\psi(t) = \sum_{n=0}^{\infty} a_n t^n \quad (22)$$

with $a_n \geq 0$ for $n \geq 0$.

Moreover, $\psi(\langle x, x' \rangle_{\mathcal{F}})$ is conditionally positive definite for any Hilbert space \mathcal{F} if and only if ψ is real entire of the form (22) with $a_n \geq 0$ for $n \geq 1$.

There are further properties of k that can be read off the coefficients a_n .

- Steinwart [2002a] showed that if all a_n are strictly positive, then the kernel of Proposition 7 is *universal* on every compact subset S of \mathbb{R}^d in the sense that its RKHS is dense in the space of continuous functions on S in the $\|\cdot\|_{\infty}$ norm. For support vector machines using universal kernels, he then shows (universal) consistency [Steinwart, 2002b]. Examples of universal kernels are (23) and (24) below.
- In Lemma 13, we will show that the a_0 term does not affect an SVM. Hence we infer that it is actually sufficient for consistency to have $a_n > 0$ for $n \geq 1$.
- Let I_{even} and I_{odd} denote the sets of even and odd indices i with the property that $a_i > 0$, respectively. Pinkus [2004] showed that for any \mathcal{F} , the kernel of Proposition 7 is strictly positive definite if and only if $a_0 > 0$ and neither I_{even} nor I_{odd} is finite.² Moreover, he states that the necessary and sufficient conditions for universality of the kernel are that in addition both $\sum_{i \in I_{\text{even}}} \frac{1}{i}$ and $\sum_{i \in I_{\text{odd}}} \frac{1}{i}$ diverge.

While the proofs of the above statements are fairly involved, one can make certain aspects plausible using simple arguments. For instance [Pinkus, 2004], suppose k is not strictly positive definite. Then we know that for some $c \neq 0$ we have $\sum_{ij} c_i c_j k(x_i, x_j) = 0$. This means that $\left\langle \sum_i c_i k(x_i, \cdot), \sum_j c_j k(x_j, \cdot) \right\rangle = 0$, implying that $\sum_i c_i k(x_i, \cdot) = 0$. For any $f = \sum_j \alpha_j k(x_j, \cdot)$, this implies that $\sum_i c_i f(x_i) = \left\langle \sum_i c_i k(x_i, \cdot), \sum_j \alpha_j k(x_j, \cdot) \right\rangle = 0$. This equality thus holds for any function in the RKHS. Therefore, the RKHS cannot lie dense in the space of continuous functions on \mathcal{X} , and k is thus not universal.

We thus know that if k is universal, it is necessarily strictly positive definite. The latter, in turn, implies that its RKHS \mathcal{H} is infinite dimensional, since otherwise the maximal rank of the Gram matrix of k with respect to a data set in \mathcal{H} would equal its dimensionality. If we assume that \mathcal{F} is finite dimensional, then infinite dimensionality of \mathcal{H} implies that infinitely many of the a_i in (22) must be strictly positive. If only finitely many of the a_i in (22) were positive, we could construct the feature space \mathcal{H} of $\psi(\langle x, x' \rangle)$ by taking finitely many tensor products and (scaled) direct sums of copies of \mathcal{F} , and \mathcal{H} would be finite.³ We conclude the section with an example of a kernel which is positive definite by Proposition 7. To this end, let \mathcal{X} be a dot product space. The power series expansion of $\psi(x) = e^x$ then tells us that

$$k(x, x') = e^{\frac{\langle x, x' \rangle}{\sigma^2}} \quad (23)$$

²One might be tempted to think that given some strictly positive definite kernel k with feature map Φ and feature space \mathcal{H} , we could set \mathcal{F} to equal \mathcal{H} and consider $\psi(\langle \Phi(x), \Phi(x') \rangle_{\mathcal{H}})$. In this case, it would seem that choosing $a_1 = 1$ and $a_n = 0$ for $n \neq 1$ should give us a strictly positive definite kernel which violates Pinkus' conditions. However, the latter kernel is only strictly positive definite on $\Phi(\mathcal{X})$, which is a weaker statement than it being strictly positive definite on all of \mathcal{H} .

³Recall that the dot product of $\mathcal{F} \otimes \mathcal{F}$ is the square of the original dot product, and the one of $\mathcal{F} \oplus \mathcal{F}'$, where \mathcal{F}' is another dot product space, is the sum of the original dot products.

is positive definite [Haussler, 1999]. If we further multiply k with the positive definite kernel $f(x)f(x')$, where $f(x) = e^{-\frac{\|x\|^2}{2\sigma^2}}$ and $\sigma > 0$, this leads to the positive definiteness of the Gaussian kernel

$$k'(x, x') = k(x, x')f(x)f(x') = e^{-\frac{\|x-x'\|^2}{2\sigma^2}}. \quad (24)$$

2.2.3 Properties of Positive Definite Functions

We now let $\mathcal{X} = \mathbb{R}^d$ and consider positive definite kernels of the form

$$k(x, x') = h(x - x'), \quad (25)$$

in which case h is called a *positive definite function*. The following characterization is due to Bochner [1933], see also Rudin [1962]. We state it in the form given by Wendland [2005].

Theorem 8 *A continuous function h on \mathbb{R}^d is positive definite if and only if there exists a finite nonnegative Borel measure μ on \mathbb{R}^d such that*

$$h(x) = \int_{\mathbb{R}^d} e^{-i\langle x, \omega \rangle} d\mu(\omega). \quad (26)$$

Whilst normally formulated for complex valued functions, the theorem also holds true for real functions. Note, however, that if we start with an arbitrary nonnegative Borel measure, its Fourier transform may not be real. Real valued positive definite functions are distinguished by the fact that the corresponding measures μ are symmetric.

We may normalize h such that $h(0) = 1$ (hence by (9) $|h(x)| \leq 1$), in which case μ is a probability measure and h is its characteristic function. For instance, if μ is a normal distribution of the form $(2\pi/\sigma^2)^{-d/2} e^{-\frac{\sigma^2\|\omega\|^2}{2}} d\omega$, then the corresponding positive definite function is the Gaussian $e^{-\frac{\|x\|^2}{2\sigma^2}}$, cf. (24).

Bochner's theorem allows us to interpret the similarity measure $k(x, x') = h(x - x')$ in the frequency domain. The choice of the measure μ determines which frequency components occur in the kernel. Since the solutions of kernel algorithms will turn out to be finite kernel expansions, the measure μ will thus determine which frequencies occur in the estimates, i.e., it will determine their regularization properties — more on that in Section 2.3.3 below.

A proof of the theorem can be found for instance in Rudin [1962]. One part of the theorem is easy to prove: if h takes the form (26), then

$$\sum_{ij} a_i a_j h(x_i - x_j) = \sum_{ij} a_i a_j \int_{\mathbb{R}^d} e^{-i\langle x_i - x_j, \omega \rangle} d\mu(\omega) = \int_{\mathbb{R}^d} \left| \sum_i a_i e^{-i\langle x_i, \omega \rangle} \right|^2 d\mu(\omega) \geq 0. \quad (27)$$

Bochner's theorem generalizes earlier work of Mathias, and has itself been generalized in various ways, e.g. by Schoenberg [1938]. An important generalization considers Abelian semigroups (e.g., Berg et al. [1984]). In that case, the theorem provides an integral representation of positive definite functions in terms of the semigroup's semicharacters. Further generalizations were given by Krein, for the cases of positive definite kernels and functions with a limited number of negative squares (see Stewart [1976] for further details and references).

As above, there are conditions that ensure that the positive definiteness becomes strict.

Proposition 9 [Wendland, 2005] *A positive definite function is strictly positive definite if the carrier of the measure in its representation (26) contains an open subset.*

This implies that the Gaussian kernel is strictly positive definite.

An important special case of positive definite functions, which includes the Gaussian, are *radial basis functions*. These are functions that can be written as $h(x) = g(\|x\|_2)$ for some function $g : [0, \infty[\rightarrow \mathbb{R}$. They have the property of being invariant under the Euclidean group. If we would like to have the additional property of compact support, which is computationally attractive in a number of large scale applications, the following result becomes relevant.

Proposition 10 [Wendland, 2005] *Assume that $g : [0, \infty[\rightarrow \mathbb{R}$ is continuous and compactly supported. Then $h(x) = g(\|x\|_2)$ cannot be positive definite on every \mathbb{R}^d .*

2.2.4 Examples of Kernels

We have already seen several instances of positive definite kernels, and now intend to complete our selection with a few more examples. In particular, we discuss polynomial kernels, convolution kernels, ANOVA expansions, and kernels on documents.

Polynomial Kernels From Proposition 5 it is clear that homogeneous polynomial kernels $k(x, x') = \langle x, x' \rangle^p$ are positive definite for $p \in \mathbb{N}$ and $x, x' \in \mathbb{R}^d$. By direct calculation we can derive the corresponding feature map [Poggio, 1975]

$$\langle x, x' \rangle^p = \left(\sum_{j=1}^d [x]_j [x']_j \right)^p = \sum_{j \in [d]^p} [x]_{j_1} \cdots [x]_{j_p} \cdot [x']_{j_1} \cdots [x']_{j_p} = \langle C_p(x), C_p(x') \rangle, \quad (28)$$

where C_p maps $x \in \mathbb{R}^d$ to the vector $C_p(x)$ whose entries are all possible p -th degree ordered products of the entries of x (note that $[d]$ is used as a shorthand for $\{1, \dots, d\}$). The polynomial kernel of degree p thus computes a dot product in the space spanned by all monomials of degree p in the input coordinates. Other useful kernels include the inhomogeneous polynomial,

$$k(x, x') = (\langle x, x' \rangle + c)^p \text{ where } p \in \mathbb{N} \text{ and } c \geq 0, \quad (29)$$

which computes all monomials up to degree p .

Spline Kernels It is possible to obtain spline functions as a result of kernel expansions [Smola, 1996, Vapnik et al., 1997] simply by noting that convolution of an even number of indicator functions yields a positive kernel function. Denote by I_X the indicator (or characteristic) function on the set X , and denote by \otimes the convolution operation, $(f \otimes g)(x) := \int_{\mathbb{R}^d} f(x')g(x' - x)dx'$. Then the B-spline kernels are given by

$$k(x, x') = B_{2p+1}(x - x') \text{ where } p \in \mathbb{N} \text{ with } B_{i+1} := B_i \otimes B_0. \quad (30)$$

Here B_0 is the characteristic function on the unit ball⁴ in \mathbb{R}^d . From the definition of (30) it is obvious that for odd m we may write B_m as inner product between functions $B_{m/2}$. Moreover, note that for even m , B_m is not a kernel.

Convolutions and Structures Let us now move to kernels defined on structured objects [Haussler, 1999, Watkins, 2000]. Suppose the object $x \in \mathcal{X}$ is composed of $x_p \in \mathcal{X}_p$, where $p \in [P]$ (note that the sets \mathcal{X}_p need not be equal). For instance, consider the string $x = ATG$, and $P = 2$. It is composed of the parts $x_1 = AT$ and $x_2 = G$, or alternatively, of $x_1 = A$ and $x_2 = TG$. Mathematically speaking, the set of “allowed” decompositions can be thought of as a *relation* $R(x_1, \dots, x_P, x)$, to be read as “ x_1, \dots, x_P constitute the composite object x .”

⁴Note that in \mathbb{R} one typically uses $I_{[-\frac{1}{2}, \frac{1}{2}]}$.

Hausler [1999] investigated how to define a kernel between composite objects by building on similarity measures that assess their respective *parts*; in other words, kernels k_p defined on $\mathcal{X}_p \times \mathcal{X}_p$. Define the *R-convolution* of k_1, \dots, k_P as

$$[k_1 \star \dots \star k_P](x, x') := \sum_{\bar{x} \in R(x), \bar{x}' \in R(x')} \prod_{p=1}^P k_p(\bar{x}_p, \bar{x}'_p), \quad (31)$$

where the sum runs over all possible ways $R(x)$ and $R(x')$ in which we can decompose x into $\bar{x}_1, \dots, \bar{x}_P$ and x' analogously.⁵ If there is only a finite number of ways, the relation R is called finite. In this case, it can be shown that the *R-convolution* is a valid kernel [Hausler, 1999].

ANOVA Kernels Specific examples of convolution kernels are Gaussians and ANOVA kernels [Wahba, 1990, Vapnik, 1998]. To construct an ANOVA kernel, we consider $\mathcal{X} = S^N$ for some set S , and kernels $k^{(i)}$ on $S \times S$, where $i = 1, \dots, N$. For $P = 1, \dots, N$, the *ANOVA kernel of order P* is defined as

$$k_P(x, x') := \sum_{1 \leq i_1 < \dots < i_P \leq N} \prod_{p=1}^P k^{(i_p)}(x_{i_p}, x'_{i_p}). \quad (32)$$

Note that if $P = N$, the sum consists only of the term for which $(i_1, \dots, i_P) = (1, \dots, N)$, and k equals the tensor product $k^{(1)} \otimes \dots \otimes k^{(N)}$. At the other extreme, if $P = 1$, then the products collapse to one factor each, and k equals the direct sum $k^{(1)} \oplus \dots \oplus k^{(N)}$. For intermediate values of P , we get kernels that lie in between tensor products and direct sums.

ANOVA kernels typically use some moderate value of P , which specifies the order of the interactions between attributes x_{i_p} that we are interested in. The sum then runs over the numerous terms that take into account interactions of order P ; fortunately, the computational cost can be reduced to $O(Pd)$ cost by utilizing recurrent procedures for the kernel evaluation. ANOVA kernels have been shown to work rather well in multi-dimensional SV regression problems [Stitson et al., 1999].

Brownian Bridge and Related Kernels The Brownian bridge kernel $\min(x, x')$ defined on \mathbb{R} likewise is a positive kernel. Note that one-dimensional linear spline kernels with an infinite number of nodes [Vapnik et al., 1997] for $x, x' \in \mathbb{R}_0^+$ are given by

$$k(x, x') = \frac{\min(x, x')^3}{3} + \frac{\min(x, x')^2 |x - x'|}{2} + 1 + xx'. \quad (33)$$

These kernels can be used as basis functions $k^{(n)}$ in ANOVA expansions. Note that it is advisable to use a $k^{(n)}$ which never or rarely takes the value zero, since a single zero term would eliminate the product in (32).

Bag of Words One way in which SVMs have been used for text categorization [Joachims, 2002] is the *bag-of-words* representation. This maps a given text to a sparse vector, where each component corresponds to a word, and a component is set to one (or some other number) whenever the related word occurs in the text. Using an efficient sparse representation, the dot product between two such vectors can be computed quickly. Furthermore, this dot product is by construction a valid kernel, referred to as a *sparse vector kernel*. One of its shortcomings, however, is that it does not take into account the word ordering of a document. Other sparse vector kernels are also conceivable, such as one that maps a text to the set of pairs of words that are in the same sentence [Joachims, 1998, Watkins, 2000], or those which look only at pairs of words within a certain vicinity with respect to each other [Sim, 2001].

⁵We use the convention that an empty sum equals zero, hence if either x or x' cannot be decomposed, then $(k_1 \star \dots \star k_P)(x, x') = 0$.

n -grams and Suffix Trees A more sophisticated way of dealing with string data was proposed [Watkins, 2000, Haussler, 1999]. The basic idea is as described above for general structured objects (31): Compare the strings by means of the substrings they contain. The more substrings two strings have in common, the more similar they are. The substrings need not always be contiguous; that said, the further apart the first and last element of a substring are, the less weight should be given to the similarity. Depending on the specific choice of a similarity measure it is possible to define more or less efficient kernels which compute the dot product in the feature space spanned by *all* substrings of documents.

Consider a finite alphabet Σ , the set of all strings of length n , Σ^n , and the set of all finite strings, $\Sigma^* := \cup_{n=0}^{\infty} \Sigma^n$. The length of a string $s \in \Sigma^*$ is denoted by $|s|$, and its elements by $s(1) \dots s(|s|)$; the concatenation of s and $t \in \Sigma^*$ is written st . Denote by

$$k(x, x') = \sum_s \#(x, s) \#(x', s) c_s$$

a string kernel computed from exact matches. Here $\#(x, s)$ is the number of occurrences of s in x and $c_s \geq 0$.

Vishwanathan and Smola [2004] provide an algorithm using suffix trees, which allows one to compute for arbitrary c_s the value of the kernel $k(x, x')$ in $O(|x| + |x'|)$ time and memory. Moreover, also $f(x) = \langle w, \Phi(x) \rangle$ can be computed in $O(|x|)$ time if preprocessing linear in the size of the support vectors is carried out. These kernels are then applied to function prediction (according to the gene ontology) of proteins using only their sequence information. Another prominent application of string kernels is in the field of splice form prediction and gene finding Ratsch et al. [2007].

For inexact matches of a limited degree, typically up to $\epsilon = 3$, and strings of bounded length a similar data structure can be built by explicitly generating a dictionary of strings and their neighborhood in terms of a Hamming distance [Leslie et al., 2002b,a]. These kernels are defined by replacing $\#(x, s)$ by a mismatch function $\#(x, s, \epsilon)$ which reports the number of approximate occurrences of s in x . By trading off computational complexity with storage (hence the restriction to small numbers of mismatches) essentially linear-time algorithms can be designed. Whether a general purpose algorithm exists which allows for efficient comparisons of strings with mismatches in linear time is still an open question. Tools from approximate string matching [Navarro and Raffinot, 1999, Cole and Hariharan, 2000, Gusfield, 1997] promise to be of help in this context.

Mismatch Kernels In the general case it is only possible to find algorithms whose complexity is linear in the lengths of the documents being compared, and the length of the substrings, i.e. $O(|x| \cdot |x'|)$ or worse. We now describe such a kernel with a specific choice of weights [Watkins, 2000, Cristianini and Shawe-Taylor, 2000].

Let us now form subsequences u of strings. Given an index sequence $\mathbf{i} := (i_1, \dots, i_{|u|})$ with $1 \leq i_1 < \dots < i_{|u|} \leq |s|$, we define $u := s(\mathbf{i}) := s(i_1) \dots s(i_{|u|})$. We call $l(\mathbf{i}) := i_{|u|} - i_1 + 1$ *the length of the subsequence in s* . Note that if \mathbf{i} is not contiguous, then $l(\mathbf{i}) > |u|$.

The feature space built from strings of length n is defined to be $\mathcal{H}_n := \mathbb{R}^{(\Sigma^n)}$. This notation means that the space has one dimension (or coordinate) for each element of Σ^n , labelled by that element (equiv- alently, we can think of it as the space of all real-valued functions on Σ^n). We can thus describe the feature map coordinate-wise for each $u \in \Sigma^n$ via

$$[\Phi_n(s)]_u := \sum_{\mathbf{i}:s(\mathbf{i})=u} \lambda^{l(\mathbf{i})}. \quad (34)$$

Here, $0 < \lambda \leq 1$ is a decay parameter: The larger the length of the subsequence in s , the smaller the respective contribution to $[\Phi_n(s)]_u$. The sum runs over all subsequences of s which equal u .

For instance, consider a dimension of \mathcal{H}_3 spanned (that is, labelled) by the string `asd`. In this case, we have $[\Phi_3(\text{Nasdaq})]_{\text{asd}} = \lambda^3$, while $[\Phi_3(\text{1ass das})]_{\text{asd}} = 2\lambda^5$.⁶ The kernel induced by the map Φ_n takes the form

$$k_n(s, t) = \sum_{u \in \Sigma^n} [\Phi_n(s)]_u [\Phi_n(t)]_u = \sum_{u \in \Sigma^n} \sum_{(\mathbf{i}, \mathbf{j}): s(\mathbf{i})=t(\mathbf{j})=u} \lambda^{l(\mathbf{i})} \lambda^{l(\mathbf{j})}. \quad (35)$$

To describe the actual computation of k_n , define

$$k'_i(s, t) := \sum_{u \in \Sigma^i} \sum_{(\mathbf{i}, \mathbf{j}): s(\mathbf{i})=t(\mathbf{j})=u} \lambda^{|\mathbf{s}|+|\mathbf{t}|-i_1-j_1+2}, \text{ for } i = 1, \dots, n-1. \quad (36)$$

Using $x \in \Sigma^1$, we then have the following recursions, which allow the computation of $k_n(s, t)$ for all $n = 1, 2, \dots$ (note that the kernels are symmetric):

$$\begin{aligned} k'_0(s, t) &= 1 \text{ for all } s, t \\ k'_i(s, t) &= 0 \text{ if } \min(|s|, |t|) < i \\ k_i(s, t) &= 0 \text{ if } \min(|s|, |t|) < i \\ k'_i(sx, t) &= \lambda k'_i(s, t) + \sum_{j: t_j=x} k'_{i-1}(s, t[1, \dots, j-1]) \lambda^{|t|-j+2}, \quad i = 1, \dots, n-1 \\ k_n(sx, t) &= k_n(s, t) + \sum_{j: t_j=x} k'_{n-1}(s, t[1, \dots, j-1]) \lambda^2 \end{aligned} \quad (37)$$

The string kernel k_n can be computed using dynamic programming, see Watkins [2000], Lodhi et al. [2000], and Durbin et al. [1998].

The above kernels on string, suffix-tree, mismatch and tree kernels have been used in sequence analysis. This includes applications in document analysis and categorization, spam filtering, function prediction in proteins, annotations of dna sequences for the detection of introns and exons, named entity tagging of documents, and the construction of parse trees.

Locality Improved Kernels It is possible to adjust kernels to the structure of spatial data. Recall the Gaussian RBF and polynomial kernels. When applied to an image, it makes no difference whether one uses as x the image or a version of x where all locations of the pixels have been permuted. This indicates that function space on \mathcal{X} induced by k does not take advantage of the *locality* properties of the data.

By taking advantage of the local structure, estimates can be improved. On biological sequences [Zien et al., 2000] one may assign more weight to the entries of the sequence close to the location where estimates should occur.

For images, local interactions between image patches need to be considered. One way is to use the *pyramidal* kernel [Schölkopf, 1997, DeCoste and Schölkopf, 2002]. It takes inner products between corresponding image patches, then raises the latter to some power p_1 , and finally raises their sum to another power p_2 . While the overall degree of this kernel is $p_1 p_2$, the first factor p_1 only captures short range interactions.

Tree Kernels We now discuss similarity measures on more structured objects. For trees Collins and Duffy [2001] propose a decomposition method which maps a tree x into its set of subtrees. The kernel between two trees x, x' is then computed by taking a weighted sum of all terms between both trees. In particular, Collins and Duffy [2001] show a quadratic time algorithm, i.e. $O(|x| \cdot |x'|)$ to compute this expression, where $|x|$ is the number of nodes of the tree. When restricting the sum to all proper rooted subtrees it is possible to reduce the computational cost to $O(|x| + |x'|)$ time by means of a tree to string conversion [Vishwanathan and Smola, 2004].

⁶In the first string, `asd` is a contiguous substring. In the second string, it appears twice as a non-contiguous substring of length 5 in `lass das`, the two occurrences are `lass das` and `lass das`.

Graph Kernels Graphs pose a twofold challenge: one may both design a kernel *on* vertices of them and also a kernel *between* them. In the former case, the graph itself becomes the object defining the metric between the vertices. See Gärtner [2003], Kashima et al. [2004], and Kashima et al. [2003] for details on the latter. In the following we discuss kernels *on* graphs.

Denote by $W \in \mathbb{R}^{n \times n}$ the adjacency matrix of a graph with $W_{ij} > 0$ if an edge between i, j exists. Moreover, assume for simplicity that the graph is undirected, that is $W^\top = W$ (see Zhou et al. [2005] for extensions to directed graphs). Denote by $L = D - W$ the graph Laplacian and by $\tilde{L} = \mathbf{1} - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$ the normalized graph Laplacian. Here D is a diagonal matrix with $D_{ii} = \sum_j W_{ij}$ denoting the degree of vertex i .

Fiedler [1973] showed that the second largest eigenvector of L approximately decomposes the graph into two parts according to their sign. The other large eigenvectors partition the graph into correspondingly smaller portions. L arises from the fact that for a function f defined on the vertices of the graph $\sum_{i,j} (f(i) - f(j))^2 = 2f^\top L f$.

Finally, Smola and Kondor [2003] show that under mild conditions and up to rescaling, L is the only quadratic permutation invariant form which can be obtained as a linear function of W .

Hence it is reasonable to consider kernel matrices K obtained from L (and \tilde{L}). Smola and Kondor [2003] suggest kernels $K = r(L)$ or $K = r(\tilde{L})$, which have desirable smoothness properties. Here $r : [0, \infty) \rightarrow [0, \infty)$ is a monotonically decreasing function. Popular choices include

$$r(\xi) = \exp(-\lambda\xi) \quad \text{diffusion kernel} \quad (38)$$

$$r(\xi) = (\xi + \lambda)^{-1} \quad \text{regularized graph Laplacian} \quad (39)$$

$$r(\xi) = (\lambda - \xi)^p \quad p\text{-step random walk} \quad (40)$$

where $\lambda > 0$ is chosen such as to reflect the amount of diffusion in (38), the degree of regularization in (39), or the weighting of steps within a random walk (40) respectively. Eq. (38) was proposed by Kondor and Lafferty [2002]. In Section 2.3.3 we will discuss the connection between regularization operators and kernels in \mathbb{R}^n . Without going into details, the function $r(\xi)$ describes the smoothness properties on the graph and L plays the role of the Laplace operator.

Kernels on Sets and Subspaces Whenever each observation x_i consists of a *set* of instances we may use a range of methods to capture the specific properties of these sets (for an overview, see Vishwanathan et al. [2006]):

- take the average of the elements of the set in feature space, that is, $\phi(x_i) = \frac{1}{n} \sum_j \phi(x_{ij})$ [Gärtner et al., 2002]. This yields good performance in the area of multi-instance learning.
- Jebara and Kondor [2003] extend the idea by dealing with distributions $p_i(x)$ such that $\phi(x_i) = \mathbf{E}[\phi(x)]$ where $x \sim p_i(x)$. They apply it to image classification with missing pixels.
- Alternatively, one can study angles enclosed by *subspaces* spanned by the observations [Wolf and Shashua, 2003, Martin, 2000, Cock and Moor, 2002]. In a nutshell, if U, U' denote the orthogonal matrices spanning the subspaces of x and x' respectively, then $k(x, x') = \det U^\top U'$.
- Vishwanathan et al. [2006] extend this to arbitrary compound matrices (i.e. matrices composed of subdeterminants of matrices) and dynamical systems. Their result exploits the Binet Cauchy theorem which states that compound matrices are a representation of the matrix group, i.e. $C_q(AB) = C_q(A)C_q(B)$. This leads to kernels of the form $k(x, x') = \text{tr} C_q(AB)$. Note that for $q = \dim A$ we recover the determinant kernel of Wolf and Shashua [2003].

Fisher Kernels Jaakkola and Haussler [1999] have designed kernels building on probability density models $p(x|\theta)$. Denote by

$$U_\theta(x) := -\partial_\theta \log p(x|\theta) \quad (41)$$

$$I := \mathbf{E}_x [U_\theta(x) U_\theta^\top(x)] \quad (42)$$

the Fisher scores and the Fisher information matrix respectively. Note that for maximum likelihood estimators $\mathbf{E}_x[U_\theta(x)] = 0$ and therefore I is the covariance of $U_\theta(x)$. The Fisher kernel is defined as

$$k(x, x') := U_\theta^\top(x)I^{-1}U_\theta(x') \text{ or } k(x, x') := U_\theta^\top(x)U_\theta(x') \quad (43)$$

depending on whether we study the normalized or the unnormalized kernel respectively.

In addition to that, it has several attractive theoretical properties: Oliver et al. [2000] show that estimation using the normalized Fisher kernel corresponds to estimation subject to a regularization on the $L_2(p(\cdot|\theta))$ norm.

Moreover, in the context of exponential families (see Section 5.1 for a more detailed discussion) where $p(x|\theta) = \exp(\langle \phi(x), \theta \rangle - g(\theta))$ we have

$$k(x, x') = [\phi(x) - \partial_\theta g(\theta)] [\phi(x') - \partial_\theta g(\theta)]. \quad (44)$$

for the unnormalized Fisher kernel. This means that up to centering by $\partial_\theta g(\theta)$ the Fisher kernel is identical to the kernel arising from the inner product of the sufficient statistics $\phi(x)$. This is not a coincidence. In fact, in our analysis of nonparametric exponential families we will encounter this fact several times (cf. Section 5 for further details). Moreover, note that the centering is immaterial, as can be seen in Lemma 13.

The above overview of kernel design is by no means complete. The reader is referred to books of Cristianini and Shawe-Taylor [2000], Schölkopf and Smola [2002], Joachims [2002], Herbrich [2002], Schölkopf et al. [2004], Shawe-Taylor and Cristianini [2004], Bakir et al. [2007] for further examples and details.

2.3 Kernel Function Classes

2.3.1 The Representer Theorem

From kernels, we now move to functions that can be expressed in terms of kernel expansions. The representer theorem [Kimeldorf and Wahba, 1971, Cox and O'Sullivan, 1990] shows that solutions of a large class of optimization problems can be expressed as kernel expansions over the sample points. We present a slightly more general version of the theorem with a simple proof [Schölkopf et al., 2001]. As above, \mathcal{H} is the RKHS associated to the kernel k .

Theorem 11 (Representer Theorem) *Denote by $\Omega : [0, \infty) \rightarrow \mathbb{R}$ a strictly monotonic increasing function, by \mathcal{X} a set, and by $c : (\mathcal{X} \times \mathbb{R}^2)^n \rightarrow \mathbb{R} \cup \{\infty\}$ an arbitrary loss function. Then each minimizer $f \in \mathcal{H}$ of the regularized risk functional*

$$c((x_1, y_1, f(x_1)), \dots, (x_n, y_n, f(x_n))) + \Omega(\|f\|_{\mathcal{H}}^2) \quad (45)$$

admits a representation of the form

$$f(x) = \sum_{i=1}^n \alpha_i k(x_i, x). \quad (46)$$

Proof We decompose any $f \in \mathcal{H}$ into a part contained in the span of the kernel functions $k(x_1, \cdot), \dots, k(x_n, \cdot)$, and one in the orthogonal complement;

$$f(x) = f_{\parallel}(x) + f_{\perp}(x) = \sum_{i=1}^n \alpha_i k(x_i, x) + f_{\perp}(x). \quad (47)$$

Here $\alpha_i \in \mathbb{R}$ and $f_{\perp} \in \mathcal{H}$ with $\langle f_{\perp}, k(x_i, \cdot) \rangle_{\mathcal{H}} = 0$ for all $i \in [n] := \{1, \dots, n\}$. By (16) we may write $f(x_j)$ (for all $j \in [n]$) as

$$f(x_j) = \langle f(\cdot), k(x_j, \cdot) \rangle = \sum_{i=1}^n \alpha_i k(x_i, x_j) + \langle f_{\perp}(\cdot), k(x_j, \cdot) \rangle_{\mathcal{H}} = \sum_{i=1}^n \alpha_i k(x_i, x_j). \quad (48)$$

Second, for all f_\perp ,

$$\Omega(\|f\|_{\mathcal{H}}^2) = \Omega\left(\left\|\sum_i^n \alpha_i k(x_i, \cdot)\right\|_{\mathcal{H}}^2 + \|f_\perp\|_{\mathcal{H}}^2\right) \geq \Omega\left(\left\|\sum_i^n \alpha_i k(x_i, \cdot)\right\|_{\mathcal{H}}^2\right). \quad (49)$$

Thus for any fixed $\alpha_i \in \mathbb{R}$ the risk functional (45) is minimized for $f_\perp = 0$. Since this also has to hold for the solution, the theorem holds. \blacksquare

Monotonicity of Ω does not prevent the regularized risk functional (45) from having multiple local minima. To ensure a global minimum, we would need to require convexity. If we discard the strictness of the monotonicity, then it no longer follows that each minimizer of the regularized risk admits an expansion (46); it still follows, however, that there is always another solution that is as good, and that *does* admit the expansion.

The significance of the Representer Theorem is that although we might be trying to solve an optimization problem in an infinite-dimensional space \mathcal{H} , containing linear combinations of kernels centered on *arbitrary* points of \mathcal{X} , it states that the solution lies in the span of n particular kernels — those centered on the training points. We will encounter (46) again further below, where it is called the *Support Vector expansion*. For suitable choices of loss functions, many of the α_i often equal 0.

2.3.2 Reduced Set Methods

Despite the finiteness of the representation in (46) it can often be the case that the number of terms in the expansion is too large in practice. This can be problematic in practical applications, since the time required to evaluate (46) is proportional to the number of terms. To deal with this issue, we thus need to express or at least approximate (46) using a smaller number of terms. Exact expressions in a reduced number of basis points are often not possible; e.g., if the kernel is strictly positive definite and the data points x_1, \dots, x_n are distinct, then no exact reduction is possible. For the purpose of approximation, we need to decide on a discrepancy measure that we want to minimize. The most natural such measure is the RKHS norm. From a practical point of view, this norm has the advantage that it can be computed using evaluations of the kernel function which opens the possibility for relatively efficient approximation algorithms. Clearly, if we were given a set of points z_1, \dots, z_m such that the function (cf. (46))

$$f = \sum_{i=1}^n \alpha_i k(x_i, \cdot). \quad (50)$$

can be expressed with small RKHS-error in the subspace spanned by $k(z_1, \cdot), \dots, k(z_m, \cdot)$, then we can compute an m -term approximation of f by projection onto that subspace.

The problem thus reduces to the choice of the expansion points z_1, \dots, z_m . In the literature, two classes of methods have been pursued. In the first, it is attempted to choose the points as a subset of some larger set x_1, \dots, x_n , which is usually the training set [Schölkopf, 1997, Frieß and Harrison, 1998]. In the second, we compute the difference in the RKHS between (50) and the *reduced set expansion* [Burges, 1996]

$$g = \sum_{p=1}^m \beta_p k(z_p, \cdot), \quad (51)$$

leading to

$$\|f - g\|^2 = \sum_{i,j=1}^n \alpha_i \alpha_j k(x_i, x_j) + \sum_{p,q=1}^m \beta_p \beta_q k(z_p, z_q) - 2 \sum_{i=1}^n \sum_{p=1}^m \alpha_i \beta_p k(x_i, z_p). \quad (52)$$

To obtain the z_1, \dots, z_m and β_1, \dots, β_m , we can minimize this objective function, which for most kernels will be a nonconvex optimization problem. Specialized methods for specific kernels exist, as well as iterative methods taking into account the fact that given the z_1, \dots, z_m , the coefficients β_1, \dots, β_m can be computed in closed form. See Schölkopf and Smola [2002] for further details.

Once we have constructed or selected a set of expansion points, we are working in a finite dimensional subspace of \mathcal{H} spanned by these points.

It is known that each closed subspace of an RKHS is itself an RKHS. Denote its kernel by l , then the projection P onto the subspace is given by (e.g., Meschkowski [1962])

$$(Pf)(x') = \langle f, l(\cdot, x') \rangle. \quad (53)$$

If the subspace is spanned by a finite set of linearly independent kernels $k(\cdot, z_1), \dots, k(\cdot, z_m)$, the kernel of the subspace takes the form

$$k_m(x, x') = (k(x, z_1), \dots, k(x, z_m)) K^{-1} (k(x', z_1), \dots, k(x', z_m))^\top, \quad (54)$$

where K is the Gram matrix of k with respect to z_1, \dots, z_m . This can be seen by noting that (i) the $k_m(\cdot, x)$ span an m -dimensional subspace, and (ii) we have $k_m(z_i, z_j) = k(z_i, z_j)$ for $i, j = 1, \dots, m$. The kernel of the subspace's orthogonal complement, on the other hand, is $k - k_m$.

The projection onto the m -dimensional subspace can be written

$$P = \sum_{i,j=1}^m (K^{-1})_{ij} k(\cdot, z_i) k(\cdot, z_j)^\top, \quad (55)$$

where $k(\cdot, z_j)^\top$ denotes the linear form mapping $k(\cdot, x)$ to $k(z_j, x)$ for $x \in \mathcal{X}$.

If we prefer to work with coordinates and the Euclidean dot product, we can use the *kernel PCA map* [Schölkopf and Smola, 2002]

$$\Phi_m : \mathcal{X} \rightarrow \mathbb{R}^m, \quad x \mapsto K^{-1/2} (k(x, z_1), \dots, k(x, z_m))^\top, \quad (56)$$

which directly projects into an orthonormal basis of the subspace.⁷

A number of articles exist discussing how to best choose the expansion points z_i for the purpose of processing a given dataset, e.g., Smola and Schölkopf [2000], Williams and Seeger [2000].

2.3.3 Regularization Properties

The regularizer $\|f\|_{\mathcal{H}_C}^2$ used in Theorem 11, which is what distinguishes SVMs from many other regularized function estimators (e.g., based on coefficient based L_1 regularizers, such as the Lasso [Tibshirani, 1996] or linear programming machines [Schölkopf and Smola, 2002]), stems from the dot product $\langle f, f \rangle_k$ in the RKHS \mathcal{H} associated with a positive definite kernel. The nature and implications of this regularizer, however, are not obvious and we shall now provide an analysis in the Fourier domain. It turns out that if the kernel is translation invariant, then its Fourier transform allows us to characterize how the different frequency components of f contribute to the value of $\|f\|_{\mathcal{H}_C}^2$. Our exposition will be informal (see also Poggio and Girosi [1990], Girosi et al. [1993], Smola et al. [1998]), and we will implicitly assume that all integrals are over \mathbb{R}^d and exist, and that the operators are well defined.

We will rewrite the RKHS dot product as

$$\langle f, g \rangle_k = \langle \Upsilon f, \Upsilon g \rangle = \langle \Upsilon^2 f, g \rangle, \quad (57)$$

⁷In case K does not have full rank, projections onto the span $\{k(\cdot, z_1), \dots, k(\cdot, z_m)\}$ are achieved by using the pseudoinverse of K instead.

where Υ is a positive (and thus symmetric) operator mapping \mathcal{H} into a function space endowed with the usual dot product

$$\langle f, g \rangle = \int f(x) \overline{g(x)} dx. \quad (58)$$

Rather than (57), we consider the equivalent condition (cf. Section 2.2.1)

$$\langle k(x, \cdot), k(x', \cdot) \rangle_k = \langle \Upsilon k(x, \cdot), \Upsilon k(x', \cdot) \rangle = \langle \Upsilon^2 k(x, \cdot), k(x', \cdot) \rangle. \quad (59)$$

If $k(x, \cdot)$ is a *Green function* of Υ^2 , we have

$$\langle \Upsilon^2 k(x, \cdot), k(x', \cdot) \rangle = \langle \delta_x, k(x', \cdot) \rangle = k(x, x'), \quad (60)$$

which by the reproducing property (16) amounts to the desired equality (59).⁸

We now consider the particular case where the kernel can be written $k(x, x') = h(x - x')$ with a continuous strictly positive definite function $h \in L_1(\mathbb{R}^d)$ (cf. Section 2.2.3). A variation of Bochner's theorem, stated by Wendland [2005], then tells us that the measure corresponding to h has a nonvanishing density v with respect to the Lebesgue measure, i.e., that k can be written as

$$k(x, x') = \int e^{-i\langle x-x', \omega \rangle} v(\omega) d\omega = \int e^{-i\langle x, \omega \rangle} \overline{e^{-i\langle x', \omega \rangle}} v(\omega) d\omega. \quad (61)$$

We would like to rewrite this as $\langle \Upsilon k(x, \cdot), \Upsilon k(x', \cdot) \rangle$ for some linear operator Υ . It turns out that a multiplication operator in the Fourier domain will do the job. To this end, recall the d -dimensional Fourier transform, given by

$$F[f](\omega) := (2\pi)^{-\frac{d}{2}} \int f(x) e^{-i\langle x, \omega \rangle} dx, \quad (62)$$

$$\text{with the inverse } F^{-1}[f](x) = (2\pi)^{-\frac{d}{2}} \int f(\omega) e^{i\langle x, \omega \rangle} d\omega. \quad (63)$$

Next, compute the Fourier transform of k as

$$F[k(x, \cdot)](\omega) = (2\pi)^{-\frac{d}{2}} \int \int \left(v(\omega') e^{-i\langle x, \omega' \rangle} \right) e^{i\langle x', \omega' \rangle} d\omega' e^{-i\langle x', \omega \rangle} dx' \quad (64)$$

$$= (2\pi)^{\frac{d}{2}} v(\omega) e^{-i\langle x, \omega \rangle}. \quad (65)$$

Hence we can rewrite (61) as

$$k(x, x') = (2\pi)^{-d} \int \frac{F[k(x, \cdot)](\omega) \overline{F[k(x', \cdot)](\omega)}}{v(\omega)} d\omega. \quad (66)$$

If our regularization operator maps

$$\Upsilon: f \mapsto (2\pi)^{-\frac{d}{2}} v^{-\frac{1}{2}} F[f], \quad (67)$$

we thus have

$$k(x, x') = \int (\Upsilon k(x, \cdot))(\omega) \overline{(\Upsilon k(x', \cdot))(\omega)} d\omega, \quad (68)$$

⁸For conditionally positive definite kernels, a similar correspondence can be established, with a regularization operator whose null space is spanned by a set of functions which are not regularized (in the case (18), which is sometimes called *conditionally positive definite of order 1*, these are the constants).

i.e., our desired identity (59) holds true.

As required in (57), we can thus interpret the dot product $\langle f, g \rangle_k$ in the RKHS as a dot product $\int (\Upsilon f)(\omega) \overline{(\Upsilon g)(\omega)} d\omega$. This allows us to understand regularization properties of k in terms of its (scaled) Fourier transform $v(\omega)$. *Small* values of $v(\omega)$ amplify the corresponding frequencies in (67). Penalizing $\langle f, f \rangle_k$ thus amounts to a *strong* attenuation of the corresponding frequencies. Hence small values of $v(\omega)$ for large $\|\omega\|$ are desirable, since high frequency components of $F[f]$ correspond to rapid changes in f . It follows that $v(\omega)$ describes the filter properties of the corresponding regularization operator Υ . In view of our comments following Theorem 8, we can translate this insight into probabilistic terms: if the probability measure $\frac{v(\omega)d\omega}{\int v(\omega)d\omega}$ describes the desired filter properties, then the natural translation invariant kernel to use is the characteristic function of the measure.

2.3.4 Remarks and Notes

The notion of kernels as dot products in Hilbert spaces was brought to the field of machine learning by Aizerman et al. [1964], Boser et al. [1992], Vapnik [1998], Schölkopf et al. [1998]. Aizerman et al. [1964] used kernels as a tool in a convergence proof, allowing them to apply the Perceptron convergence theorem to their class of potential function algorithms. To the best of our knowledge, Boser et al. [1992] were the first to use kernels to construct a nonlinear estimation algorithm, the hard margin predecessor of the Support Vector Machine, from its linear counterpart, the *generalized portrait* [Vapnik and Lerner, 1963, Vapnik, 1982]. Whilst all these uses were limited to kernels defined on vectorial data, Schölkopf [1997] observed that this restriction is unnecessary, and nontrivial kernels on other data types were proposed by Haussler [1999], Watkins [2000]. Schölkopf et al. [1998] applied the kernel trick to generalize principal component analysis and pointed out the (in retrospect obvious) fact that any algorithm which only uses the data via dot products can be generalized using kernels.

In addition to the above uses of positive definite kernels in machine learning, there has been a parallel, and partly earlier development in the field of statistics, where such kernels have been used for instance for time series analysis [Parzen, 1970] as well as regression estimation and the solution of inverse problems [Wahba, 1990].

In probability theory, positive definite kernels have also been studied in depth since they arise as covariance kernels of stochastic processes, see e.g. [Loève, 1978]. This connection is heavily being used in a subset of the machine learning community interested in prediction with Gaussian Processes [Rasmussen and Williams, 2006].

In functional analysis, the problem of Hilbert space representations of kernels has been studied in great detail; a good reference is Berg et al. [1984]; indeed, a large part of the material in the present section is based on that work. Interestingly, it seems that for a fairly long time, there have been two separate strands of development [Stewart, 1976]. One of them was the study of positive definite functions, which started later but seems to have been unaware of the fact that it considered a special case of positive definite kernels. The latter was initiated by Hilbert [1904], Mercer [1909], and pursued for instance by Schoenberg [1938]. Hilbert calls a kernel k *definit* if

$$\int_a^b \int_a^b k(x, x') f(x) f(x') dx dx' > 0 \quad (69)$$

for all nonzero continuous functions f , and shows that all eigenvalues of the corresponding integral operator $f \mapsto \int_a^b k(x, \cdot) f(x) dx$ are then positive. If k satisfies the condition (69) subject to the constraint that $\int_a^b f(x) g(x) dx = 0$, for some fixed function g , Hilbert calls it *relativ definit*. For that case, he shows that k has at most one negative eigenvalue. Note that if f is chosen to be constant, then this notion is closely related to the one of conditionally positive definite kernels, cf. (18). For further historical details see the review of Stewart [1976] or Berg et al. [1984].

3 Convex Programming Methods for Estimation

As we saw, kernels can be used both for the purpose of describing nonlinear functions subject to smoothness constraints and for the purpose of computing inner products in some feature space efficiently. In this section we focus on the latter and how it allows us to design methods of estimation based on the geometry of the problems at hand.

Unless stated otherwise $\mathbf{E}[\cdot]$ denotes the expectation with respect to all random variables of the argument. Subscripts, such as $\mathbf{E}_X[\cdot]$, indicate that the expectation is taken over X . We will omit them wherever obvious. Finally we will refer to $\mathbf{E}_{\text{emp}}[\cdot]$ as the empirical average with respect to an n -sample. Given a sample $\mathcal{S} := \{(x_1, y_1), \dots, (x_n, y_n)\} \subseteq \mathcal{X} \times \mathcal{Y}$ we now aim at finding an affine function $f(x) = \langle w, \phi(x) \rangle + b$ or in some cases a function $f(x, y) = \langle \phi(x, y), w \rangle$ such that the empirical risk on \mathcal{S} is minimized. In the binary classification case this means that we want to maximize the agreement between $\text{sgn } f(x)$ and y .

- Minimization of the empirical risk with respect to (w, b) is NP-hard [Minsky and Papert, 1969]. In fact, Ben-David et al. [2003] show that even approximately minimizing the empirical risk is NP-hard, not only for linear function classes but also for spheres and other simple geometrical objects. This means that even if the statistical challenges could be solved, we still would be confronted with a formidable algorithmic problem.
- The indicator function $\{yf(x) < 0\}$ is discontinuous and even small changes in f may lead to large changes in both empirical and expected risk. Properties of such functions can be captured by the VC-dimension [Vapnik and Chervonenkis, 1971], that is, the maximum number of observations which can be labeled in an arbitrary fashion by functions of the class. Necessary and sufficient conditions for estimation can be stated in these terms [Vapnik and Chervonenkis, 1991]. However, much tighter bounds can be obtained by also using the scale of the class [Alon et al., 1993, Bartlett et al., 1996, Williamson et al., 2001]. In fact, there exist function classes parameterized by a single scalar which have infinite VC-dimension [Vapnik, 1995].

Given the difficulty arising from minimizing the empirical risk we now discuss algorithms which minimize an upper bound on the empirical risk, while providing good computational properties and consistency of the estimators. The statistical analysis is relegated to Section 4.

3.1 Support Vector Classification

Assume that \mathcal{S} is linearly separable, i.e. there exists a linear function $f(x)$ such that $\text{sgn } yf(x) = 1$ on \mathcal{S} . In this case, the task of finding a large margin separating hyperplane can be viewed as one of solving [Vapnik and Lerner, 1963]

$$\underset{w, b}{\text{minimize}} \quad \frac{1}{2} \|w\|^2 \quad \text{s.t.} \quad y_i (\langle w, x_i \rangle + b) \geq 1. \quad (70)$$

Note that $\|w\|^{-1} f(x_i)$ is the distance of the point x_i to the hyperplane $H(w, b) := \{x \mid \langle w, x \rangle + b = 0\}$. The condition $y_i f(x_i) \geq 1$ implies that the margin of separation is at least $2 \|w\|^{-1}$. The bound becomes exact if equality is attained for some $y_i = 1$ and $y_j = -1$. Consequently minimizing $\|w\|$ subject to the constraints maximizes the margin of separation. Eq. (70) is a quadratic program which can be solved efficiently [Luenberger, 1984, Fletcher, 1989, Boyd and Vandenberghe, 2004].

Mangasarian [1965] devised a similar optimization scheme using $\|w\|_1$ instead of $\|w\|_2$ in the objective function of (70). The result is a *linear* program. In general, one can show [Smola et al., 2000] that minimizing the ℓ_p norm of w leads to the maximizing of the margin of separation in the ℓ_q norm where $\frac{1}{p} + \frac{1}{q} = 1$. The ℓ_1 norm leads to sparse approximation schemes (see also Chen et al. [1999]), whereas the ℓ_2 norm can be extended to Hilbert spaces and kernels.

To deal with nonseparable problems, i.e. cases when (70) is infeasible, we need to relax the constraints of the optimization problem. Bennett and Mangasarian [1992] and Cortes and Vapnik [1995] impose a linear penalty on the violation of the large-margin constraints to obtain:

$$\underset{w,b,\xi}{\text{minimize}} \quad \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n \xi_i \quad \text{s.t.} \quad y_i (\langle w, x_i \rangle + b) \geq 1 - \xi_i \text{ and } \xi_i \geq 0, \forall i \in [n]. \quad (71)$$

Eq. (71) is a quadratic program which is always feasible (e.g. $w, b = 0$ and $\xi_i = 1$ satisfy the constraints). $C > 0$ is a regularization constant trading off the violation of the constraints vs. maximizing the overall margin.

Whenever the dimensionality of \mathcal{X} exceeds n , direct optimization of (71) is computationally inefficient. This is particularly true if we map from \mathcal{X} into an RKHS. To address these problems one may solve the problem in dual space as follows. The Lagrange function of (71) is given by

$$\mathcal{L}(w, b, \xi, \alpha, \eta) = \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n \xi_i + \sum_{i=1}^n \alpha_i (1 - \xi_i - y_i (\langle w, x_i \rangle + b)) - \sum_{i=1}^n \eta_i \xi_i \quad (72)$$

where $\alpha_i, \eta_i \geq 0$ for all $i \in [n]$. To compute the dual of \mathcal{L} we need to identify the first order conditions in w, b . They are given by

$$\partial_w \mathcal{L} = w - \sum_{i=1}^n \alpha_i y_i x_i = 0 \text{ and } \partial_b \mathcal{L} = - \sum_{i=1}^n \alpha_i y_i = 0 \text{ and } \partial_{\xi_i} \mathcal{L} = C - \alpha_i + \eta_i = 0. \quad (73)$$

This translates into $w = \sum_{i=1}^n \alpha_i y_i x_i$, the linear constraint $\sum_{i=1}^n \alpha_i y_i = 0$, and the box-constraint $\alpha_i \in [0, C]$ arising from $\eta_i \geq 0$. Substituting (73) into \mathcal{L} yields the Wolfe [1961] dual

$$\underset{\alpha}{\text{minimize}} \quad \frac{1}{2} \alpha^\top Q \alpha - \alpha^\top \mathbf{1} \quad \text{s.t.} \quad \alpha^\top y = 0 \text{ and } \alpha_i \in [0, C], \forall i \in [n]. \quad (74)$$

$Q \in \mathbb{R}^{n \times n}$ is the matrix of inner products $Q_{ij} := y_i y_j \langle x_i, x_j \rangle$. Clearly this can be extended to feature maps and kernels easily via $K_{ij} := y_i y_j \langle \Phi(x_i), \Phi(x_j) \rangle = y_i y_j k(x_i, x_j)$. Note that w lies in the span of the x_i . This is an instance of the Representer Theorem (Theorem 11). The KKT conditions [Karush, 1939, Kuhn and Tucker, 1951, Boser et al., 1992, Cortes and Vapnik, 1995] require that at optimality $\alpha_i (y_i f(x_i) - 1) = 0$. This means that only those x_i may appear in the expansion (73) for which $y_i f(x_i) \leq 1$, as otherwise $\alpha_i = 0$. The x_i with $\alpha_i > 0$ are commonly referred to as support vectors.

Note that $\sum_{i=1}^n \xi_i$ is an upper bound on the empirical risk, as $y_i f(x_i) \leq 0$ implies $\xi_i \geq 1$ (see also Lemma 12). The number of misclassified points x_i itself depends on the configuration of the data and the value of C . The result of Ben-David et al. [2003] suggests that finding even an approximate minimum classification error solution is difficult. That said, it is possible to modify (71) such that a desired target number of observations violates $y_i f(x_i) \geq \rho$ for some $\rho \in \mathbb{R}$ by making the threshold itself a variable of the optimization problem [Schölkopf et al., 2000]. This leads to the following optimization problem (ν -SV classification):

$$\underset{w,b,\xi}{\text{minimize}} \quad \frac{1}{2} \|w\|^2 + \sum_{i=1}^n \xi_i - n\nu\rho \quad \text{subject to} \quad y_i (\langle w, x_i \rangle + b) \geq \rho - \xi_i \text{ and } \xi_i \geq 0. \quad (75)$$

The dual of (75) is essentially identical to (74) with the exception of an additional constraint:

$$\underset{\alpha}{\text{minimize}} \quad \frac{1}{2} \alpha^\top Q \alpha \quad \text{subject to} \quad \alpha^\top y = 0 \text{ and } \alpha^\top \mathbf{1} = n\nu \text{ and } \alpha_i \in [0, 1]. \quad (76)$$

One can show that for every C there exists a ν such that the solution of (76) is a multiple of the solution of (74). Schölkopf et al. [2000] prove that solving (76) for which $\rho > 0$ satisfies:

1. ν is an upper bound on the fraction of margin errors.
2. ν is a lower bound on the fraction of SVs.

Moreover under mild conditions, with probability 1, asymptotically, ν equals both the fraction of SVs and the fraction of errors.

This statement implies that whenever the data are sufficiently well separable (that is $\rho > 0$), ν -SV Classification finds a solution with a fraction of at most ν margin errors. Also note that for $\nu = 1$, all $\alpha_i = 1$, that is, f becomes an affine copy of the Parzen windows classifier (5).

3.2 Estimating the Support of a Density

We now extend the notion of linear separation to that of estimating the support of a density [Schölkopf et al., 2001, Tax and Duin, 1999]. Denote by $X = \{x_1, \dots, x_n\} \subseteq \mathcal{X}$ the sample drawn from $P(x)$. Let \mathcal{C} be a class of measurable subsets of \mathcal{X} and let λ be a real-valued function defined on \mathcal{C} . The *quantile function* [Einmal and Mason, 1992] with respect to $(P, \lambda, \mathcal{C})$ is defined as

$$U(\mu) = \inf \{ \lambda(C) | P(C) \geq \mu, C \in \mathcal{C} \} \text{ where } \mu \in (0, 1]. \quad (77)$$

We denote by $C_\lambda(\mu)$ and $C_\lambda^m(\mu)$ the (not necessarily unique) $C \in \mathcal{C}$ that attain the infimum (when it is achievable) on $P(x)$ and on the empirical measure given by X respectively. A common choice of λ is Lebesgue measure, in which case $C_\lambda(\mu)$ is the minimum volume set $C \in \mathcal{C}$ that contains at least a fraction μ of the probability mass.

Support estimation requires us to find some $C_\lambda^m(\mu)$ such that $|P(C_\lambda^m(\mu)) - \mu|$ is small. This is where the complexity trade-off enters: On the one hand, we want to use a rich class \mathcal{C} to capture all possible distributions, on the other hand large classes lead to large deviations between μ and $P(C_\lambda^m(\mu))$. Therefore, we have to consider classes of sets which are suitably restricted. This can be achieved using an SVM regularizer.

In the case where $\mu < 1$, it seems the first work was reported in Sager [1979] and Hartigan [1987], in which $\mathcal{X} = \mathbb{R}^2$, with \mathcal{C} being the class of closed convex sets in \mathcal{X} . Nolan [1991] considered higher dimensions, with \mathcal{C} being the class of ellipsoids. Tsybakov [1997] studied an estimator based on piecewise polynomial approximation of $C_\lambda(\mu)$ and showed it attains the asymptotically minimax rate for certain classes of densities. Polonik [1997] studied the estimation of $C_\lambda(\mu)$ by $C_\lambda^m(\mu)$. He derived asymptotic rates of convergence in terms of various measures of richness of \mathcal{C} . More information on minimum volume estimators can be found in that work, and in Schölkopf et al. [2001].

SV support estimation works by using SV support estimation relates to previous work as follows: set $\lambda(C_w) = \|w\|^2$, where $C_w = \{x | f_w(x) \geq \rho\}$, $f_w(x) = \langle w, x \rangle$, and (w, ρ) are respectively a weight vector and an offset. Stated as a convex optimization problem we want to separate the data from the origin with maximum margin via:

$$\underset{w, \xi, \rho}{\text{minimize}} \quad \frac{1}{2} \|w\|^2 + \sum_{i=1}^n \xi_i - n\nu\rho \text{ subject to } \langle w, x_i \rangle \geq \rho - \xi_i \text{ and } \xi_i \geq 0. \quad (78)$$

Here, $\nu \in (0, 1]$ plays the same role as in (75), controlling the number of observations x_i for which $f(x_i) \leq \rho$. Since nonzero slack variables ξ_i are penalized in the objective function, if w and ρ solve this problem, then the decision function $f(x)$ will attain or exceed ρ for at least a fraction $1 - \nu$ of the x_i contained in X while the regularization term $\|w\|$ will still be small. The dual of (78) yields:

$$\underset{\alpha}{\text{minimize}} \quad \frac{1}{2} \alpha^\top K \alpha \text{ subject to } \alpha^\top \mathbf{1} = \nu n \text{ and } \alpha_i \in [0, 1] \quad (79)$$

To compare (79) to a Parzen windows estimator assume that k is such that it can be normalized as a density in input space, such as a Gaussian. Using $\nu = 1$ in (79) the constraints automatically imply

$\alpha_i = 1$. Thus f reduces to a Parzen windows estimate of the underlying density. For $\nu < 1$, the equality constraint (79) still ensures that f is a thresholded density, now depending only on a *subset* of X — those which are important for deciding whether $f(x) \leq \rho$.

3.3 Regression Estimation

SV regression was first proposed in Vapnik [1995], Vapnik et al. [1997], and Drucker et al. [1997] using the so-called ϵ -insensitive loss function. It is a direct extension of the soft-margin idea to regression: instead of requiring that $yf(x)$ exceeds some margin value, we now require that the values $y - f(x)$ are bounded by a margin on both sides. That is, we impose the soft constraints

$$y_i - f(x_i) \leq \epsilon_i - \xi_i \text{ and } f(x_i) - y_i \leq \epsilon_i - \xi_i^* \quad (80)$$

where $\xi_i, \xi_i^* \geq 0$. If $|y_i - f(x_i)| \leq \epsilon$ no penalty occurs. The objective function is given by the sum of the slack variables ξ_i, ξ_i^* penalized by some $C > 0$ and a measure for the slope of the function $f(x) = \langle w, x \rangle + b$, that is $\frac{1}{2} \|w\|^2$.

Before computing the dual of this problem let us consider a somewhat more general situation where we use a range of different convex penalties for the deviation between y_i and $f(x_i)$. One may check that minimizing $\frac{1}{2} \|w\|^2 + C \sum_{i=1}^n \xi_i + \xi_i^*$ subject to (80) is equivalent to solving

$$\text{minimize}_{w,b,\xi} \frac{1}{2} \|w\|^2 + \sum_{i=1}^n \psi(y_i - f(x_i)) \text{ where } \psi(\xi) = \max(0, |\xi| - \epsilon). \quad (81)$$

Choosing different loss functions ψ leads to a rather rich class of estimators:

- $\psi(\xi) = \frac{1}{2} \xi^2$ yields penalized least squares (LS) regression [Hoerl and Kennard, 1970, Tikhonov, 1963, Morozov, 1984, Wahba, 1990]. The corresponding optimization problem can be minimized by solving a linear system.
- For $\psi(\xi) = |\xi|$ we obtain the penalized least absolute deviations (LAD) estimator [Bloomfield and Steiger, 1983]. That is, we obtain a quadratic program to estimate the conditional median.
- A combination of LS and LAD loss yields a penalized version of Huber's robust regression [Huber, 1981, Smola and Schölkopf, 1998]. In this case we have $\psi(\xi) = \frac{1}{2\sigma} \xi^2$ for $|\xi| \leq \sigma$ and $\psi(\xi) = |\xi| - \frac{\sigma}{2}$ for $|\xi| \geq \sigma$.
- Note that also quantile regression [Koenker, 2005] can be modified to work with kernels [Schölkopf et al., 2000, Takeuchi et al., 2006] by using as loss function the “pinball” loss, that is $\psi(\xi) = (1-\tau)\psi$ if $\psi < 0$ and $\psi(\xi) = \tau\psi$ if $\psi > 0$.

All the optimization problems arising from the above five cases are convex quadratic programs. Their dual resembles that of (80), namely

$$\text{minimize}_{\alpha, \alpha^*} \frac{1}{2} (\alpha - \alpha^*)^\top K (\alpha - \alpha^*) + \epsilon^\top (\alpha + \alpha^*) - y^\top (\alpha - \alpha^*) \quad (82a)$$

$$\text{subject to } (\alpha - \alpha^*)^\top \mathbf{1} = 0 \text{ and } \alpha_i, \alpha_i^* \in [0, C]. \quad (82b)$$

Here $K_{ij} = \langle x_i, x_j \rangle$ for linear models and $K_{ij} = k(x_i, x_j)$ if we map $x \rightarrow \Phi(x)$. The ν -trick, as described in (75) [Schölkopf et al., 2000], can be extended to regression, allowing one to choose the margin of approximation automatically. In this case (82a) drops the terms in ϵ . In its place, we add a linear constraint $(\alpha - \alpha^*)^\top \mathbf{1} = \nu n$. Likewise, LAD is obtained from (82) by dropping the terms in ϵ without additional constraints. Robust regression leaves (82) unchanged, however, in the definition of K we have an additional term of σ^{-1} on the main diagonal. Further details can be found in Schölkopf and Smola [2002]. For Quantile Regression we drop ϵ and we obtain different constants $C(1-\tau)$ and $C\tau$ for the constraints on α^* and α . We will discuss uniform convergence properties of the empirical risk estimates with respect to various $\psi(\xi)$ in Section 4.

3.4 Multicategory Classification, Ranking and Ordinal Regression

Many estimation problems cannot be described by assuming that $\mathcal{Y} = \{\pm 1\}$. In this case it is advantageous to go beyond simple functions $f(x)$ depending on x only. Instead, we can encode a larger degree of information by estimating a function $f(x, y)$ and subsequently obtaining a prediction via $\hat{y}(x) := \operatorname{argmax}_{y \in \mathcal{Y}} f(x, y)$. In other words, we study problems where y is obtained as the solution of an optimization problem over $f(x, y)$ and we wish to find f such that y matches y_i as well as possible for relevant inputs x .

Note that the loss may be more than just a simple 0 – 1 loss. In the following we denote by $\Delta(y, y')$ the loss incurred by estimating y' instead of y . Without loss of generality we require that $\Delta(y, y) = 0$ and that $\Delta(y, y') \geq 0$ for all $y, y' \in \mathcal{Y}$. Key in our reasoning is the following:

Lemma 12 *Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ and assume that $\Delta(y, y') \geq 0$ with $\Delta(y, y) = 0$. Moreover let $\xi \geq 0$ such that $f(x, y) - f(x, y') \geq \Delta(y, y') - \xi$ for all $y' \in \mathcal{Y}$. In this case $\xi \geq \Delta(y, \operatorname{argmax}_{y' \in \mathcal{Y}} f(x, y'))$.*

Proof Denote by $y^* := \operatorname{argmax}_{y \in \mathcal{Y}} f(x, y)$. By assumption we have $\xi \geq \Delta(y, y^*) + f(x, y^*) - f(x, y)$. Since $f(x, y^*) \geq f(x, y')$ for all $y' \in \mathcal{Y}$ the inequality holds. ■

The construction of the estimator was suggested in Taskar et al. [2003] and Tsochantaridis et al. [2005], and a special instance of the above lemma is given by Joachims [2005]. While the bound appears quite innocuous, it allows us to describe a much richer class of estimation problems as a convex program.

To deal with the added complexity we assume that f is given by $f(x, y) = \langle \Phi(x, y), w \rangle$. Given the possibly nontrivial connection between x and y the use of $\Phi(x, y)$ cannot be avoided. Corresponding kernel functions are given by $k(x, y, x', y') = \langle \Phi(x, y), \Phi(x', y') \rangle$. We have the following optimization problem [Tsochantaridis et al., 2005]:

$$\underset{w, \xi}{\text{minimize}} \quad \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n \xi_i \quad \text{s.t.} \quad \langle w, \Phi(x_i, y_i) - \Phi(x_i, y) \rangle \geq \Delta(y_i, y) - \xi_i, \quad \forall i \in [n], y \in \mathcal{Y}. \quad (83)$$

This is a convex optimization problem which can be solved efficiently if the constraints can be evaluated without high computational cost. One typically employs column-generation methods [Hettich and Kortanek, 1993, Rätsch, 2001, Bennett et al., 2000, Tsochantaridis et al., 2005, Fletcher, 1989] which identify one violated constraint at a time to find an approximate minimum of the optimization problem.

To describe the flexibility of the framework set out by (83) we give several examples below:

- Binary classification can be recovered by setting $\Phi(x, y) = y\Phi(x)$, in which case the constraint of (83) reduces to $2y_i \langle \Phi(x_i), w \rangle \geq 1 - \xi_i$. Ignoring constant offsets and a scaling factor of 2, this is exactly the standard SVM optimization problem.
- Multicategory classification problems [Crammer and Singer, 2001, Collins, 2000, Allwein et al., 2000, Rätsch et al., 2003] can be encoded via $\mathcal{Y} = [N]$, where N is the number of classes and $\Delta(y, y') = 1 - \delta_{y, y'}$. In other words, the loss is 1 whenever we predict the wrong class and 0 for correct classification. Corresponding kernels are typically chosen to be $\delta_{y, y'} k(x, x')$.
- We can deal with joint labeling problems by setting $\mathcal{Y} = \{\pm 1\}^n$. In other words, the error measure does not depend on a single observation but on an entire set of labels. Joachims [2005] shows that the so-called F_1 score [van Rijsbergen, 1979] used in document retrieval and the area under the ROC curve [Bamber, 1975, Gribskov and Robinson, 1996] fall into this category of problems. Moreover, Joachims [2005] derives an $O(n^2)$ method for evaluating the inequality constraint over \mathcal{Y} .
- Multilabel estimation problems deal with the situation where we want to find the best subset of labels $\mathcal{Y} \subseteq 2^{[N]}$ which correspond to some observation x . The problem is described in Elisseeff and Weston [2001], where the authors devise a ranking scheme such that $f(x, i) > f(x, j)$ if label $i \in \mathcal{Y}$ and $j \notin \mathcal{Y}$. It is a special case of a general ranking approach described next.

Note that (83) is invariant under translations $\Phi(x, y) \leftarrow \Phi(x, y) + \Phi_0$ where Φ_0 is constant, as $\Phi(x_i, y_i) - \Phi(x_i, y)$ remains unchanged. In practice this means that transformations $k(x, y, x', y') \leftarrow k(x, y, x', y') + \langle \Phi_0, \Phi(x, y) \rangle + \langle \Phi_0, \Phi(x', y') \rangle + \|\Phi_0\|^2$ do not affect the outcome of the estimation process. Since Φ_0 was arbitrary, we have the following lemma:

Lemma 13 *Let \mathcal{H} be an RKHS on $\mathcal{X} \times \mathcal{Y}$ with kernel k . Moreover let $g \in \mathcal{H}$. Then the function $k(x, y, x', y') + f(x, y) + f(x', y') + \|g\|_{\mathcal{H}}^2$ is a kernel and it yields the same estimates as k .*

We need a slight extension to deal with general ranking problems. Denote by $\mathcal{Y} = \text{Graph}[N]$ the set of all directed graphs on N vertices which do not contain loops of less than three nodes. Here an edge $(i, j) \in y$ indicates that i is preferred to j with respect to the observation x . It is the goal to find some function $f : \mathcal{X} \times [N] \rightarrow \mathbb{R}$ which imposes a total order on $[N]$ (for a given x) by virtue of the function values $f(x, i)$ such that the total order and y are in good agreement.

More specifically, Dekel et al. [2003] and Crammer and Singer [2005] propose a decomposition algorithm \mathcal{A} for the graphs y such that the estimation error is given by the number of subgraphs of y which are in disagreement with the total order imposed by f . As an example, multiclass classification can be viewed as a graph y where the correct label i is at the root of a directed graph and all incorrect labels are its children. Multilabel classification is then a bipartite graph where the correct labels only contain outgoing arcs and the incorrect labels only incoming ones.

This setting leads to a form similar to (83) except for the fact that we now have constraints over each subgraph $G \in \mathcal{A}(y)$. We solve

$$\begin{aligned} & \underset{w, \xi}{\text{minimize}} \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n |\mathcal{A}(y_i)|^{-1} \sum_{G \in \mathcal{A}(y_i)} \xi_{iG} & (84) \\ & \text{subject to } \langle w, \Phi(x_i, u) - \Phi(x_i, v) \rangle \geq 1 - \xi_{iG} \text{ and } \xi_{iG} \geq 0 \text{ for all } (u, v) \in G \in \mathcal{A}(y_i). \end{aligned}$$

That is, we test for all $(u, v) \in G$ whether the ranking imposed by $G \in y_i$ is satisfied.

Finally, ordinal regression problems which perform ranking not over labels y but rather over observations x were studied by Herbrich et al. [2000] and Chapelle and Harchaoui [2005] in the context of ordinal regression and conjoint analysis respectively. In ordinal regression x is preferred to x' if $f(x) > f(x')$ and hence one minimizes an optimization problem akin to (83), with constraint $\langle w, \Phi(x_i) - \Phi(x_j) \rangle \geq 1 - \xi_{ij}$. In conjoint analysis the same operation is carried out for $\Phi(x, u)$, where u is the user under consideration. Similar models were also studied by Basilico and Hofmann [2004]. Further models will be discussed in Section 5, in particular situations where \mathcal{Y} is of exponential size. These models allow one to deal with sequences and more sophisticated structures.

3.5 Applications of SVM Algorithms

When SVMs were first presented, they initially met with scepticism in the statistical community. Part of the reason was that as described, SVMs construct their decision rules in potentially very high-dimensional feature spaces associated with kernels. Although there was a fair amount of theoretical work addressing this issue (see Section 4 below), it was probably to a larger extent the empirical success of SVMs that paved its way to become a standard method of the statistical toolbox. The first successes of SVMs on practical problem were in handwritten digit recognition, which was the main benchmark task considered in the Adaptive Systems Department at AT&T Bell Labs where SVMs were developed (see e.g. LeCun et al. [1998]). Using methods to incorporate transformation invariances, SVMs were shown to beat the world record on the MNIST benchmark set, at the time the gold standard in the field [DeCoste and Schölkopf, 2002]. There has been a significant number of further computer vision applications of SVMs since then, including tasks such as object recognition [Blanz et al., 1996, Chapelle et al., 1999, Holub et al., 2005, Everingham et al., 2005] and detection [Romdhani et al., 2004]. Nevertheless, it is probably fair to say

that two other fields have been more influential in spreading the use of SVMs: bioinformatics and natural language processing. Both of them have generated a spectrum of challenging high-dimensional problems on which SVMs excel, such as microarray processing tasks [Brown et al., 2000] and text categorization [Dumais, 1998]. For further references, see Schölkopf et al. [2004] and Joachims [2002].

Many successful applications have been implemented using SV classifiers; however, also the other variants of SVMs have led to very good results, including SV regression [Müller et al., 1997], SV novelty detection [Hayton et al., 2001], SVMs for ranking [Herbrich et al., 2000] and more recently, problems with interdependent labels [Tsochantaridis et al., 2005, McCallum et al., 2005].

At present there exists a large number of readily available software packages for SVM optimization. For instance, SVMStruct, based on [Tsochantaridis et al., 2005] solves structured estimation problems. LibSVM is an open source solver which excels on binary problems. The Torch package contains a number of estimation methods, including SVM solvers. Several SVM implementations are also available via statistical packages, such as R.

4 Margins and Uniform Convergence Bounds

So far we motivated the algorithms by means of their practicality and the fact that $0 - 1$ loss functions yield hard-to-control estimators. We now follow up on the analysis by providing uniform convergence bounds for large margin classifiers. We focus on the case of scalar valued functions applied to classification for two reasons: The derivation is well established and it can be presented in a concise fashion. Secondly, the derivation of corresponding bounds for the vectorial case is by and large still an open problem. Preliminary results exist, such as the bounds by Collins [2000] for the case of perceptrons, Taskar et al. [2003] who derive capacity bounds in terms of covering numbers by an explicit covering construction, and Bartlett and Mendelson [2002], who give Gaussian average bounds for vectorial functions. We believe that the scaling behavior of these bounds in the number of classes $|\mathcal{Y}|$ is currently not optimal, when applied to the problems of type (83).

Our analysis is based on the following ideas: firstly the $0 - 1$ loss is upper bounded by some function $\psi(yf(x))$ which can be efficiently minimized, such as the soft margin function $\max(0, 1 - yf(x))$ of the previous section. Secondly we prove that the empirical average of the ψ -loss is concentrated close to its expectation. This will be achieved by means of Rademacher averages. Thirdly we show that under rather general conditions the minimization of the ψ -loss is consistent with the minimization of the expected risk. Finally, we combine these bounds to obtain rates of convergence which only depend on the Rademacher average and the approximation properties of the function class under consideration.

4.1 Margins and Empirical Risk

While the sign of $yf(x)$ can be used to assess the accuracy of a binary classifier we saw that for algorithmic reasons one rather optimizes (a smooth function of) $yf(x)$ directly. In the following we assume that the binary loss $\chi(\xi) = \frac{1}{2}(1 - \text{sgn}\xi)$ is majorized by some function $\psi(\xi) \geq \chi(\xi)$, e.g. via the construction of Lemma 12. Consequently $\mathbf{E}[\chi(yf(x))] \leq \mathbf{E}[\psi(yf(x))]$ and likewise $\mathbf{E}_{\text{emp}}[\chi(yf(x))] \leq \mathbf{E}_{\text{emp}}[\psi(yf(x))]$. The hope is (as will be shown in Section 4.3) that minimizing the upper bound leads to consistent estimators.

There is a long-standing tradition of minimizing $yf(x)$ rather than the number of misclassifications. $yf(x)$ is known as “margin” (based on the geometrical reasoning) in the context of SVMs [Vapnik and Lerner, 1963, Mangasarian, 1965], as “stability” in the context of Neural Networks [Krauth and Mézard, 1987, Ruján, 1993], and as the “edge” in the context of arcing [Breiman, 1999]. One may show [Makovoz, 1996, Barron, 1993, Herbrich and Williamson, 2002] that functions f in an RKHS achieving a large margin can be approximated by another function f' achieving almost the same empirical error using a much smaller number of kernel functions.

Note that by default, uniform convergence bounds are expressed in terms of minimization of the empirical risk average with respect to a *fixed* function class \mathcal{F} , e.g. Vapnik and Chervonenkis [1971]. This is

very much unlike what is done in practice: in SVM (83) the sum of empirical risk and a regularizer is minimized. However, one may check that minimizing $\mathbf{E}_{\text{emp}}[\psi(yf(x))]$ subject to $\|w\|^2 \leq W$ is equivalent to minimizing $\mathbf{E}_{\text{emp}}[\psi(yf(x))] + \lambda \|w\|^2$ for suitably chosen values of λ . The equivalence is immediate by using Lagrange multipliers. For numerical reasons, however, the second formulation is much more convenient [Tikhonov, 1963, Morozov, 1984], as it acts as a regularizer. Finally, for the design of adaptive estimators, so-called luckiness results exist, which provide risk bounds in a data-dependent fashion [Shawe-Taylor et al., 1998, Herbrich and Williamson, 2002].

4.2 Uniform Convergence and Rademacher Averages

The next step is to bound the deviation $\mathbf{E}_{\text{emp}}[\psi(yf(x))] - \mathbf{E}[\psi(yf(x))]$ by means of Rademacher averages. For details see Bousquet et al. [2005], Mendelson [2003], Bartlett et al. [2002], and Koltchinskii [2001]. Denote by $g : \mathcal{X}^n \rightarrow \mathbb{R}$ a function of n variables and let $c > 0$ such that $|g(x_1, \dots, x_n) - g(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)| \leq c$ for all $x_1, \dots, x_n, x'_i \in \mathcal{X}$ and for all $i \in [n]$, then [McDiarmid, 1989]

$$\mathbf{P}\{\mathbf{E}[g(x_1, \dots, x_n)] - g(x_1, \dots, x_n) > \epsilon\} \leq \exp(-2\epsilon^2/nc^2). \quad (85)$$

Assume that $f(x) \in [0, B]$ for all $f \in \mathcal{F}$ and let $g(x_1, \dots, x_n) := \sup_{f \in \mathcal{F}} |\mathbf{E}_{\text{emp}}[f(x)] - \mathbf{E}[f(x)]|$. Then it follows that $c \leq \frac{B}{n}$. Solving (85) for g we obtain that with probability at least $1 - \delta$

$$\sup_{f \in \mathcal{F}} \mathbf{E}[f(x)] - \mathbf{E}_{\text{emp}}[f(x)] \leq \mathbf{E}\left[\sup_{f \in \mathcal{F}} \mathbf{E}[f(x)] - \mathbf{E}_{\text{emp}}[f(x)]\right] + B\sqrt{-\frac{\log \delta}{2n}}. \quad (86)$$

This means that with high probability the largest deviation between the sample average and its expectation is concentrated around its mean and within an $O(n^{-\frac{1}{2}})$ term. The expectation can be bounded by a classical symmetrization argument [Vapnik and Chervonenkis, 1971] as follows:

$$\begin{aligned} \mathbf{E}_X \left[\sup_{f \in \mathcal{F}} \mathbf{E}[f(x')] - \mathbf{E}_{\text{emp}}[f(x)] \right] &\leq \mathbf{E}_{X, X'} \left[\sup_{f \in \mathcal{F}} \mathbf{E}_{\text{emp}}[f(x')] - \mathbf{E}_{\text{emp}}[f(x)] \right] \\ = \mathbf{E}_{X, X', \sigma} \left[\sup_{f \in \mathcal{F}} \mathbf{E}_{\text{emp}}[\sigma f(x')] - \mathbf{E}_{\text{emp}}[\sigma f(x)] \right] &\leq 2\mathbf{E}_{X, \sigma} \left[\sup_{f \in \mathcal{F}} \mathbf{E}_{\text{emp}}[\sigma f(x)] \right] =: 2R_n[\mathcal{F}]. \end{aligned}$$

The first inequality follows from the convexity of the argument of the expectation, the second equality follows from the fact that x_i and x'_i are drawn i.i.d. from the same distribution, hence we may swap terms. Here σ_i are independent ± 1 -valued zero-mean Rademacher random variables. $R_n[\mathcal{F}]$ is referred as the Rademacher average [Mendelson, 2001, Bartlett and Mendelson, 2002, Koltchinskii, 2001] of \mathcal{F} wrt. sample size n .

For linear function classes $R_n[\mathcal{F}]$ takes on a particularly nice form. We begin with $\mathcal{F} := \{f|f(x) = \langle x, w \rangle \text{ and } \|w\| \leq 1\}$. It follows that $\sup_{\|w\| \leq 1} \sum_{i=1}^n \sigma_i \langle w, x_i \rangle = \|\sum_{i=1}^n \sigma_i x_i\|$. Hence

$$nR_n[\mathcal{F}] = \mathbf{E}_{X, \sigma} \left\| \sum_{i=1}^n \sigma_i x_i \right\| \leq \mathbf{E}_X \left[\mathbf{E}_\sigma \left[\left\| \sum_{i=1}^n \sigma_i x_i \right\|^2 \right] \right]^{\frac{1}{2}} = \mathbf{E}_X \left[\sum_{i=1}^n \|x_i\|^2 \right]^{\frac{1}{2}} \leq \sqrt{n\mathbf{E}[\|x\|^2]}. \quad (87)$$

Here the first inequality is a consequence of Jensen's inequality, the second equality follows from the fact that σ_i are i.i.d. zero-mean random variables, and the last step again is a result of Jensen's inequality. Corresponding tight lower bounds by a factor of $1/\sqrt{2}$ exist and they are a result of the Khintchine-Kahane inequality [Kahane, 1968].

Note that (87) allows us to bound $R_n[\mathcal{F}] \leq n^{-\frac{1}{2}}r$ where r is the average norm of vectors in the sample. An extension to kernel functions is straightforward: by design of the inner product we have $r = \sqrt{\mathbf{E}_x[k(x, x)]}$. Note that this bound is *independent* of the dimensionality of the data but rather only depends on the expected norm of the input vectors. Moreover r can also be interpreted as the trace of the integral operator with kernel $k(x, x')$ and probability measure on \mathcal{X} .

Since we are computing $\mathbf{E}_{\text{emp}}[\psi(yf(x))]$ we are interested in the Rademacher complexity of $\psi \circ \mathcal{F}$. Bartlett and Mendelson [2002] show that $R_n[\psi \circ \mathcal{F}] \leq LR_n[\mathcal{F}]$ for any Lipschitz continuous function ψ with Lipschitz constant L and with $\psi(0) = 0$. Secondly, for $\{yb \text{ where } |b| \leq B\}$ the Rademacher average can be bounded by $B\sqrt{2 \log 2/n}$, as follows from [Bousquet et al., 2005, eq. (4)]. This takes care of the offset b . For sums of function classes \mathcal{F} and \mathcal{G} we have $R_n[\mathcal{F} + \mathcal{G}] \leq R_n[\mathcal{F}] + R_n[\mathcal{G}]$. This means that for linear functions with $\|w\| \leq W$, $|b| \leq B$, and ψ Lipschitz continuous with constant L we have $R_n \leq \frac{L}{\sqrt{n}}(Wr + B\sqrt{2 \log 2})$.

4.3 Upper Bounds and Convex Functions

We briefly discuss consistency of minimization of the surrogate loss function $\psi : \mathbb{R} \rightarrow [0, \infty)$, which we assume to be convex and χ -majorizing, i.e. $\psi \geq \chi$ [Jordan et al., 2003, Zhang, 2004]. Examples of such functions are the soft-margin loss $\max(0, 1 - \gamma\xi)$, which we discussed in Section 3, the Boosting Loss $e^{-\xi}$, which is commonly used in AdaBoost [Schapire et al., 1998, Rätsch et al., 2001], and the logistic loss $\ln(1 + e^{-\xi})$ (see Section 5).

Denote by f_χ^* the minimizer of the expected risk and let f_ψ^* be the minimizer of $\mathbf{E}[\psi(yf(x))]$ with respect to f . Then under rather general conditions on ψ [Zhang, 2004] for all f the following inequality holds:

$$\mathbf{E}[\chi(yf(x))] - \mathbf{E}[\chi(yf_\chi^*(x))] \leq c(\mathbf{E}[\psi(yf(x))] - \mathbf{E}[\psi(yf_\psi^*(x))])^s. \quad (88)$$

In particular we have $c = 4$ and $s = 1$ for soft margin loss, whereas for Boosting and logistic regression $c = \sqrt{8}$ and $s = \frac{1}{2}$. Note that (88) implies that the minimizer of the ψ loss is consistent, i.e. $\mathbf{E}[\chi(yf_\psi(x))] = \mathbf{E}[\chi(yf_\chi(x))]$.

4.4 Rates of Convergence

We now have all tools at our disposition to obtain rates of convergence to the minimizer of the expected risk which depend only on the complexity of the function class and its approximation properties in terms of the ψ -loss. Denote by $f_{\psi, \mathcal{F}}^*$ the minimizer of $\mathbf{E}[\psi(yf(x))]$ restricted to \mathcal{F} , let $f_{\psi, \mathcal{F}}^n$ be the minimizer of the empirical ψ -risk, and let $\delta(\mathcal{F}, \psi) := \mathbf{E}[yf_{\psi, \mathcal{F}}^*(x)] - \mathbf{E}[yf_\psi^*(x)]$ be the approximation error due to the restriction of f to \mathcal{F} . Then a simple telescope sum yields

$$\begin{aligned} \mathbf{E}[\chi(yf_{\psi, \mathcal{F}}^n)] &\leq \mathbf{E}[\chi(yf_\chi^*)] + 4[\mathbf{E}[\psi(yf_{\psi, \mathcal{F}}^n)] - \mathbf{E}_{\text{emp}}[\psi(yf_{\psi, \mathcal{F}}^n)]] \\ &\quad + 4[\mathbf{E}_{\text{emp}}[\psi(yf_{\psi, \mathcal{F}}^*)] - \mathbf{E}[\psi(yf_{\psi, \mathcal{F}}^*)]] + \delta(\mathcal{F}, \psi) \\ &\leq \mathbf{E}[\chi(yf_\chi^*)] + \delta(\mathcal{F}, \psi) + 4\frac{RW\gamma}{\sqrt{n}}[\sqrt{-2 \log \delta} + r/R + \sqrt{8 \log 2}]. \end{aligned} \quad (89)$$

Here γ is the effective margin of the soft-margin loss $\max(0, 1 - \gamma yf(x))$, W is an upper bound on $\|w\|$, $R \geq \|x\|$, r is the average radius, as defined in the previous section, and we assumed that b is bounded by the largest value of $\langle w, x \rangle$. A similar reasoning for logistic and exponential loss is given in Bousquet et al. [2005].

Note that we get an $O(1/\sqrt{n})$ rate of convergence regardless of the dimensionality of x . Moreover note that the rate is dominated by $RW\gamma$, that is, the classical radius-margin bound [Vapnik, 1995]. Here R is the radius of an enclosing sphere for the data and $1/(W\gamma)$ is an upper bound on the radius of the data — the soft-margin loss becomes active only for $yf(x) \leq \gamma$.

4.5 Localization and Noise Conditions

In many cases it is possible to obtain better rates of convergence than $O(1/\sqrt{n})$ by exploiting information about the magnitude of the error of misclassification and about the variance of f on \mathcal{X} . Such bounds use Bernstein-type inequalities and they lead to localized Rademacher averages [Bartlett et al., 2002, Mendelson, 2003, Bousquet et al., 2005].

Basically the slow $O(1/\sqrt{n})$ rates arise whenever the region around the Bayes optimal decision boundary is large. In this case, determining this region produces the slow rate, whereas the well-determined region could be estimated at an $O(1/n)$ rate.

Tsybakov's noise condition [Tsybakov, 2003] requires that there exist $\beta, \gamma \geq 0$ such that

$$\mathbf{P} \left\{ \left| \mathbf{P} \{y = 1|x\} - \frac{1}{2} \right| \leq t \right\} \leq \beta t^\gamma \text{ for all } t \geq 0. \quad (90)$$

Note that for $\gamma = \infty$ the condition implies that there exists some s such that $|\mathbf{P} \{y = 1|x\} - \frac{1}{2}| \geq s > 0$ almost surely. This is also known as Massart's noise condition.

The key benefit of (90) is that it implies a relationship between variance and expected value of classification loss. More specifically for $\alpha = \frac{\gamma}{1+\gamma}$ and $g : \mathcal{X} \rightarrow \mathcal{Y}$ we have

$$\mathbf{E} \left[[\{g(x) \neq y\} - \{g^*(x) \neq y\}]^2 \right] \leq c [\mathbf{E} [\{g(x) \neq y\} - \{g^*(x) \neq y\}]]^\alpha. \quad (91)$$

Here $g^*(x) := \operatorname{argmax}_y \mathbf{P}(y|x)$ denotes the Bayes optimal classifier. This is sufficient to obtain faster rates for finite sets of classifiers. For more complex function classes localization is used. See, e.g. Bousquet et al. [2005] and Bartlett et al. [2002] for more details.

5 Statistical Models and RKHS

As we have argued so far, the Reproducing Kernel Hilbert Space approach offers many advantages in machine learning: (i) powerful and flexible models can be defined, (ii) many results and algorithms for linear models in Euclidean spaces can be generalized to RKHS, (iii) learning theory assures that effective learning in RKHS is possible, for instance, by means of regularization.

In this chapter, we will show how kernel methods can be utilized in the context of *statistical models*. There are several reasons to pursue such an avenue. First of all, in conditional modeling, it is often insufficient to compute a prediction without assessing confidence and reliability. Second, when dealing with multiple or structured responses, it is important to model *dependencies between responses* in addition to the dependence on a set of covariates. Third, incomplete data, be it due to missing variables, incomplete training targets, or a model structure involving latent variables, needs to be dealt with in a principled manner. All of these issues can be addressed by using the RKHS approach to define statistical models and by combining kernels with statistical approaches such as exponential models, generalized linear models, and Markov networks.

5.1 Exponential RKHS Models

5.1.1 Exponential Models

Exponential models or *exponential families* are among the most important class of parametric models studied in statistics. Given a canonical vector of statistics Φ and a σ -finite measure ν over the sample space \mathcal{X} , an exponential model can be defined via its probability density with respect to ν (cf. Barndorff-Nielsen [1978]),

$$p(x; \theta) = \exp [\langle \theta, \Phi(x) \rangle - g(\theta)], \quad \text{where} \quad g(\theta) := \ln \int_{\mathcal{X}} e^{\langle \theta, \Phi(x) \rangle} d\nu(x). \quad (92)$$

The m -dimensional vector $\theta \in \Theta$ with $\Theta := \{\theta \in \mathbb{R}^m : g(\theta) < \infty\}$ is also called the *canonical parameter* vector. In general, there are multiple exponential representations of the same model via canonical parameters that are affinely related to one another (Murray and Rice [1993]). A representation with minimal m is called a minimal representation, in which case m is the *order* of the exponential model. One of the most important properties of exponential families is that they have sufficient statistics of fixed dimensionality, i.e. the joint density for i.i.d. random variables X_1, X_2, \dots, X_n is also exponential, the corresponding canonical statistics simply being $\sum_{i=1}^n \Phi(X_i)$. It is well-known that much of the structure of exponential models can be derived from the log partition function $g(\theta)$, in particular (cf. Lauritzen [1996])

$$\nabla_{\theta} g(\theta) = \mu(\theta) := \mathbf{E}_{\theta} [\Phi(X)], \quad \partial_{\theta}^2 g(\theta) = \mathbf{V}_{\theta} [\Phi(X)], \quad (93)$$

where μ is known as the mean-value map. Being a covariance matrix, the Hessian of g is positive semi-definite and consequently g is convex.

Maximum likelihood estimation (MLE) in exponential families leads to a particularly elegant form for the MLE equations: the expected and the observed canonical statistics agree at the MLE $\hat{\theta}$. This means, given an i.i.d. sample $\mathcal{S} = (x_i)_{i \in [n]}$,

$$\mathbf{E}_{\hat{\theta}} [\Phi(X)] = \mu(\hat{\theta}) = \frac{1}{n} \sum_{i=1}^n \Phi(x_i) =: \mathbf{E}_{\mathcal{S}} [\Phi(X)]. \quad (94)$$

5.1.2 Exponential RKHS Models

One can extend the parametric exponential model in (92) by defining a statistical model via an RKHS \mathcal{H} with generating kernel k . Linear function $\langle \theta, \Phi(\cdot) \rangle$ over \mathcal{X} are replaced with functions $f \in \mathcal{H}$, which yields an exponential RKHS model

$$p(x; f) = \exp[f(x) - g(f)], \quad f \in \mathcal{H} := \left\{ f : f(\cdot) = \sum_{x \in \mathcal{S}} \alpha_x k(\cdot, x), \mathcal{S} \subseteq \mathcal{X}, |\mathcal{S}| < \infty \right\}. \quad (95)$$

A justification for using exponential RKHS families with rich canonical statistics as a generic way to define non-parametric models stems from the fact that if the chosen kernel k is powerful enough, the associated exponential families become universal density estimators. This can be made precise using the concept of universal kernels (Steinwart [2002a], cf. Section 2).

Proposition 14 (Dense Densities) *Let \mathcal{X} be a measurable set with a fixed σ -finite measure ν and denote by \mathcal{P} a family of densities on \mathcal{X} with respect to ν such that $p \in \mathcal{P}$ is uniformly bounded from above and continuous. Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a universal kernel for \mathcal{H} . Then the exponential RKHS family of densities generated by k according to Eq. (95) are dense in \mathcal{P} in the L_{∞} sense.*

Proof Let $D := \sup_p \|p\|_{\infty}$ and chose $\eta = \eta(\epsilon, D) > 0$ appropriately (see below). By the universality assumption for every $p \in \mathcal{P}$ there exists $f \in \mathcal{H}$ such that $\|f - \ln p\|_{\infty} < \eta$. Exploiting the strict convexity of the exp-function one gets that $e^{-\eta} p < e^f < e^{\eta} p$, which directly implies $e^{-\eta} < e^{g(f)} = \int_{\mathcal{X}} e^{f(x)} d\nu(x) < e^{\eta}$. Moreover $\|e^f - p\|_{\infty} < D(e^{\eta} - 1)$, and hence

$$\|p(f) - p\|_{\infty} \leq \|e^f - p\|_{\infty} + \|e^{f-g(f)} - e^f\|_{\infty} < D(e^{\eta} - 1)(1 + e^{\eta}) < 2De^{\eta}(e^{\eta} - 1),$$

where we utilized the upper bound

$$\|e^{f-g(f)} - e^f\|_{\infty} \leq |e^{-g(f)} - 1| \cdot \|e^f\|_{\infty} < (e^{\eta} - 1)\|e^f\|_{\infty} < De^{\eta}(e^{\eta} - 1).$$

Equating with ϵ and solving for η results in $\eta(\epsilon, D) = \ln \left(\frac{1}{2} + \frac{1}{2} \sqrt{1 + \frac{2\epsilon}{D}} \right) > 0$. ■

5.1.3 Conditional Exponential Models

For the rest of the paper, we will focus on the case of predictive or conditional modeling with a – potentially compound or structured – response variable Y and predictor variables X . Taking up the concept of joint kernels introduced in the previous section, we will investigate conditional models that are defined by functions $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ from some RKHS \mathcal{H} over $\mathcal{X} \times \mathcal{Y}$ with kernel k as follows

$$p(y|x; f) = \exp [f(x, y) - g(x, f)], \quad \text{where} \quad g(x, f) := \ln \int_{\mathcal{Y}} e^{f(x, y)} d\nu(y). \quad (96)$$

Notice that in the finite-dimensional case we have a feature map $\Phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^m$ from which parametric models are obtained via $\mathcal{H} := \{f : \exists w, f(x, y) = f(x, y; w) := \langle w, \Phi(x, y) \rangle\}$ and each f can be identified with its parameter w . Let us discuss some concrete examples to illustrate the rather general model equation (96).

- Let Y be univariate and define $\Phi(x, y) = y\Phi(x)$. Then simply $f(x, y; w) = \langle w, \Phi(x, y) \rangle = y\tilde{f}(x; w)$, with $\tilde{f}(x; w) := \langle w, \Phi(x) \rangle$ and the model equation in (96) reduces to

$$p(y|x; w) = \exp [y \langle w, \Phi(x) \rangle - g(x, w)]. \quad (97)$$

This is a *generalized linear model* (GLM) (Nelder and Wedderburn [1972], McCullagh and Nelder [1983]) with a canonical link, i.e. the canonical parameters depend linearly on the covariates $\Phi(x)$. For different response scales we get several well-known models such as, for instance, logistic regression where $y \in \{-1, 1\}$.

- In the non-parametric extension of generalized linear models following [Green and Yandell, 1985, O’Sullivan et al., 1986] the parametric assumption on the linear predictor $\tilde{f}(x; w) = \langle w, \Phi(x) \rangle$ in GLMs is relaxed by requiring that \tilde{f} comes from some sufficiently smooth class of functions, namely a RKHS defined over \mathcal{X} . In combination with a parametric part, this can also be used to define semi-parametric models. Popular choices of kernels include the ANOVA kernel investigated in Wahba et al. [1995]. This is a special case of defining joint kernels from an existing kernel k over inputs via $k((x, y), (x', y')) := yy'k(x, x')$.
- Joint kernels provide a powerful framework for prediction problems with structured outputs. An illuminating example is statistical natural language parsing with lexicalized probabilistic context free grammars (cf. Magerman [1996]). Here x will be an English sentence and y a parse tree for x , i.e. a highly structured and complex output. The productions of the grammar are known, but the conditional probability $p(y|x)$ needs to be estimated based on training data of parsed/annotated sentences. In the simplest case, the extracted statistics Φ may encode the frequencies of the use of different productions in a sentence with known parse tree. More sophisticated feature encodings are discussed in Taskar et al. [2004], Zettlemoyer and Collins [2005]. The conditional modeling approach provide alternatives to state-of-the art approaches that estimate joint models $p(x, y)$ with maximum likelihood or maximum entropy (Charniak [2000]) and obtain predictive models by conditioning on x .

5.1.4 Risk Functions for Model Fitting

There are different inference principles to determine the optimal function $f \in \mathcal{H}$ for the conditional exponential model in (96). One standard approach to parametric model fitting is to maximize the conditional log-likelihood – or equivalently – minimize a logarithmic loss, a strategy pursued in the Conditional Random Field (CRF) approach of Lafferty et al. [2001]. Here we consider the more general case of minimizing a functional that includes a monotone function of the Hilbert space norm $\|f\|_{\mathcal{H}}$ as a stabilizer (cf. Wahba [1990]). This reduces to penalized log-likelihood estimation in the finite dimensional case,

$$C^{\text{ll}}(f; \mathcal{S}) := -\frac{1}{n} \sum_{i=1}^n \ln p(y_i|x_i; f), \quad \hat{f}^{\text{ll}}(\mathcal{S}) := \underset{f \in \mathcal{H}}{\operatorname{argmin}} \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2 + C^{\text{ll}}(f; \mathcal{S}). \quad (98)$$

- For the parametric case, Lafferty et al. [2001] have employed variants of improved iterative scaling (Della Pietra et al. [1997], Darroch and Ratchiff [1972]) to optimize Eq.(98) whereas Sha and Pereira [2003] have investigated preconditioned conjugate gradient descent and limited memory quasi-Newton methods.
- In order to optimize Eq.(98) one usually needs to compute expectations of the canonical statistics $\mathbf{E}_f[\Phi(Y, x)]$ at sample points $x = x_i$, which requires the availability of efficient inference algorithms.

As we have seen in the case of classification and regression, likelihood-based criteria are by no means the only justifiable choice and large margin methods offer an interesting alternative. To that extend, we will present a general formulation of large margin methods for response variables over finite sample spaces that is based on the approach suggested by Altun et al. [2003] and Taskar et al. [2003]. Define

$$r(x, y; f) := f(x, y) - \max_{y' \neq y} f(x, y') = \min_{y' \neq y} \log \frac{p(y|x; f)}{p(y'|x; f)} \quad \text{and} \quad r(\mathcal{S}; f) := \min_{i=1}^n r(x_i, y_i; f). \quad (99)$$

Here $r(\mathcal{S}; f)$ generalizes the notion of separation margin used in SVMs. Since the log-odds ratio is sensitive to rescaling of f , i.e. $r(x, y; \beta f) = \beta r(x, y; f)$, we need to constrain $\|f\|_{\mathcal{H}}$ to make the problem well-defined. We thus replace f by $\phi^{-1}f$ for some fixed dispersion parameter $\phi > 0$ and define the maximum margin problem $\hat{f}^{\text{mm}}(\mathcal{S}) := \phi^{-1} \operatorname{argmax}_{\|f\|_{\mathcal{H}}=1} r(\mathcal{S}; f/\phi)$. For the sake of the presentation, we will drop ϕ in the following.⁹ Using the same line of arguments as was used in Section 3, the maximum margin problem can be re-formulated as a constrained optimization problem

$$\hat{f}^{\text{mm}}(\mathcal{S}) := \operatorname{argmin}_{f \in \mathcal{H}} \frac{1}{2} \|f\|_{\mathcal{H}}^2, \quad \text{s.t.} \quad r(x_i, y_i; f) \geq 1, \forall i \in [n], \quad (100)$$

provided the latter is feasible, i.e. if there exists $f \in \mathcal{H}$ such that $r(\mathcal{S}; f) > 0$. To make the connection to SVMs consider the case of binary classification with $\Phi(x, y) = y\Phi(x)$, $f(x, y; w) = \langle w, y\Phi(x) \rangle$, where $r(x, y; f) = \langle w, y\Phi(x) \rangle - \langle w, -y\Phi(x) \rangle = 2y \langle w, \Phi(x) \rangle = 2\rho(x, y; w)$. The latter is twice the standard margin for binary classification in SVMs.

A soft margin version can be defined based on the Hinge loss as follows

$$C^{\text{hl}}(f; \mathcal{S}) := \frac{1}{n} \sum_{i=1}^n \min\{1 - r(x_i, y_i; f), 0\}, \quad \hat{f}^{\text{sm}}(\mathcal{S}) := \operatorname{argmin}_{f \in \mathcal{H}} \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2 + C^{\text{hl}}(f, \mathcal{S}). \quad (101)$$

- An equivalent formulation using slack variables ξ_i as discussed in Section 3 can be obtained by introducing soft-margin constraints $r(x_i, y_i; f) \geq 1 - \xi_i$, $\xi_i \geq 0$ and by defining $C^{\text{hl}} = \frac{1}{n} \xi_i$. Each non-linear constraint can be further expanded into $|\mathcal{Y}|$ linear constraints $f(x_i, y_i) - f(x_i, y) \geq 1 - \xi_i$ for all $y \neq y_i$.
- Prediction problems with structured outputs often involve task-specific loss function $\Delta : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ discussed in Section 3.4. As suggested in Taskar et al. [2003] cost sensitive large margin methods can be obtained by defining re-scaled margin constraints $f(x_i, y_i) - f(x_i, y) \geq \Delta(y_i, y) - \xi_i$.
- Another sensible option in the parametric case is to minimize an exponential risk function of the following type

$$\hat{f}^{\text{exp}}(\mathcal{S}) := \operatorname{argmin}_w \frac{1}{n} \sum_{i=1}^n \sum_{y \neq y_i} \exp [f(x_i, y_i; w) - f(x_i, y; w)]. \quad (102)$$

This is related to the exponential loss used in the AdaBoost method of Freund and Schapire [1996]. Since we are mainly interested in kernel-based methods here, we refrain from further elaborating on this connection.

⁹We will not deal with the problem of how to estimate ϕ here; note, however, that one does need to know ϕ in order to make an optimal deterministic prediction.

5.1.5 Generalized Representer Theorem and Dual Soft-Margin Formulation

It is crucial to understand how the representer theorem applies in the setting of arbitrary discrete output spaces, since a finite representation for the optimal $\hat{f} \in \{\hat{f}^{\text{ll}}, \hat{f}^{\text{sm}}\}$ is the basis for constructive model fitting. Notice that the regularized log-loss as well as the soft margin functional introduced above depend not only on the values of f on the sample \mathcal{S} , but rather on the evaluation of f on the augmented sample $\tilde{\mathcal{S}} := \{(x_i, y) : i \in [n], y \in \mathcal{Y}\}$. This is the case, because for each x_i , output values $y \neq y_i$ not observed with x_i show up in the log-partition function $g(x_i, f)$ in (96) as well as in the log-odds ratios in (99). This adds an additional complication compared to binary classification.

Corollary 15 *Denote by \mathcal{H} an RKHS on $\mathcal{X} \times \mathcal{Y}$ with kernel k and let $\mathcal{S} = ((x_i, y_i))_{i \in [n]}$. Furthermore let $C(f; \mathcal{S})$ be a functional depending on f only via its values on the augmented sample $\tilde{\mathcal{S}}$. Let Ω be a strictly monotonically increasing function. Then the solution of the optimization problem $\hat{f}(\mathcal{S}) := \operatorname{argmin}_{f \in \mathcal{H}} C(f; \tilde{\mathcal{S}}) + \Omega(\|f\|_{\mathcal{H}})$ can be written as:*

$$\hat{f}(\cdot) = \sum_{i=1}^n \sum_{y \in \mathcal{Y}} \beta_{iy} k(\cdot, (x_i, y)) \quad (103)$$

This follows directly from Theorem 11.

Let us focus on the soft margin maximizer \hat{f}^{sm} . Instead of solving (101) directly, we first derive the dual program, following essentially the derivation in Section 3.

Proposition 16 (Tsochantaridis et al. [2005]) *The minimizer $\hat{f}^{\text{sm}}(\mathcal{S})$ can be written as in Corollary 15 where the expansion coefficients can be computed from the solution of the following convex quadratic program*

$$\alpha^* = \operatorname{argmin}_{\alpha} \left\{ \frac{1}{2} \sum_{i,j=1}^n \sum_{y \neq y_i} \sum_{y' \neq y_j} \alpha_{iy} \alpha_{jy'} K_{iy, jy'} - \sum_{i=1}^n \sum_{y \neq y_i} \alpha_{iy} \right\} \quad (104a)$$

$$\text{s.t. } \lambda n \sum_{y \neq y_i} \alpha_{iy} \leq 1, \quad \forall i \in [n]; \quad \alpha_{iy} \geq 0, \quad \forall i \in [n], y \in \mathcal{Y}, \quad (104b)$$

where $K_{iy, jy'} := k((x_i, y_i), (x_j, y_j)) + k((x_i, y), (x_j, y')) - k((x_i, y_i), (x_j, y')) - k((x_i, y), (x_j, y_j))$.

Proof *The primal program for soft-margin maximization can be expressed as*

$$\min_{f \in \mathcal{H}, \xi} \mathcal{R}(f, \xi, \mathcal{S}) := \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2 + \frac{1}{n} \sum_{i=1}^n \xi_i, \quad \text{s.t. } f(x_i, y_i) - f(x_i, y) \geq 1 - \xi_i, \quad \forall y \neq y_i; \quad \xi_i \geq 0, \quad \forall i \in [n].$$

Introducing dual variables α_{iy} for the margin constraints and η_i for the non-negativity constraints on ξ_i yields the Lagrange function (the objective has been divided by λ)

$$\mathcal{L}(f, \alpha) := \frac{1}{2} \langle f, f \rangle_{\mathcal{H}} + \frac{1}{\lambda n} \sum_{i=1}^n \xi_i - \sum_{i=1}^n \sum_{y \neq y_i} \alpha_{iy} [f(x_i, y_i) - f(x_i, y) - 1 + \xi_i] - \sum_{i=1}^n \xi_i \eta_i.$$

Solving for f results in

$$\nabla_f \mathcal{L}(f, \alpha) = 0 \iff f(\cdot) = \sum_{i=1}^n \sum_{y \neq y_i} \alpha_{iy} [k(\cdot, (x_i, y_i)) - k(\cdot, (x_i, y))]$$

since for $f \in \mathcal{H}$, $\nabla_f f(x, y) = k(\cdot, (x, y))$. Solving for ξ_i implies $\lambda n \sum_{y \neq y_i} \alpha_{iy} \leq 1$ and plugging the solution back into \mathcal{L} yields the (negative) objective function in the claim. Finally note that the representation in (103) can be obtained by identifying

$$\beta_{iy} := \begin{cases} -\alpha_{iy} & \text{if } y \neq y_i \\ \sum_{y \neq y_i} \alpha_{iy} & \text{if } y = y_i \end{cases}$$

■

- The multiclass SVM formulation of Crammer and Singer [2001] can be recovered as a special case for kernels that are diagonal with respect to the outputs, i.e. $k((x, y), (x', y')) = \delta_{y, y'} k(x, x')$. Notice that in this case the quadratic part in Eq. (104a) simplifies to

$$\sum_{i, j} k(x_i, x_j) \sum_y \alpha_{iy} \alpha_{jy} [1 + \delta_{y_i, y} \delta_{y_j, y} - \delta_{y_i, y} - \delta_{y_j, y}].$$

- The pairs (x_i, y) for which $\alpha_{iy} > 0$ are the *support pairs*, generalizing the notion of support vectors. As in binary SVMs their number can be much smaller than the total number of constraints. Notice also that in the final expansion contributions $k(\cdot, (x_i, y_i))$ will get non-negative weights whereas $k(\cdot, (x_i, y))$ for $y \neq y_i$ will get non-positive weights. Overall one gets a balance equation $\beta_{iy_i} - \sum_{y \neq y_i} \beta_{iy} = 0$ for every data point.

5.1.6 Sparse Approximation

Proposition 16 shows that sparseness in the representation of \hat{f}^{sm} is linked to the fact that only few α_{iy} in the solution to the dual problem in Eq. (104) are non-zero. Note that each of these Lagrange multipliers is linked to the corresponding soft margin constraint $f(x_i, y_i) - f(x_i, y) \geq 1 - \xi_i$. Hence, sparseness is achieved, if only few constraints are active at the optimal solution. While this may or may not be the case for a given sample, one can still exploit this observation to define a nested sequence of relaxations, where margin constraint are incrementally added. This corresponds to a constraint selection algorithm (cf. Bertsimas and Tsitsiklis [1997]) for the primal or – equivalently – a variable selection or column generation method for the dual program and has been investigated in Tsochantaridis et al. [2005]. Solving a sequence of increasingly tighter relaxations to a mathematical problem is also known as an *outer approximation*. In particular, one may iterate through the training examples according to some (fair) visitation schedule and greedily select constraints that are most violated at the current solution f , i.e. for the i -th instance one computes

$$\hat{y}_i = \operatorname{argmax}_{y \neq y_i} f(x_i, y) = \operatorname{argmax}_{y \neq y_i} p(y|x_i; f), \quad (105)$$

and then strengthens the current relaxation by including $\alpha_{i\hat{y}_i}$ in the optimization of the dual if $f(x_i, y_i) - f(x_i, \hat{y}_i) < 1 - \xi_i - \epsilon$. Here $\epsilon > 0$ is a pre-defined tolerance parameter. It is important to understand how many strengthening steps are necessary to achieve a reasonable close approximation to the original problem. The following theorem provides an answer:

Theorem 17 (Tsochantaridis et al. [2005]) *Let $\bar{R} = \max_{i, y} K_{iy, iy}$ and chose $\epsilon > 0$. A sequential strengthening procedure, which optimizes Eq. (101) by greedily selecting ϵ -violated constraints, will find an approximate solution where all constraints are fulfilled within a precision of ϵ , i.e. $r(x_i, y_i; f) \geq 1 - \xi_i - \epsilon$ after at most $\frac{2n}{\epsilon} \cdot \max \left\{ 1, \frac{4\bar{R}^2}{\lambda n^2 \epsilon} \right\}$ steps.*

Corollary 18 Denote by $(\hat{f}, \hat{\xi})$ the optimal solution of a relaxation of the problem in Proposition 16, minimizing $\mathcal{R}(f, \xi, \mathcal{S})$ while violating no constraint by more than ϵ (cf. Theorem 17). Then

$$\mathcal{R}(\hat{f}, \hat{\xi}, \mathcal{S}) \leq \mathcal{R}(\hat{f}^{sm}, \xi^*, \mathcal{S}) \leq \mathcal{R}(\hat{f}, \hat{\xi}, \mathcal{S}) + \epsilon$$

where (\hat{f}^{sm}, ξ^*) is the optimal solution of the original problem.

Proof The first inequality follows from the fact that $(\hat{f}, \hat{\xi})$ is an optimal solution to a relaxation, the second from the observation that by setting $\tilde{\xi} = \hat{\xi} + \epsilon$ one gets an admissible solution $(\hat{f}, \tilde{\xi})$ such that $\mathcal{R}(\hat{f}, \tilde{\xi}, \mathcal{S}) = \frac{\lambda}{2} \|\hat{f}\|_{\mathcal{H}}^2 + \frac{1}{n} \sum_{i=1}^n (\hat{\xi}_i + \epsilon) = \mathcal{R}(\hat{f}, \hat{\xi}, \mathcal{S}) + \epsilon$. \blacksquare

- Combined with an efficient QP solver, the above theorem guarantees a runtime polynomial in n , ϵ^{-1} , \bar{R} , and λ^{-1} . This holds irrespective of special properties of the data set utilized, the only exception being the dependency on the sample points x_i is through the radius \bar{R} .
- The remaining key problem is how to compute Eq. (105) efficiently. The answer depends on the specific form of the joint kernel k and/or the feature map Φ . In many cases, efficient dynamic programming techniques exists, whereas in other cases one has to resort to approximations or use other methods to identify a set of candidate distractors $\mathcal{Y}_i \subset \mathcal{Y}$ for a training pair (x_i, y_i) (cf. Collins [2000]). Sometimes one may also have search heuristics available that may not find the solution to Eq. (105), but that find (other) ϵ -violating constraints with a reasonable computational effort.

5.1.7 Generalized Gaussian Processes Classification

The model Eq. (96) and the minimization of the regularized log-loss can be interpreted as a generalization of Gaussian process classification (Rasmussen and Williams [2006], Altun et al. [2004b]) by assuming that $(f(x, \cdot))_{x \in \mathcal{X}}$ is a vector-valued zero mean Gaussian process; note that the covariance function C is defined over pairs $\mathcal{X} \times \mathcal{Y}$. For a given sample \mathcal{S} , define a multi-index vector $F(\mathcal{S}) := (f(x_i, y))_{i,y}$ as the restriction of the stochastic process f to the augmented sample $\tilde{\mathcal{S}}$. Denote the kernel matrix by $K = (K_{iy, jy'})$, where $K_{iy, jy'} := C((x_i, y), (x_j, y'))$ with indices $i, j \in [n]$ and $y, y' \in \mathcal{Y}$, so that in summary: $F(\mathcal{S}) \sim \mathcal{N}(0, K)$. This induces a predictive model via Bayesian model integration according to

$$p(y|x; \mathcal{S}) = \int p(y|F(x, \cdot))p(F|\mathcal{S})dF, \quad (106)$$

where x is a test point that has been included in the sample (transductive setting). For an i.i.d. sample the log-posterior for F can be written as

$$\ln p(F|\mathcal{S}) = -\frac{1}{2}F^T \mathbf{K}^{-1}F + \sum_{i=1}^n [f(x_i, y_i) - g(x_i, F)] + const. \quad (107)$$

Invoking the representer theorem for $\hat{F}(\mathcal{S}) := \operatorname{argmax}_F \ln p(F|\mathcal{S})$, we know that

$$\hat{F}(\mathcal{S})_{iy} = \sum_{j=1}^n \sum_{y' \in \mathcal{Y}} \alpha_{iy} K_{iy, jy'}, \quad (108)$$

which we plug into Eq. (107) to arrive at

$$\min_{\alpha} \alpha^T \mathbf{K} \alpha - \sum_{i=1}^n \left(\alpha^T \mathbf{K} e_{iy_i} + \log \sum_{y \in \mathcal{Y}} \exp [\alpha^T \mathbf{K} e_{iy}] \right), \quad (109)$$

were e_{iy} denotes the respective unit vector. Notice that for $f(\cdot) = \sum_{i,y} \alpha_{iy} k(\cdot, (x_i, y))$ the first term is equivalent to the squared RKHS norm of $f \in \mathcal{H}$ since $\langle f, f \rangle_{\mathcal{H}} = \sum_{i,j} \sum_{y,y'} \alpha_{iy} \alpha_{jy'} \langle k(\cdot, (x_i, y)), k(\cdot, (x_j, y')) \rangle$. The latter inner product reduces to $k((x_i, y), (x_j, y'))$ due to the reproducing property. Again, the key issue in solving Eq. (109) is how to achieve sparseness in the expansion for \hat{F} .

5.2 Markov Networks and Kernels

In Section 5.1 no assumptions about the specific structure of the joint kernel defining the model in Eq. (96) has been made. In the following, we will focus on a more specific setting with multiple outputs, where dependencies are modeled by a conditional independence graph. This approach is motivated by the fact that independently predicting individual responses based on marginal response models will often be suboptimal and explicitly modelling these interactions can be of crucial importance.

5.2.1 Markov Networks and Factorization Theorem

Denote predictor variables by X , response variables by Y and define $Z := (X, Y)$ with associated sample space \mathcal{Z} . We use Markov networks as the modeling formalism for representing dependencies between covariates and response variables as well as interdependencies among response variables.

Definition 19 A **conditional independence graph** (or *Markov network*) is an undirected graph $\mathcal{G} = (Z, E)$ such that for any pair of variables $(Z_i, Z_j) \notin E$ if and only if $Z_i \perp\!\!\!\perp Z_j \mid Z - \{Z_i, Z_j\}$.

The above definition is based on the pairwise Markov property, but by virtue of the separation theorem (cf. e.g. Whittaker [1990]) this implies the global Markov property for distributions with full support. The global Markov property says that for disjoint subsets $U, V, W \subseteq Z$ where W separates U from V in \mathcal{G} one has that $U \perp\!\!\!\perp V \mid W$. Even more important in the context of this paper is the factorization result due to Hammersley and Clifford [1971] and Besag [1974].

Theorem 20 Given a random vector Z with conditional independence graph \mathcal{G} . Any density function for Z with full support factorizes over $\mathcal{C}(\mathcal{G})$, the set of maximal cliques of \mathcal{G} as follows:

$$p(z) = \exp \left[\sum_{c \in \mathcal{C}(\mathcal{G})} f_c(z_c) \right] \quad (110)$$

where f_c are clique compatibility functions dependent on z only via the restriction on clique configurations z_c .

The significance of this result is that in order to specify a distribution for Z , one only needs to specify or estimate the simpler functions f_c .

5.2.2 Kernel Decomposition over Markov Networks

It is of interest to analyse the structure of kernels k that generate Hilbert spaces \mathcal{H} of functions that are consistent with a graph.

Definition 21 A function $f : \mathcal{Z} \rightarrow \mathbb{R}$ is **compatible** with a conditional independence graph \mathcal{G} , if f decomposes additively as $f(z) = \sum_{c \in \mathcal{C}(\mathcal{G})} f_c(z_c)$ with suitably chosen functions f_c . A Hilbert space \mathcal{H} over \mathcal{Z} is compatible with \mathcal{G} , if every function $f \in \mathcal{H}$ is compatible with \mathcal{G} . Such f and \mathcal{H} are also called \mathcal{G} -compatible.

Proposition 22 Let \mathcal{H} with kernel k be a \mathcal{G} -compatible RKHS. Then there are functions $k_{cd} : \mathcal{Z}_c \times \mathcal{Z}_d \rightarrow \mathbb{R}$ such that the kernel decomposes as

$$k(u, z) = \sum_{c, d \in \mathcal{C}} k_{cd}(u_c, z_d).$$

Proof Let $z \in \mathcal{Z}$, then $k(\cdot, z) = k(z, \cdot) \in \mathcal{H}$ and by assumption for all $u, z \in \mathcal{Z}$ there exist functions $k_d(\cdot; z)$ such that

$$k(u, z) = \sum_d k_d(u_d; z) = \sum_c k_c(z_c; u).$$

The left hand sum involves functions restricted to cliques of the first argument, whereas the right hand sum involves functions restricted to cliques of the second argument. Then the following simple lemma shows that there have to be function $k_{cd}(u_c, z_d)$ as claimed. \blacksquare

Lemma 23 Let \mathcal{X} be a set of n -tuples and $f_i, g_i : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ for $i \in [n]$ functions such that $f_i(x, y) = f_i(x_i, y)$ and $g_i(x, y) = g_i(x, y_i)$. If $\sum_i f_i(x_i, y) = \sum_j g_j(x, y_j)$ for all x, y , then there exist functions h_{ij} such that $\sum_i f_i(x_i, y) = \sum_{i,j} h_{ij}(x_i, y_j)$.

- Proposition 22 is useful for the design of kernels, since it states that only kernels allowing an additive decomposition into local functions k_{cd} are compatible with a given Markov network \mathcal{G} . Lafferty et al. [2004] have pursued a similar approach by considering kernels for RKHS with functions defined over $\mathcal{Z}_{\mathcal{C}} := \{(c, z_c) : c \in \mathcal{C}, z_c \in \mathcal{Z}_c\}$. In the latter case one can even deal with cases where the conditional dependency graph is (potentially) different for every instance.
- An illuminating example of how to design kernels via the decomposition in Proposition 22 is the case of *conditional Markov chains*, for which models based on joint kernels have been proposed in Lafferty et al. [2001], Collins [2000], Altun et al. [2003], Taskar et al. [2003]. Given an input sequences $X = (X_t)_{t \in [T]}$, the goal is to predict a sequence of labels or class variables $Y = (Y_t)_{t \in [T]}$, $Y_t \in \Sigma$. Dependencies between class variables are modeled in terms of a Markov chain, whereas outputs Y_t are assumed to depend (directly) on an observation window $(X_{t-r}, \dots, X_t, \dots, X_{t+r})$. Notice that this goes beyond the standard Hidden Markov Model structure (cf. Rabiner [1989]) by allowing for overlapping features ($r \geq 1$). For simplicity we focus on a window size of $r = 1$, in which case the clique set is given by $\mathcal{C} := \{c_t := (x_t, y_t, y_{t+1}), c'_t := (x_{t+1}, y_t, y_{t+1}) : t \in [T-1]\}$. We assume an input kernel k is given and introduce indicator vectors (or dummy variates) $I(Y_{\{t, t+1\}}) := (I_{\omega, \omega'}(Y_{\{t, t+1\}}))_{\omega, \omega' \in \Sigma}$. Now we can define the local kernel functions as

$$k_{cd}(z_c, z'_d) := \left\langle I(y_{\{s, s+1\}}), I(y'_{\{t, t+1\}}) \right\rangle \times \begin{cases} k(x_s, x_t) & \text{if } c = c_s \text{ and } d = c_t \\ k(x_{s+1}, x_{t+1}) & \text{if } c = c'_s \text{ and } d = c'_t \end{cases} \quad (111)$$

Notice that the inner product between indicator vectors is zero, unless the variable pairs are in the same configuration.

Conditional Markov chain models have found widespread applications, in natural language processing (e.g. for part of speech tagging and shallow parsing, cf. Sha and Pereira [2003]), in information retrieval (e.g. for information extraction, cf. McCallum et al. [2005]), or in computational biology (e.g. for gene prediction, cf. Culotta et al. [2005]).

5.2.3 Clique-based Sparse Approximation

Proposition 22 immediately leads to an alternative version of the representer theorem as observed by Lafferty et al. [2004] and Altun et al. [2004b].

Corollary 24 *If \mathcal{H} is \mathcal{G} -compatible then in the same setting as in Corollary 15, the optimizer \hat{f} can be written as*

$$\hat{f}(u) = \sum_{i=1}^n \sum_{c \in \mathcal{C}} \sum_{y_c \in \mathcal{Y}_c} \beta_{c,y_c}^i \sum_{d \in \mathcal{C}} k_{cd}((x_{ic}, y_c), u_d) \quad (112)$$

here x_{ic} are the variables of x_i belonging to clique c and \mathcal{Y}_c is the subspace of \mathcal{Z}_c that contains response variables.

Proof According to Proposition 23 the kernel k of \mathcal{H} can be written as $k(z, u) = \sum_{c,d} k_{cd}(z_c, u_d)$ plugging this into the expansion of Corollary 15 yields

$$\hat{f}(u) = \sum_{i=1}^n \sum_{y \in \mathcal{Y}} \beta_{iy} \sum_{c,d \in \mathcal{C}} k_{cd}((x_{ic}, y_c), u_d) = \sum_{i=1}^n \sum_{c,d \in \mathcal{C}} \sum_{y_c \in \mathcal{Y}_c} k_{cd}((x_{ic}, y_c), u_d) \sum_{y'} \delta_{y_c, y'_c} \beta_{iy'}.$$

Setting $\beta_{c,y_c}^i := \sum_{y'} \delta_{y_c, y'_c} \beta_{iy'}$ completes the proof. \blacksquare

- Notice that the number of parameters in the representation Eq. (112) scales with $n \cdot \sum_{c \in \mathcal{C}} |\mathcal{Y}_c|$ as opposed to $n \cdot |\mathcal{Y}|$ in Eq. (103). For cliques with reasonably small state spaces this will be a significantly more compact representation. Notice also that the evaluation of functions k_{cd} will typically be more efficient than evaluating k .
- In spite of this improvement, the number of terms in the expansion in Eq. (112) may in practice still be too large. In this case, one can pursue a reduced set approach, which selects a subset of variables to be included in a sparsified expansion. This has been proposed in Taskar et al. [2003] for the soft margin maximization problem as well as in Lafferty et al. [2004], Altun et al. [2004a] for conditional random fields and Gaussian processes. For instance, in Lafferty et al. [2004] parameters $\beta_{cy_c}^i$ that maximize the functional gradient of the regularized log-loss are greedily included in the reduced set. In Taskar et al. [2003] a similar selection criterion is utilized with respect to margin violations, leading to an SMO-like optimization algorithm (cf. Platt [1999]).

5.2.4 Probabilistic Inference

In dealing with structured or interdependent response variables, computing marginal probabilities of interest or computing the most probable response (cf. Eq. (105)) may be non-trivial. However, for dependency graphs with small tree width efficient inference algorithms exist, such as the junction tree algorithm (cf. Jensen et al. [1990], Dawid [1992]) and variants thereof. Notice that in the case of the conditional or hidden Markov chain, the junction tree algorithm is equivalent to the well-known forward-backward algorithm Baum [1972]. Recently, a number of approximate inference algorithms has been developed to deal with dependency graphs for which exact inference is not tractable (cf. e.g. Wainwright and Jordan [2003]).

6 Kernel Methods for Unsupervised Learning

This section discusses various methods of data analysis by modeling the distribution of data in feature space. To that extent, we study the behavior of $\Phi(x)$ by means of rather simple linear methods, which has implications for nonlinear methods on the original data space \mathcal{X} . In particular, we will discuss the extension of PCA to Hilbert spaces, which allows for image denoising, clustering, and nonlinear dimensionality reduction, the study of covariance operators for the measure of independence, the study of mean operators for the design of two-sample tests, and the modeling of complex dependencies between sets of random variables via kernel dependency estimation and canonical correlation analysis.

6.1 Kernel Principal Component Analysis

Principal Component Analysis (PCA) is a powerful technique for extracting structure from possibly high-dimensional data sets. It is readily performed by solving an eigenvalue problem, or by using iterative algorithms which estimate principal components. For reviews of the existing literature, see Jolliffe [1986] and Diamantaras and Kung [1996]; some of the classical papers are Pearson [1901], Hotelling [1933], and Karhunen [1946].

PCA is an orthogonal transformation of the coordinate system in which we describe our data. The new coordinate system is obtained by projection onto the so-called principal axes of the data. A small number of principal components is often sufficient to account for most of the structure in the data, e.g. for the purpose of principal component regression [Draper and Smith, 1981].

The basic idea is strikingly simple: denote by $X = \{x_1, \dots, x_n\}$ an n -sample drawn from $P(x)$. Then the covariance operator C is given by $C = \mathbf{E} [(x - \mathbf{E}[x])(x - \mathbf{E}[x])^\top]$. PCA aims at estimating leading eigenvectors of C via the empirical estimate $C_{\text{emp}} = \mathbf{E}_{\text{emp}} [(x - \mathbf{E}_{\text{emp}}[x])(x - \mathbf{E}_{\text{emp}}[x])^\top]$. If \mathcal{X} is d -dimensional, then the eigenvectors can be computed in $O(d^3)$ time [Press et al., 1994].

The problem can also be posed in feature space [Schölkopf et al., 1998] by replacing x with $\Phi(x)$. In this case, however, it is impossible to compute the eigenvectors directly. Yet, note that the image of C_{emp} lies in the span of $\{\Phi(x_1), \dots, \Phi(x_n)\}$. Hence it is sufficient to diagonalize C_{emp} in that subspace. In other words, we replace the *outer* product C_{emp} by an *inner* product matrix, leaving the eigenvalues unchanged, which can be computed efficiently. Using $w = \sum_{i=1}^n \alpha_i \Phi(x_i)$ it follows that α needs to satisfy $PKP\alpha = \lambda\alpha$, where P is the projection operator with $P_{ij} = \delta_{ij} - n^{-2}$ and K is the kernel matrix on X .

Note that the problem can also be recovered as one of maximizing some $\text{Contrast}[f, X]$ subject to $f \in \mathcal{F}$. This means that the projections onto the leading eigenvectors correspond to the most reliable features. This optimization problem also allows us to unify various feature extraction methods as follows:

- For $\text{Contrast}[f, X] = \text{Var}_{\text{emp}}[f, X]$ and $\mathcal{F} = \{\langle w, x \rangle \text{ subject to } \|w\| \leq 1\}$ we recover PCA.
- Changing \mathcal{F} to $\mathcal{F} = \{\langle w, \Phi(x) \rangle \text{ subject to } \|w\| \leq 1\}$ we recover Kernel PCA.
- For $\text{Contrast}[f, X] = \text{Curtosis}[f, X]$ and $\mathcal{F} = \{\langle w, x \rangle \text{ subject to } \|w\| \leq 1\}$ we have Projection Pursuit [Huber, 1985, Friedman and Tukey, 1974]. Other contrasts lead to further variants, e.g. the Epanechnikov kernel, entropic contrasts, etc. [Cook et al., 1993, Friedman, 1987, Jones and Sibson, 1987].
- If \mathcal{F} is a convex combination of basis functions and the contrast function is convex in w one obtains computationally efficient algorithms, as the solution of the optimization problem can be found at one of the vertices [Rockafellar, 1970, Schölkopf and Smola, 2002].

Subsequent projections are obtained, e.g. by seeking directions orthogonal to f or other computationally attractive variants thereof.

Kernel PCA has been applied to numerous problems, from preprocessing and invariant feature extraction [Mika et al., 2003] by simultaneous diagonalization of the data and a noise covariance matrix, to image denoising [Mika et al., 1999] and super-resolution [Kim et al., 2005]. The basic idea in the latter case is to obtain a set of principal directions in feature space w_1, \dots, w_l , obtained from noise-free data, and to project the image $\Phi(x)$ of a noisy observation x onto the space spanned by w_1, \dots, w_l . This yields a “denoised” solution $\tilde{\Phi}(x)$ in feature space. Finally, to obtain the pre-image of this denoised solution one minimizes $\|\Phi(x') - \tilde{\Phi}(x)\|$. The fact that projections onto the leading principal components turn out to be good starting points for pre-image iterations is further exploited in kernel dependency estimation (Section 6.4). Kernel PCA can be shown to contain several popular dimensionality reduction algorithms as special cases, including LLE, Laplacian Eigenmaps, and (approximately) Isomap [Ham et al., 2004].

6.2 Canonical Correlation and Measures of Independence

Given two samples X, Y canonical correlation analysis [Hotelling, 1936] aims at finding directions of projection u, v such that the correlation coefficient between X and Y is maximized. That is, (u, v) are given by

$$\operatorname{argmax}_{u,v} \operatorname{Var}_{\text{emp}} [\langle u, x \rangle]^{-1} \operatorname{Var}_{\text{emp}} [\langle v, y \rangle]^{-1} \mathbf{E}_{\text{emp}} [\langle u, x - \mathbf{E}_{\text{emp}} [x] \rangle \langle v, y - \mathbf{E}_{\text{emp}} [y] \rangle]. \quad (113)$$

This problem can be solved by finding the eigensystem of $C_x^{-\frac{1}{2}} C_{xy} C_y^{-\frac{1}{2}}$, where C_x, C_y are the covariance matrices of X and Y and C_{xy} is the covariance matrix between X and Y , respectively. Multivariate extensions are discussed in [Kettenring, 1971].

CCA can be extended to kernels by means of replacing linear projections $\langle u, x \rangle$ by projections in feature space $\langle u, \Phi(x) \rangle$. More specifically, Bach and Jordan [2002] used the so-derived contrast to obtain a measure of independence and applied it to Independent Component Analysis with great success. However, the formulation requires an additional regularization term to prevent the resulting optimization problem from becoming distribution independent. The problem is that bounding the variance of each projections is not a good way to control the capacity of the function class. Instead, one may modify the correlation coefficient by normalizing by the norms of u and v . Hence we maximize

$$\|u\|^{-1} \|v\|^{-1} \mathbf{E}_{\text{emp}} [\langle u, x - \mathbf{E}_{\text{emp}} [x] \rangle \langle v, y - \mathbf{E}_{\text{emp}} [y] \rangle], \quad (114)$$

which can be solved by diagonalization of the covariance matrix C_{xy} . One may check that in feature space, this amounts to finding the eigenvectors of $(PK_x P)^{\frac{1}{2}} (PK_y P)^{\frac{1}{2}}$ [Gretton et al., 2005b].

6.3 Measures of Independence

Rényi [1959] showed that independence between random variables is equivalent to the condition of vanishing covariance $\operatorname{Cov} [f(x), g(y)] = 0$ for all C^1 functions f, g bounded by L_∞ norm 1 on \mathcal{X} and \mathcal{Y} . In Das and Sen [1994], Dauxois and Nkiet [1998], Bach and Jordan [2002], Gretton et al. [2005c], and Gretton et al. [2005a] a constrained empirical estimate of the above criterion is used. That is, one studies

$$\Lambda(X, Y, \mathcal{F}, \mathcal{G}) := \sup_{f,g} \operatorname{Cov}_{\text{emp}} [f(x), g(y)] \text{ subject to } f \in \mathcal{F} \text{ and } g \in \mathcal{G}. \quad (115)$$

This statistic is often extended to use the entire series $\Lambda_1, \dots, \Lambda_d$ of maximal correlations where each of the function pairs (f_i, g_i) are orthogonal to the previous set of terms. More specifically Dauxois and Nkiet [1998] restrict \mathcal{F}, \mathcal{G} to finite-dimensional linear function classes subject to their L_2 norm bounded by 1, Bach and Jordan [2002] use functions in the RKHS for which some sum of the ℓ_2^n and the RKHS norm on the sample is bounded. Martin [2000], Cock and Moor [2002], and Vishwanathan et al. [2006] show that degrees of maximum correlation can also be viewed as an inner product between subspaces spanned by the observations.

Gretton et al. [2005a] use functions with bounded RKHS norm only, which provides necessary and sufficient criteria if kernels are universal. That is, $\Lambda(X, Y, \mathcal{F}, \mathcal{G}) = 0$ if and only if x and y are independent. Moreover $\operatorname{tr} PK_x PK_y P$ has the same theoretical properties and it can be computed much more easily in linear time, as it allows for incomplete Cholesky factorizations. Here K_x and K_y are the kernel matrices on X and Y respectively.

The above criteria can be used to derive algorithms for Independent Component Analysis [Bach and Jordan, 2002, Gretton et al., 2005a]. In this case one seeks rotations of the data $S = VX$ where $V \in \operatorname{SO}(d)$ such that S is coordinate-wise independent, i.e. $P(S) = \prod_{i=1}^d P(s_i)$. This approach provides the currently best performing independence criterion, although it comes at a considerable computational cost. For fast algorithms, one should consider the work of Cardoso [1998], Hyvärinen et al. [2001], Lee

et al. [2000]. Also the work of Chen and Bickel [2005] and Yang and Amari [1997] is of interest in this context.

Note that a similar approach can be used to develop two-sample tests based on kernel methods. The basic idea is that for universal kernels the map between distributions and points on the marginal polytope $\mu : p \rightarrow \mathbf{E}_{x \sim p}[\phi(x)]$ is bijective and consequently it imposes a norm on distributions. This builds on the ideas of Fortet and Mourier [1953]. The corresponding distance $d(p, q) := \|\mu[p] - \mu[q]\|$ leads to a U-statistic which allows one to compute empirical estimates of distances between distributions efficiently Borgwardt et al. [2006].

6.4 Kernel Dependency Estimation

A large part of the previous discussion revolved around estimating dependencies between sample \mathcal{X} and \mathcal{Y} for rather structured spaces \mathcal{Y} , in particular (83). In general, such dependencies can be hard to compute. Weston et al. [2003] proposed an algorithm which allows one to extend standard regularized LS regression models, as described in Section 3.3, to cases where \mathcal{Y} has complex structure.

It works by recasting the estimation problem as a *linear* estimation problem for the map $f : \Phi(x) \rightarrow \Phi(y)$ and then as a nonlinear pre-image estimation problem for finding $\hat{y} := \operatorname{argmin}_y \|f(x) - \Phi(y)\|$ as the point in \mathcal{Y} closest to $f(x)$.

Since $\Phi(y)$ is a possibly infinite dimensional space, which may lead to problems of capacity of the estimator, Weston et al. [2003] restrict f to the leading principal components of $\Phi(y)$ and subsequently they perform regularized LS regression onto each of the projections $\langle v_j, \Phi(y) \rangle$ separately, yielding functions f_j . This yields an estimate of f via $f(x) = \sum_j v_j f_j(x)$. This works well on strings and other structured data. This problem can be solved directly [Cortes et al., 2005] without the need for subspace projections. The authors apply it to the analysis of sequence data.

7 Conclusion

We have summarized some of the advances in the field of machine learning with positive definite kernels. Due to lack of space this article is by no means comprehensive, in particular, we were not able to cover statistical learning theory, which is often cited as providing theoretical support for kernel methods. However, we nevertheless hope that the main ideas that make kernel methods attractive became clear. In particular, these include the fact that kernels address the following three major issues of learning and inference:

- they formalize the notion of *similarity* of data
- they provide a *representation* of the data in an associated reproducing kernel Hilbert space
- they characterize the function class used for estimation via the representer theorem (see Eqs. (46) and (112))

We have explained a number of approaches where kernels are useful. Many of them involve the substitution of kernels for dot products, thus turning a linear geometric algorithm into a nonlinear one. This way, obtains gets SVMs from hyperplane classifiers, and kernel PCA from linear PCA. There is, however, a more recent method of constructing kernel algorithms, where the starting point is not a linear algorithm, but a linear criterion (e.g., that two random variables have zero covariance, or that the means of two samples are identical), which can be turned into a condition involving an optimization over a large function class using kernels, thus yielding tests for independence of random variables, or tests for solving the two-sample problem. We believe that these works, as well as the increasing amount of work on the use of kernel methods for structured data, illustrates that we can expect significant further progress in the field in the years to come.

Acknowledgments

The main body of this paper was written at the Mathematisches Forschungsinstitut Oberwolfach. The authors would like to express thanks for the support during that time. National ICT Australia is funded through the Australian Government's *Backing Australia's Ability* initiative, in part through the Australian Research Council. This work was supported by grants of the ARC and by the Pascal Network of Excellence. We thank Florian Steinke, Matthias Hein, Jakob Macke, Conrad Sanderson and Holger Wendland for comments and suggestions.

References

- M. A. Aizerman, É. M. Braverman, and L. I. Rozonoér. Theoretical foundations of the potential function method in pattern recognition learning. *Automation and Remote Control*, 25:821–837, 1964.
- E. L. Allwein, R. E. Schapire, and Y. Singer. Reducing multiclass to binary: a unifying approach for margin classifiers. In P. Langley, editor, *Proc. Intl. Conf. Machine Learning*, pages 9–16, San Francisco, California, 2000. Morgan Kaufmann Publishers.
- N. Alon, S. Ben-David, N. Cesa-Bianchi, and D. Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. In *Proc. of the 34rd Annual Symposium on Foundations of Computer Science*, pages 292–301. IEEE Computer Society Press, Los Alamitos, CA, 1993.
- Y. Altun, A. J. Smola, and T. Hofmann. Exponential families for conditional random fields. In *Uncertainty in Artificial Intelligence (UAI)*, pages 2–9, 2004a.
- Y. Altun, A. J. Smola, and T. Hofmann. Gaussian process classification for segmenting and annotating sequences. In *Proc. Intl. Conf. Machine Learning*, 2004b.
- Y. Altun, I. Tsochantaridis, and T. Hofmann. Hidden Markov support vector machines. In *Proc. Intl. Conf. Machine Learning*, 2003.
- N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68: 337–404, 1950.
- F. R. Bach and M. I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002.
- G. Bakir, T. Hofmann, B. Schölkopf, A. Smola, B. Taskar, and S. V. N. Vishwanathan, editors. *Predicting Structured Data*. MIT Press, Cambridge, Massachusetts, 2007.
- D. Bamber. The area above the ordinal dominance graph and the area below the receiver operating characteristic graph. *Journal of Mathematical Psychology*, 12:387–415, 1975.
- O. E. Barndorff-Nielsen. *Information and Exponential Families in Statistical Theory*. Wiley, Chichester, 1978.
- A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, May 1993.
- P. Bartlett, P. M. Long, and R. C. Williamson. Fat-shattering and the learnability of real-valued functions. *Journal of Computer and System Sciences*, 52(3):434–452, 1996.
- P. L. Bartlett, O. Bousquet, and S. Mendelson. Localized rademacher averages. In *Proc. Annual Conf. Computational Learning Theory*, pages 44–58, 2002.
- P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- J. Basilico and T. Hofmann. Unifying collaborative and content-based filtering. In *Proc. Intl. Conf. Machine Learning*, 2004.

- L. E. Baum. An inequality and associated maximization technique in statistical estimation of probabilistic functions of a Markov process. *Inequalities*, 3, 1972.
- S. Ben-David, N. Eiron, and P.M. Long. On the difficulty of approximately maximizing agreements. *J. Comput. Syst. Sci.*, 66(3):496–514, 2003.
- K. P. Bennett, A. Demiriz, and J. Shawe-Taylor. A column generation algorithm for boosting. In P. Langley, editor, *Proc. Intl. Conf. Machine Learning*, San Francisco, 2000. Morgan Kaufmann Publishers.
- K. P. Bennett and O. L. Mangasarian. Robust linear programming discrimination of two linearly inseparable sets. *Optimization Methods and Software*, 1:23–34, 1992.
- C. Berg, J. P. R. Christensen, and P. Ressel. *Harmonic Analysis on Semigroups*. Springer, New York, 1984.
- C. Berg and G. Forst. *Potential Theory on Locally Compact Abelian Groups*. Springer-Verlag, Berlin, 1975.
- D. Bertsimas and J. Tsitsiklis. *Introduction to linear programming*. Athena Scientific, 1997.
- J. Besag. Spatial interaction and the statistical analysis of lattice systems (with discussion). *Journal of the Royal Statistical Society*, 36(B):192–326, 1974.
- V. Blanz, B. Schölkopf, H. Bülthoff, C. Burges, V. Vapnik, and T. Vetter. Comparison of view-based object recognition algorithms using realistic 3D models. In C. von der Malsburg, W. von Seelen, J. C. Vorbrüggen, and B. Sendhoff, editors, *Artificial Neural Networks ICANN'96*, volume 1112 of *Lecture Notes in Computer Science*, pages 251–256, Berlin, 1996. Springer-Verlag.
- P. Bloomfield and W.L. Steiger. *Least absolute deviations: theory, applications and algorithms*. Birkhäuser, Boston, 1983.
- S. Bochner. Monotone Funktionen, Stieltjessche Integrale und harmonische Analyse. *Mathematische Annalen*, 108:378–410, 1933.
- K. M. Borgwardt, A. Gretton, M.J. Rasch, H.P. Kriegel, B. Schölkopf, and A.J. Smola. Integrating structured biological data by kernel maximum mean discrepancy. In *Proc. of Intelligent Systems in Molecular Biology (ISMB)*, Fortaleza, Brazil, 2006.
- B. Boser, I. Guyon, and V. Vapnik. A training algorithm for optimal margin classifiers. In D. Haussler, editor, *Proc. Annual Conf. Computational Learning Theory*, pages 144–152, Pittsburgh, PA, July 1992. ACM Press.
- O. Bousquet, S. Boucheron, and G. Lugosi. Theory of classification: a survey of recent advances. *ESAIM: Probability and Statistics*, 9:323 – 375, 2005.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, England, 2004.
- L. Breiman. Prediction games and arcing algorithms. *Neural Computation*, 11(7):1493–1518, 1999.
- M. P. Brown, W. N. Grundy, D. Lin, N. Cristianini, C. W. Sugnet, T. S. Furey, J. r. Ares M, and D. Haussler. Knowledge-based analysis of microarray gene expression data by using support vector machines. *Proc Natl Acad Sci U S A*, 97(1):262–267, Jan 2000.
- C. J. C. Burges. Simplified support vector decision rules. In L. Saitta, editor, *Proc. Intl. Conf. Machine Learning*, pages 71–77, San Mateo, CA, 1996. Morgan Kaufmann Publishers.
- C. J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, 1998.
- J.-F. Cardoso. Blind signal separation: statistical principles. *Proceedings of the IEEE*, 90(8):2009–2026, 1998.

- O. Chapelle, P. Haffner, and V. Vapnik. SVMs for histogram-based image classification. *IEEE Transactions on Neural Networks*, 10(5), 1999.
- O. Chapelle and Z. Harchaoui. A machine learning approach to conjoint analysis. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 17*, Cambridge, MA, 2005. MIT Press.
- E. Charniak. A maximum-entropy-inspired parser. In *North American Association of Computational Linguistics (NAACL 1)*, pages 132–139, 2000.
- A. Chen and P. Bickel. Consistent independent component analysis and prewhitening. *IEEE Transactions on Signal Processing*, 53(10):3625–3632, 2005.
- S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *Siam Journal of Scientific Computing*, 20(1):33–61, 1999.
- K. De Cock and B. De Moor. Subspace angles between ARMA models. *Systems and Control Letter*, 46:265–270, 2002.
- R. Cole and R. Hariharan. Faster suffix tree construction with missing suffix links. In *Proceedings of the Thirty Second Annual Symposium on the Theory of Computing*, Portland, OR, USA, May 2000. ACM.
- M. Collins. Discriminative reranking for natural language parsing. In *Proc. Intl. Conf. Machine Learning*, pages 175–182. Morgan Kaufmann, San Francisco, CA, 2000.
- M. Collins and N. Duffy. Convolution kernels for natural language. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, Cambridge, MA, 2001. MIT Press.
- D. Cook, A. Buja, and J. Cabrera. Projection pursuit indices based on orthonormal function expansions. *Journal of Computational and Graphical Statistics*, 2:225–250, 1993.
- C. Cortes, M. Mohri, and J. Weston. A general regression technique for learning transductions. In *ICML '05: Proceedings of the 22nd international conference on Machine learning*, pages 153–160, New York, NY, USA, 2005. ACM Press.
- C. Cortes and V. Vapnik. Support vector networks. *Machine Learning*, 20(3):273–297, 1995.
- D. Cox and F. O’Sullivan. Asymptotic analysis of penalized likelihood and related estimators. *Annals of Statistics*, 18:1676–1695, 1990.
- K. Crammer and Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2001.
- K. Crammer and Y. Singer. Loss bounds for online category ranking. In *Proc. Annual Conf. Computational Learning Theory*, 2005.
- N. Cristianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines*. Cambridge University Press, Cambridge, UK, 2000.
- N. Cristianini, J. Shawe-Taylor, A. Elisseeff, and J. Kandola. On kernel-target alignment. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 367–373, Cambridge, MA, 2002. MIT Press.
- A. Culotta, D. Kulp, and A. McCallum. Gene prediction with conditional random fields. Technical Report UM-CS-2005-028, University of Massachusetts, Amherst, 2005.
- J. N. Darroch and D. Ratcliff. Generalized iterative scaling for log-linear models. *The Annals of Mathematical Statistics*, 43, 1972.
- D. Das and P. Sen. Restricted canonical correlations. *Linear Algebra and its Applications*, 210:29–47, 1994.

- J. Dauxois and G. M. Nkiet. Nonlinear canonical analysis and independence tests. *Annals of Statistics*, 26(4):1254–1278, 1998.
- A. P. Dawid. Applications of a general propagation algorithm for probabilistic expert systems. *Statistics and Computing*, 2:25–36, 1992.
- D. DeCoste and B. Schölkopf. Training invariant support vector machines. *Machine Learning*, 46:161–190, 2002.
- O. Dekel, C.D. Manning, and Y. Singer. Log-linear models for label ranking. In S. Becker, Sebastian Thrun, and Klaus Obermayer, editors, *Advances in Neural Information Processing Systems 15*, 2003.
- S. Della Pietra, V. Della Pietra, and J. Lafferty. Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:380–393, 1997.
- K. I. Diamantaras and S. Y. Kung. *Principal Component Neural Networks*. John Wiley and Sons, New York, 1996.
- N. R. Draper and H. Smith. *Applied Regression Analysis*. John Wiley and Sons, New York, NY, 1981.
- H. Drucker, C. J. C. Burges, L. Kaufman, A. J. Smola, and V. Vapnik. Support vector regression machines. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems 9*, pages 155–161, Cambridge, MA, 1997. MIT Press.
- S. Dumais. Using SVMs for text categorization. *IEEE Intelligent Systems*, 13(4), 1998. In: M. A. Hearst, B. Schölkopf, S. Dumais, E. Osuna, and J. Platt: Trends and Controversies - Support Vector Machines.
- R. Durbin, S. Eddy, A. Krogh, and G. Mitchison. *Biological Sequence Analysis: Probabilistic models of proteins and nucleic acids*. Cambridge University Press, 1998.
- J. H. J. Einmal and D. M. Mason. Generalized quantile processes. *Annals of Statistics*, 20(2):1062–1078, 1992.
- A. Elisseeff and J. Weston. A kernel method for multi-labeled classification. In *Advances in Neural Information Processing Systems 14*. MIT press, 2001.
- M. Everingham, L. Van Gool, C. Williams, and A. Zisserman. Pascal visual object classes challenge results. Available from www.pascal-network.org, 2005.
- M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.
- C. H. FitzGerald, C. A. Micchelli, and A. Pinkus. Functions that preserve families of positive semidefinite matrices. *Linear Algebra and its Applications*, 221:83–102, 1995.
- R. Fletcher. *Practical Methods of Optimization*. John Wiley and Sons, New York, 1989.
- R. Fortet and E. Mourier. Convergence de la réparation empirique vers la réparation théorique. *Ann. Scient. École Norm. Sup.*, 70:266–285, 1953.
- Y. Freund and R. E. Schapire. Experiments with a new boosting algorithm. In *Proceedings of the International Conference on Machine Learning*, pages 148–146. Morgan Kaufmann Publishers, 1996.
- J. H. Friedman. Exploratory projection pursuit. *Journal of the American Statistical Association*, 82: 249–266, 1987.
- J. H. Friedman and J. W. Tukey. A projection pursuit algorithm for exploratory data analysis. *IEEE Transactions on Computers*, C-23(9):881–890, 1974.
- T. T. Frieß and R. F. Harrison. Linear programming support vector machines for pattern classification and regression estimation and the set reduction algorithm. TR RR-706, University of Sheffield, Sheffield, UK, 1998.

- T. Gärtner. A survey of kernels for structured data. *SIGKDD Explorations*, 5(1):49–58, July 2003.
- T. Gärtner, P. A. Flach, A. Kowalczyk, and A. J. Smola. Multi-instance kernels. In *Proc. Intl. Conf. Machine Learning*, 2002.
- F. Girosi, M. Jones, and T. Poggio. Priors, stabilizers and basis functions: From regularization to radial, tensor and additive splines. A.I. Memo 1430, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1993.
- P. Green and B. Yandell. Semi-parametric generalized linear models. In *Proceedings 2nd International GLIM Conference*, number 32 in Lecture Notes in Statistics, pages 44–55. Springer Verlag, New York, 1985.
- A. Gretton, O. Bousquet, A.J. Smola, and B. Schölkopf. Measuring statistical dependence with Hilbert-Schmidt norms. In S. Jain and W.-S. Lee, editors, *Proceedings Algorithmic Learning Theory*, 2005a.
- A. Gretton, R. Herbrich, A. Smola, O. Bousquet, and B. Schölkopf. Kernel methods for measuring independence. *Journal of Machine Learning Research*, 6:2075–2129, 2005b.
- A. Gretton, A. Smola, O. Bousquet, R. Herbrich, A. Belitski, M. Augath, Y. Murayama, J. Pauls, B. Schölkopf, and N. Logothetis. Kernel constrained covariance for dependence measurement. In *Proceedings of International Workshop on Artificial Intelligence and Statistics*, volume 10, 2005c.
- M. Gribskov and N. L. Robinson. Use of receiver operating characteristic (ROC) analysis to evaluate sequence matching. *Computers and Chemistry*, 20(1):25–33, 1996.
- D. Gusfield. *Algorithms on Strings, Trees and Sequences: Computer Science and Computational Biology*. Cambridge University Press, June 1997. ISBN 0–521–58519–8.
- J. Ham, D. Lee, S. Mika, and B. Schölkopf. A kernel view of the dimensionality reduction of manifolds. In *Proceedings of the Twenty-First International Conference on Machine Learning*, pages 369–376. 2004.
- J. M. Hammersley and P. E. Clifford. Markov fields on finite graphs and lattices. unpublished manuscript, 1971.
- J. A. Hartigan. Estimation of a convex density contour in two dimensions. *Journal of the American Statistical Association*, 82:267–270, 1987.
- D. Haussler. Convolutional kernels on discrete structures. Technical Report UCSC-CRL-99-10, Computer Science Department, UC Santa Cruz, 1999.
- P. Hayton, B. Schölkopf, L. Tarassenko, and P. Anuzis. Support vector novelty detection applied to jet engine vibration spectra. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 946–952, Cambridge, MA, 2001. MIT Press.
- M. Hein, O. Bousquet, and B. Schölkopf. Maximal margin classification for metric spaces. *Journal of Computer and System Sciences*, 71:333 – 359, 2005.
- R. Herbrich. *Learning Kernel Classifiers: Theory and Algorithms*. MIT Press, 2002.
- R. Herbrich, T. Graepel, and K. Obermayer. Large margin rank boundaries for ordinal regression. In A. J. Smola, P. L. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 115–132, Cambridge, MA, 2000. MIT Press.
- R. Herbrich and R. C. Williamson. Algorithmic luckiness. *Journal of Machine Learning Research*, 3: 175–212, 2002.
- R. Hettich and K. O. Kortanek. Semi-infinite programming: theory, methods, and applications. *SIAM Rev.*, 35(3):380–429, 1993.

- D. Hilbert. Grundzüge einer allgemeinen Theorie der linearen Integralgleichungen. *Nachrichten der Akademie der Wissenschaften in Göttingen, Mathematisch-Physikalische Klasse*, pages 49–91, 1904.
- A. E. Hoerl and R. W. Kennard. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67, 1970.
- T. Hofmann, B. Schölkopf, and A. J. Smola. A review of kernel methods in machine learning. Technical Report 156, Max-Planck-Institut für biologische Kybernetik, 2006.
- A. D. Holub, M. Welling, and P. Perona. Combining generative models and fisher kernels for object recognition. In *Proceedings of the International Conference on Computer Vision*, 2005.
- H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24:417–441 and 498 – 520, 1933.
- H. Hotelling. Relations between two sets of variates. *Biometrika*, 28:321–377, 1936.
- P. J. Huber. *Robust Statistics*. John Wiley and Sons, New York, 1981.
- P. J. Huber. Projection pursuit. *Annals of Statistics*, 13(2):435–475, 1985.
- A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. John Wiley, 2001.
- T. S. Jaakkola and D. Haussler. Probabilistic kernel regression models. In *Proceedings of the 1999 Conference on AI and Statistics*, 1999.
- T. Jebara and I. Kondor. Bhattacharyya and expected likelihood kernels. In B. Schölkopf and M. Warmuth, editors, *Proceedings of the Sixteenth Annual Conference on Computational Learning Theory*, number 2777 in LNAI, pages 57–71, 2003.
- F. V. Jensen, S. L. Lauritzen, and K. G. Olesen. Bayesian updates in causal probabilistic networks by local computation. *Computational Statistics Quarterly*, 4:269–282, 1990.
- T. Joachims. Text categorization with support vector machines: Learning with many relevant features. In *Proceedings of the European Conference on Machine Learning*, pages 137–142, Berlin, 1998. Springer.
- T. Joachims. *Learning to Classify Text Using Support Vector Machines: Methods, Theory, and Algorithms*. Kluwer Academic Publishers, Boston, 2002.
- T. Joachims. A support vector method for multivariate performance measures. In *Proc. Intl. Conf. Machine Learning*, 2005.
- I. T. Jolliffe. *Principal Component Analysis*. Springer, New York, 1986.
- M. C. Jones and R. Sibson. What is projection pursuit? *Journal of the Royal Statistical Society A*, 150(1):1–36, 1987.
- M. I. Jordan, P. L. Bartlett, and J. D. McAuliffe. Convexity, classification, and risk bounds. Technical Report 638, UC Berkeley, 2003.
- J. P. Kahane. *Some Random Series of Functions*. Cambridge University Press, 1968.
- K. Karhunen. Zur Spektraltheorie stochastischer Prozesse. *Annales Academiae Scientiarum Fennicae*, 37, 1946.
- W. Karush. Minima of functions of several variables with inequalities as side constraints. Master’s thesis, Dept. of Mathematics, Univ. of Chicago, 1939.
- H. Kashima, K. Tsuda, and A. Inokuchi. Marginalized kernels between labeled graphs. In *Proc. Intl. Conf. Machine Learning*, Washington, DC, United States, 2003.
- H. Kashima, K. Tsuda, and A. Inokuchi. Kernels on graphs. In K. Tsuda, B. Schölkopf, and J.P. Vert, editors, *Kernels and Bioinformatics*, pages 155–170, Cambridge, MA, 2004. MIT Press.
- J. R. Kettenring. Canonical analysis of several sets of variables. *Biometrika*, 58:433–451, 1971.

- K. Kim, M. O. Franz, and B. Schölkopf. Iterative kernel principal component analysis for image modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(9):1351–1366, 2005.
- G. S. Kimeldorf and G. Wahba. Some results on Tchebycheffian spline functions. *J. Math. Anal. Applic.*, 33:82–95, 1971.
- R. Koenker. *Quantile Regression*. Cambridge University Press, 2005.
- V. Koltchinskii. Rademacher penalties and structural risk minimization. *IEEE Transactions on Information Theory*, 47:1902–1914, 2001.
- I. R. Kondor and J. D. Lafferty. Diffusion kernels on graphs and other discrete structures. In *Proc. Intl. Conf. Machine Learning*, pages 315–322, 2002.
- W. Krauth and M. Mézard. Learning algorithms with optimal stability in neural networks. *Journal of Physics A*, 20:745–752, 1987.
- H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proc. 2nd Berkeley Symposium on Mathematical Statistics and Probabilistics*, pages 481–492, Berkeley, 1951. University of California Press.
- J. Lafferty, X. Zhu, and Y. Liu. Kernel conditional random fields: Representation and clique selection. In *Proc. Intl. Conf. Machine Learning*, volume 21, 2004.
- J. D. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic modeling for segmenting and labeling sequence data. In *Proc. Intl. Conf. Machine Learning*, volume 18, 2001.
- S. L. Lauritzen. *Graphical Models*. Oxford University Press, 1996.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.
- T.-W. Lee, M. Girolami, A. Bell, and T. Sejnowski. A unifying framework for independent component analysis. *Computers and Mathematics with Applications*, 39:1–21, 2000.
- C. Leslie, E. Eskin, and W. S. Noble. The spectrum kernel: A string kernel for SVM protein classification. In *Proceedings of the Pacific Symposium on Biocomputing*, pages 564–575, 2002a.
- C. Leslie, E. Eskin, J. Weston, and W. S. Noble. Mismatch string kernels for SVM protein classification. In S. Becker, Sebastian Thrun, and Klaus Obermayer, editors, *Advances in Neural Information Processing Systems 15*, volume 15, Cambridge, MA, 2002b. MIT Press.
- H. Lodhi, J. Shawe-Taylor, N. Cristianini, and C. Watkins. Text classification using string kernels. Technical Report 2000 - 79, NeuroCOLT, 2000. Published in: T. K. Leen, T. G. Dietterich and V. Tresp (eds.), *Advances in Neural Information Processing Systems 13*, MIT Press, 2001.
- M. Loève. *Probability Theory II*. Springer-Verlag, New York, 4 edition, 1978.
- D. G. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley, Reading, second edition, May 1984. ISBN 0 - 201 - 15794 - 2.
- D. M. Magerman. Learning grammatical structure using statistical decision-trees. In *Proceedings ICGI*, volume 1147 of *Lecture Notes in Artificial Intelligence*, pages 1–21. Springer, Berlin, 1996.
- Y. Makovoz. Random approximants and neural networks. *Journal of Approximation Theory*, 85:98–109, 1996.
- O. L. Mangasarian. Linear and nonlinear separation of patterns by linear programming. *Operations Research*, 13:444–452, 1965.
- R.J. Martin. A metric for ARMA processes. *IEEE Transactions on Signal Processing*, 48(4):1164–1170, 2000.
- A. McCallum, K. Bellare, and F. Pereira. A conditional random field for discriminatively-trained finite-state string edit distance. In *Conference on Uncertainty in AI (UAI)*, 2005.

- P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman and Hall, London, 1983.
- C. McDiarmid. On the method of bounded differences. In *Survey in Combinatorics*, pages 148–188. Cambridge University Press, 1989.
- S. Mendelson. Rademacher averages and phase transitions in glivenko-cantelli classes. *IEEE Transactions on Information Theory*, 2001. submitted.
- S. Mendelson. A few notes on statistical learning theory. In S. Mendelson and A. J. Smola, editors, *Advanced Lectures on Machine Learning*, number 2600 in LNAI, pages 1–40. Springer, 2003.
- J. Mercer. Functions of positive and negative type and their connection with the theory of integral equations. *Philosophical Transactions of the Royal Society, London, A* 209:415–446, 1909.
- H. Meschkowski. *Hilbertsche Räume mit Kernfunktion*. Springer-Verlag, Berlin, 1962.
- S. Mika, G. Rätsch, J. Weston, B. Schölkopf, A. J. Smola, and K.-R. Müller. Learning discriminative and invariant nonlinear features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(5): 623–628, 2003.
- S. Mika, B. Schölkopf, A. J. Smola, K.-R. Müller, M. Scholz, and G. Rätsch. Kernel PCA and denoising in feature spaces. In M. S. Kearns, S. A. Solla, and D. A. Cohn, editors, *Advances in Neural Information Processing Systems 11*, pages 536–542. MIT Press, 1999.
- M. Minsky and S. Papert. *Perceptrons: An Introduction To Computational Geometry*. MIT Press, Cambridge, MA, 1969.
- V. A. Morozov. *Methods for Solving Incorrectly Posed Problems*. Springer, 1984.
- K.-R. Müller, A. J. Smola, G. Rätsch, B. Schölkopf, J. Kohlmorgen, and V. Vapnik. Predicting time series with support vector machines. In W. Gerstner, A. Germond, M. Hasler, and J.-D. Nicoud, editors, *Artificial Neural Networks ICANN'97*, volume 1327 of *Lecture Notes in Computer Science*, pages 999–1004, Berlin, 1997. Springer-Verlag.
- M. K. Murray and J.W. Rice. *Differential geometry and statistics*, volume 48 of *Monographs on statistics and applied probability*. Chapman and Hall, 1993.
- G. Navarro and M. Raffinot. Fast regular expression search. In *Proceedings of WAE*, number 1668 in LNCS, pages 198–212. Springer, 1999.
- J. A. Nelder and R. W. M. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society, A*, 135:370–384, 1972.
- D. Nolan. The excess mass ellipsoid. *Journal of Multivariate Analysis*, 39:348–371, 1991.
- N. Oliver, B. Schölkopf, and A. J. Smola. Natural regularization in SVMs. In A. J. Smola, P. L. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 51–60, Cambridge, MA, 2000. MIT Press.
- F. O’Sullivan, B. Yandell, and W. Raynor. Automatic smoothing of regression functions in generalized linear models. *J. Amer. Statist. Assoc.*, 81:96–103, 1986.
- E. Parzen. Statistical inference on time series by RKHS methods. In R. Pyke, editor, *Proceedings 12th Biennial Seminar*, Montreal, 1970. Canadian Mathematical Congress. 1 - 37.
- K. Pearson. On lines and planes of closest fit to points in space. *Philosophical Magazine*, 2 (sixth series): 559–572, 1901.
- A. Pinkus. Strictly positive definite functions on a real inner product space. *Advances in Computational Mathematics*, 20:263–271, 2004.

- J. Platt. Fast training of support vector machines using sequential minimal optimization. In B. Schölkopf, C. J. C. Burges, and A. J. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, pages 185–208, Cambridge, MA, 1999. MIT Press.
- T. Poggio. On optimal nonlinear associative recall. *Biological Cybernetics*, 19:201–209, 1975.
- T. Poggio and F. Girosi. Networks for approximation and learning. *Proceedings of the IEEE*, 78(9), September 1990.
- W. Polonik. Minimum volume sets and generalized quantile processes. *Stochastic Processes and their Applications*, 69:1–24, 1997.
- W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C. The Art of Scientific Computation*. Cambridge University Press, 1994.
- L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–285, February 1989.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- G. Rätsch. *Robust Boosting via Convex Optimization: Theory and Applications*. PhD thesis, University of Potsdam, 2001.
- G. Rätsch, S. Mika, and A. J. Smola. Adapting codes and embeddings for polychotomies. In S. Thrun, S. Becker and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 513–520. MIT Press, Cambridge, MA, 2003.
- G. Rätsch, T. Onoda, and K. R. Müller. Soft margins for adaboost. *Machine Learning*, 42(3):287–320, 2001.
- G. Rätsch, S. Sonnenburg, J. Srinivasan, H. Witte, K.-R. Müller, R. J. Sommer, and B. Schölkopf. Improving the *Caenorhabditis elegans* genome annotation using machine learning. *PLoS Computational Biology*, 2007. In press.
- A. Rényi. On measures of dependence. *Acta Math. Acad. Sci. Hungar.*, 10:441–451, 1959.
- R. T. Rockafellar. *Convex Analysis*, volume 28 of *Princeton Mathematics Series*. Princeton University Press, 1970.
- S. Romdhani, P. H. S. Torr, B. Schölkopf, and A. Blake. Efficient face detection by a cascaded support vector machine expansion. *Proceedings of the Royal Society A*, 460:3283–3297, 2004.
- W. Rudin. *Fourier analysis on groups*. Interscience Publishers, New York, 1962.
- P. Ruján. A fast method for calculating the perceptron with maximal stability. *Journal de Physique I France*, 3:277–290, 1993.
- T. W. Sager. An iterative method for estimating a multivariate mode and isopleth. *Journal of the American Statistical Association*, 74(366):329–339, 1979.
- R. Schapire, Y. Freund, P. L. Bartlett, and W. S. Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. *Annals of Statistics*, 26:1651–1686, 1998.
- I. J. Schoenberg. Metric spaces and completely monotone functions. *Annals of Mathematics*, 39:811–841, 1938.
- B. Schölkopf. *Support Vector Learning*. R. Oldenbourg Verlag, Munich, 1997. Download: <http://www.kernel-machines.org>.
- B. Schölkopf, R. Herbrich, and A. J. Smola. A generalized representer theorem. In D. P. Helmbold and B. Williamson, editors, *Proc. Annual Conf. Computational Learning Theory*, number 2111 in Lecture Notes in Computer Science, pages 416 – 426, London, UK, 2001. Springer-Verlag.

- B. Schölkopf, J. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.
- B. Schölkopf and A. Smola. *Learning with Kernels*. MIT Press, Cambridge, MA, 2002.
- B. Schölkopf, A. J. Smola, and K.-R. Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10:1299–1319, 1998.
- B. Schölkopf, A. J. Smola, R. C. Williamson, and P. L. Bartlett. New support vector algorithms. *Neural Computation*, 12:1207–1245, 2000.
- B. Schölkopf, K. Tsuda, and J.-P. Vert. *Kernel Methods in Computational Biology*. MIT Press, Cambridge, MA, 2004.
- F. Sha and F. Pereira. Shallow parsing with conditional random fields. In *Proceedings of HLT-NAACL*, pages 213–220. Association for Computational Linguistics, 2003.
- J. Shawe-Taylor, P. L. Bartlett, R. C. Williamson, and M. Anthony. Structural risk minimization over data-dependent hierarchies. *IEEE Transactions on Information Theory*, 44(5):1926–1940, 1998.
- J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, Cambridge, UK, 2004.
- K. Sim. Context kernels for text categorization. Master’s thesis, The Australian National University, Canberra, Australia, ACT 0200, June 2001.
- A. J. Smola. Regression estimation with support vector learning machines. Diplomarbeit, Technische Universität München, 1996.
- A. J. Smola, P. L. Bartlett, B. Schölkopf, and D. Schuurmans, editors. *Advances in Large Margin Classifiers*. MIT Press, Cambridge, MA, 2000.
- A. J. Smola and I. R. Kondor. Kernels and regularization on graphs. In B. Schölkopf and M. K. Warmuth, editors, *Proc. Annual Conf. Computational Learning Theory*, Lecture Notes in Computer Science, pages 144–158. Springer, 2003.
- A. J. Smola and B. Schölkopf. On a kernel-based method for pattern recognition, regression, approximation and operator inversion. *Algorithmica*, 22:211–231, 1998.
- A. J. Smola and B. Schölkopf. Sparse greedy matrix approximation for machine learning. In P. Langley, editor, *Proc. Intl. Conf. Machine Learning*, pages 911–918, San Francisco, 2000. Morgan Kaufmann Publishers.
- A. J. Smola, B. Schölkopf, and K.-R. Müller. The connection between regularization operators and support vector kernels. *Neural Networks*, 11(5):637–649, 1998.
- I. Steinwart. On the influence of the kernel on the consistency of support vector machines. *Journal of Machine Learning Research*, 2:67–93, 2002a.
- I. Steinwart. Support vector machines are universally consistent. *J. Complexity*, 18:768–791, 2002b.
- J. Stewart. Positive definite functions and generalizations, an historical survey. *Rocky Mountain J. Math.*, 6:409–434, 1976.
- M. Stitson, A. Gammelman, V. Vapnik, V. Vovk, C. Watkins, and J. Weston. Support vector regression with ANOVA decomposition kernels. In B. Schölkopf, C. J. C. Burges, and A. J. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, pages 285–292, Cambridge, MA, 1999. MIT Press.
- I. Takeuchi, Q.V. Le, T. Sears, and A.J. Smola. Nonparametric quantile estimation. *Journal of Machine Learning Research*, 2006. forthcoming.

- B. Taskar, C. Guestrin, and D. Koller. Max-margin Markov networks. In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, 2003.
- B. Taskar, D. Klein, M. Collins, D. Koller, and C. Manning. Max-margin parsing. In *Empirical Methods in Natural Language Processing*, 2004.
- D. M. J. Tax and R. P. W. Duin. Data domain description by support vectors. In M. Verleysen, editor, *Proceedings ESANN*, pages 251–256, Brussels, 1999. D Facto.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society - Series B*, 58:267–288, 1996.
- A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Math. Dokl.*, 4:1035–1038, 1963.
- I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 2005.
- A. B. Tsybakov. On nonparametric estimation of density level sets. *Annals of Statistics*, 25(3):948–969, 1997.
- A.B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *Annals of Statistics*, 32(1): 135–166, 2003.
- C.J. van Rijsbergen. *Information Retrieval*. Butterworths, London, 2 edition, 1979.
- V. Vapnik. *Estimation of Dependences Based on Empirical Data*. Springer, Berlin, 1982.
- V. Vapnik. *The Nature of Statistical Learning Theory*. Springer, New York, 1995.
- V. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, New York, 1998.
- V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–281, 1971.
- V. Vapnik and A. Chervonenkis. The necessary and sufficient conditions for consistency in the empirical risk minimization method. *Pattern Recognition and Image Analysis*, 1(3):283–305, 1991.
- V. Vapnik, S. Golowich, and A. J. Smola. Support vector method for function approximation, regression estimation, and signal processing. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems 9*, pages 281–287, Cambridge, MA, 1997. MIT Press.
- V. Vapnik and A. Lerner. Pattern recognition using generalized portrait method. *Automation and Remote Control*, 24:774–780, 1963.
- S. V. N. Vishwanathan and A. J. Smola. Fast kernels for string and tree matching. In K. Tsuda, B. Schölkopf, and J.P. Vert, editors, *Kernels and Bioinformatics*, Cambridge, MA, 2004. MIT Press.
- S. V. N. Vishwanathan, A. J. Smola, and R. Vidal. Binet-cauchy kernels on dynamical systems and its application to the analysis of dynamic scenes. *International Journal of Computer Vision*, 2006. to appear.
- G. Wahba. *Spline Models for Observational Data*, volume 59 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, Philadelphia, 1990.
- G. Wahba, Y. Wang, C. Gu, R. Klein, and B. Klein. Smoothing spline ANOVA for exponential families, with application to the Wisconsin Epidemiological Study of Diabetic Retinopathy. *Annals of Statistics*, 23:1865–1895, 1995.
- M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. Technical Report 649, UC Berkeley, Department of Statistics, September 2003.
- C. Watkins. Dynamic alignment kernels. In A. J. Smola, P. L. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 39–50, Cambridge, MA, 2000. MIT Press.

- H. Wendland. *Scattered Data Approximation*. Cambridge University Press, Cambridge, UK, 2005.
- J. Weston, O. Chapelle, A. Elisseeff, B. Schölkopf, and V. Vapnik. Kernel dependency estimation. In S. Thrun, S. Becker and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 873–880. MIT Press, Cambridge, MA, 2003.
- J. Whittaker. *Graphical Models in Applied Multivariate Statistics*. John Wiley & Sons, 1990.
- C. K. I. Williams and M. Seeger. Using the Nystrom method to speed up kernel machines. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, 2000.
- R. C. Williamson, A. J. Smola, and B. Schölkopf. Generalization bounds for regularization networks and support vector machines via entropy numbers of compact operators. *IEEE Transaction on Information Theory*, 47(6):2516–2532, 2001.
- L. Wolf and A. Shashua. Learning over sets using kernel principal angles. *Journal of Machine Learning Research*, 4:913–931, 2003.
- P. Wolfe. A duality theorem for nonlinear programming. *Quarterly of Applied Mathematics*, 19:239–244, 1961.
- H. H. Yang and S.-I. Amari. Adaptive on-line learning algorithms for blind separation – maximum entropy and minimum mutual information. *Neural Computation*, 9(7):1457–1482, 1997.
- L. S. Zettlemoyer and M. Collins. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *Uncertainty in Artificial Intelligence UAI*, 2005.
- T. Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, 32(1):56–85, 2004.
- D. Zhou, J. Huang, and B. Schölkopf. Learning from labeled and unlabeled data on a directed graph. In *Proc. Intl. Conf. Machine Learning*, 2005.
- A. Zien, G. Rätsch, S. Mika, B. Schölkopf, T. Lengauer, and K.-R. Müller. Engineering Support Vector Machine Kernels That Recognize Translation Initiation Sites. *Bioinformatics*, 16(9):799–807, 2000.