Entropy Numbers, Operators and Support Vector Kernels

Robert C. Williamson, ANU [1] Alex J. Smola, GMD [2]
Bernhard Schölkopf, GMD [3]

NeuroCOLT2 Technical Report Series

NC2-TR-1998-023

July, 1998[4]

Produced as part of the ESPRIT Working
Group in Neural and Computational
Learning II, NeuroCOLT2 27150

For more information see the NeuroCOLT website
http://www.neurocolt.com
or email neurocolt@neurocolt.com

[1]Department of Engineering, Australian National University, Canberra, ACT 0200, Australia Bob.Williamson@anu.edu.au

[2]{smola@first.gmd.de} GMD FIRST, Rudower Chaussee 5, 12489 Berlin, Germany

[3]bs@first.gmd.de GMD FIRST, Rudower Chaussee 5, 12489 Berlin, Germany

[4]Received 07-JUL-98, Revised 22-OCT-98

Abstract

We derive new bounds for the generalization error of feature space machines, such as support vector machines and related regularization networks by obtaining new bounds on their covering numbers. The proofs are based on a viewpoint that is apparently novel in the field of statistical learning theory. The hypothesis class is described in terms of a linear operator mapping from a possibly infinite dimensional unit ball in feature space into a finite dimensional space. The covering numbers of the class are then determined via the entropy numbers of the operator. These numbers, which characterize the degree of compactness of the operator, can be bounded in terms of the eigenvalues of an integral operator induced by the kernel function used by the machine. As a consequence we are able to theoretically explain the effect of the choice of kernel functions on the generalization performance of support vector machines.

1 Introduction, Definitions and Notation

In this paper we give new bounds on the covering numbers for feature space machines. This leads to improved bounds on their generalization performance. Feature space machines perform a mapping from input space into a feature space construct regression functions or decision boundaries based on this mapping, and use constraints in feature space for capacity control. Support Vector (SV) machines, which have recently been proposed as a new class of learning algorithms solving problems of pattern recognition, regression estimation, and operator inversion [13] are a well known example of this class.[1]

A key feature of the present paper is the manner in which we *directly* bound the covering numbers of interest rather than making use of a Combinatorial dimension (such as the VC-dimension or the fat-shattering dimension) and subsequent application of a general result relating such dimensions to covering numbers. We bound covering numbers by viewing the relevant class of functions as the image of a unit ball under a particular compact operator. The results can be applied to bound the generalization performance of SV regression machines, although we do not explictly indicate the results so obtained in this brief paper.

Capacity control. In order to perform pattern recognition using linear hyperplanes, often a maximum margin of separation between the classes is sought for, as this leads to good generalization ability independent of the dimensionality [10]. It can be shown that for separable training data

$$(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \{\pm 1\},$$

[1] An extended version of this paper is available as NC2-TR-1998-019.

this is achieved by minimizing $||\mathbf{w}||_2$ subject to the constraints $y_j(\langle \mathbf{w}, \mathbf{x}_j \rangle + b) \geq 1$ for $j = 1, \ldots, m$, and some $b \in \mathbb{R}$. The decision function then takes the form

$$f(\mathbf{x}) = \mathrm{sgn}(\langle \mathbf{w}, \mathbf{x} \rangle + b).$$

Similarly, a linear regression

$$f(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle + b$$

can be estimated from data

$$(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \mathbb{R}$$

by finding the flattest function which approximates the data within some margin of error: in this case, one minimizes $||\mathbf{w}||_2$ subject to $|f(\mathbf{x}_j) - y_j| \leq \varepsilon$, where the parameter $\varepsilon > 0$ plays the role of the margin, albeit not in the space of the inputs $\mathbf{x}$, but in that of the outputs y.

Nonlinear kernels. In order to apply the above reasoning to a rather general class of *nonlinear* functions, one can use kernels computing dot products in high-dimensional spaces nonlinearly related to input space [1, 4]. Under certain conditions on a kernel k, to be stated below (Theorem 2), there exists a nonlinear map Φ into a reproducing kernel Hilbert space F such that k computes the dot product in F, i.e.

$$k(\mathbf{x}, \mathbf{y}) = \langle \Phi(\mathbf{x}), \Phi(\mathbf{y}) \rangle_F.$$

Given any algorithm which can be expressed in terms of dot products exclusively, one can thus construct a nonlinear version of it by substituting a kernel for the dot product.

By using the kernel trick for SV machines, the maximum margin idea is thus extended to a large variety of nonlinear function classes (e.g. radial basis function networks, polynomial networks, neural networks), which in the case of regression estimation comprise functions written as kernel expansions

$$f(\mathbf{x}) = \sum_{i=1}^{m} \alpha_j k(\mathbf{x}_j, \mathbf{x}) + b,$$

with $\alpha_j \in \mathbb{R},\ j = 1, \ldots, m$. It has been noticed that different kernels can be characterized by their regularization properties [12]. This provides insight into the regularization properties of SV kernels. However, it does not give us a comprehensive understanding of how to select a kernel for a given learning problem, and how using a specific kernel might influence the performance of a SV machine.

Definitions and Notation For $d \in \mathbb{N}$, $\mathbb{R}^d$ denotes the d-dimensional space of vectors $\mathbf{x} = (x_1, \ldots, x_d)$. We define spaces ℓ_p^d as follows: as vector spaces, they are identical to $\mathbb{R}^d$, in addition, they are endowed with p-norms:

$$\text{for} \quad 0 < p < \infty \quad \|\mathbf{x}\|_{\ell_p^d} := \|\mathbf{x}\|_p = \left(\sum_{j=1}^{d} |x_j|^p\right)^{1/p}$$
$$\text{for} \quad p = \infty \quad \|\mathbf{x}\|_{\ell_\infty^d} := \|\mathbf{x}\|_\infty = \max_{j=1,\dots,d} |x_j|.$$

Given m points $\mathbf{x}_1, \dots, \mathbf{x}_m \in \ell_p^d$, we use the shorthand $\mathbf{X}^m = (\mathbf{x}_1^T, \dots, \mathbf{x}_m^T)$. Suppose $\mathcal{F}$ is a class of functions defined on $\mathbb{R}^d$. The ℓ_∞^d norm *with respect to* $\mathbf{X}^m$ of $f \in \mathcal{F}$ is defined as $\|f\|_{\ell_\infty^{\mathbf{X}^m}} := \max_{i=1,\dots,m} |f(\mathbf{x}_i)|$. Given some set $\mathcal{X}$, a measure μ on $\mathcal{X}$, some $1 \le p < \infty$ and a function $f\colon \mathcal{X} \to \mathbb{R}$ we define $\|f\|_{L_p} := (\int |f(x)|^p d\mu(x))^{1/p}$ if the integral exists and $\|f\|_{L_\infty} := \operatorname{ess\,sup}_{x\in\mathcal{X}} |f(x)|$. For $1 \le p \le \infty$, we let $L_p(\mathcal{X}) := \{f\colon \mathcal{X} \to \mathbb{R}\colon \|f\|_{L_p} < \infty\}$.

Let $\mathfrak{L}(E,F)$ be the set of all bounded linear operators T between the normed spaces $(E, \|\cdot\|_E)$ and $(F, \|\cdot\|_F)$, i.e. operators such that the image of the (closed) unit ball

$$U_E := \{x \in E\colon \|x\|_E \le 1\}$$

is bounded. The smallest such bound is called the *operator norm*,

$$\|T\| := \sup_{x\in U_E} \|Tx\|_F.$$

The nth *entropy number of a set* $M \subset E$, for $n \in \mathbb{N}$, is

$$\varepsilon_n(M) := \inf\{\varepsilon > 0\colon \text{there exists an } \varepsilon\text{-cover for } M \text{ in } E \text{ containing } n \text{ or fewer points}\}.$$

The *entropy numbers of an operator* $T \in \mathfrak{L}(E,F)$ are defined as

$$\varepsilon_n(T) := \varepsilon_n(T(U_E)).$$

Note that $\varepsilon_1(T) = \|T\|$, and that $\varepsilon_n(T)$ certainly is well defined for all $n \in \mathbb{N}$ if T is a *compact operator*, i.e. if $T(U_E)$ is compact. The *dyadic entropy numbers of an operator* are defined by

$$e_n(T) := \varepsilon_{2^{n-1}}(T),\ n \in \mathbb{N}.$$

A very nice introduction to entropy numbers of operators is [6]. The ε-*covering number of* $\mathcal{F}$ *with respect to the metric* d denoted $\mathcal{N}(\varepsilon, \mathcal{F}, d)$ is the size of the smallest ε-cover for $\mathcal{F}$ using the metric d. Let

$$\mathcal{N}^m(\varepsilon, \mathcal{F}) := \sup_{\mathbf{x}_1,\dots,\mathbf{x}_m \in \mathcal{X}^m} \mathcal{N}(\varepsilon, \mathcal{F}, \ell_\infty^{\mathbf{X}^m}).$$

By log and ln, we denote the logarithms to base 2 and e, respectively. By i, we denote the imaginary unit $i = \sqrt{-1}$, k will always be a kernel, and d and m will be the input dimensionality and the number of examples

$$(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m) \in \mathbb{R}^d \times \mathbb{R},$$

respectively. We will map the input data into a feature space via a mapping $\mathbf{\Phi}$. We let $\tilde{\mathbf{x}} := \mathbf{\Phi}(\mathbf{x})$.

2 Generalization Bounds via Uniform Convergence

The generalization performance of learning machines can be bounded via uniform convergence results as in [15]. The key thing about these results is the role of the covering numbers of the hypothesis class — the focus of the present paper. Results for both classification and regression are now known. For the sake of concreteness, we quote below a result suitable for regression which was proved in [2]. Let $P_m(f) := \frac{1}{m}\sum_{i=1}^m f(\mathbf{x}_i)$ denote the *empirical mean* of f on the sample $\mathbf{x}_1, \ldots, \mathbf{x}_m$.

Lemma 1 (Alon, Ben–David, Cesa–Bianchi, and Haussler, 1997) *Let $\mathcal{F}$ be a class of functions from $\mathcal{X}$ into $[0,1]$ and let P be a distribution over $\mathcal{X}$. then, for all $\varepsilon > 0$ and all $m \geq \frac{2}{\varepsilon^2}$,*

$$\Pr\left\{\sup_{f\in\mathcal{F}} |P_m(f) - P(f)| > \varepsilon\right\} \leq 12m \cdot \mathbf{E}\left[\mathcal{N}\left(\frac{\varepsilon}{6}, \mathcal{F}, \ell_\infty^{\tilde{\mathbf{X}}^{2m}}\right)\right] e^{-\varepsilon^2 m/36} \tag{1}$$

where $\Pr$ *denotes the probability w.r.t. the sample $\mathbf{x}_1, \ldots, \mathbf{x}_m$ drawn i.i.d. from P, and* $\mathbf{E}$ *the expectation w.r.t. a second sample $\tilde{\mathbf{X}}^m = (\tilde{\mathbf{x}}_1^T, \ldots, \tilde{\mathbf{x}}_{2m}^T)$ also drawn i.i.d. from P.*

In order to use this lemma one usually makes use of the fact that

$$\mathbf{E}\left[\mathcal{N}(\varepsilon, \mathcal{F}, \ell_\infty^{\tilde{\mathbf{X}}^m})\right] \leq \mathcal{N}^m(\varepsilon, \mathcal{F}).$$

The above result can be used to give a generalization error result by applying it to the loss-function induced class using standard techniques.

3 Entropy Numbers for Kernel Machines

In the following we will mainly consider machines where the mapping into feature space is defined by Mercer kernels $k(\mathbf{x}, \mathbf{y})$ as they are easier to deal with using functional analytic methods. Such machines have become very popular due to the success of SV machines. Nonetheless in Subsection 3.3 we will show how a more direct approach could be taken towards upper–bounding entropy numbers.

3.1 Mercer's Theorem, Feature Spaces and Scaling

Our goal is to make statements about the shape of the image of the input space $\mathcal{X}$ under the feature map $\Phi(\cdot)$. We will make use of Mercer's theorem. The version stated below is a special case of the theorem proven in [9, p. 145]. In the following we will assume $(\mathcal{X}, \mu)$ to be a finite measure space, i.e. $\mu(\mathcal{X}) < \infty$. As usual, by "almost all" we mean for all elements of $\mathcal{X}^n$ except a set of μ^n-measure zero.

Theorem 2 (Mercer) *Suppose $k \in L_\infty(\mathcal{X}^2)$ such that the integral operator $T_k : L_2(\mathcal{X}) \to L_2(\mathcal{X})$,*

$$T_k f(\cdot) := \int_{\mathcal{X}} k(\cdot, \mathbf{y}) f(\mathbf{y}) d\mu(\mathbf{y}) \tag{2}$$

is positive. Let $\psi_j \in L_2(\mathcal{X})$ be the eigenfunction of T_k associated with the eigenvalue $\lambda_j \neq 0$ and normalized such that $\|\psi_j\|_{L_2} = 1$ and let $\overline{\psi_j}$ denote its complex conjugate. Then

1. $(\lambda_j(T))_j \in \ell_1$.
2. $\psi_j \in L_\infty(\mathcal{X})$ *and* $\sup_j \|\psi_j\|_{L_\infty} < \infty$.
3. $k(\mathbf{x}, \mathbf{y}) = \sum_{j \in \mathbb{N}} \lambda_j \overline{\psi_j(\mathbf{x})} \psi_j(\mathbf{y})$ *holds for almost all* $(\mathbf{x}, \mathbf{y})$*, where the series converges absolutely and uniformly for almost all* $(\mathbf{x}, \mathbf{y})$.

We will call a kernel satisfying the conditions of this theorem a *Mercer kernel*. From statement 2 of Mercer's theorem there exists some constant $C_k \in \mathbb{R}^+$ depending on $k(\cdot, \cdot)$ such that

$$|\psi_j(\mathbf{x})| \leq C_k \text{ for all } j \in \mathbb{N} \text{ and } \mathbf{x} \in \mathcal{X}. \tag{3}$$

(Actually (3) holds only for almost all $\mathbf{x} \in \mathcal{X}$, but from here on we gloss over these measure-theoretic niceties in the exposition.) Moreover from statement 3 it follows that $k(\mathbf{x}, \mathbf{y})$ corresponds to a dot product in ℓ_2 i.e. $k(\mathbf{x}, \mathbf{y}) = \langle \Phi(\mathbf{x}), \Phi(\mathbf{y}) \rangle_{\ell_2}$ with

$$\begin{array}{lll} \Phi : \mathcal{X} & \to & \ell_2 \\ \mathbf{x} & \mapsto & (\phi_j(\mathbf{x}))_j := (\sqrt{\lambda_j} \psi_j(\mathbf{x}))_j \end{array} \tag{4}$$

for almost all $\mathbf{x} \in \mathcal{X}$. In the following we will (without loss of generality) assume the sequence of $(\lambda_j)_j$ be sorted in nonincreasing order. From the argument above one can see that $\Phi(\mathcal{X})$ lives not only in ℓ_2 but in an axis parallel parallelepiped with lengths $2C_k\sqrt{\lambda_j}$.

It will be useful to consider maps that map $\Phi(\mathcal{X})$ into balls of some radius R centered at the origin. The following proposition shows that the class of all these maps is determined by elements of ℓ_2 and the sequence of eigenvalues $(\lambda_j)_j$.

Proposition 3 (Mapping $\Phi(\mathbf{x})$ into ℓ_2) *Let S be the diagonal map*

$$\begin{array}{l} S : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}^{\mathbb{N}} \\ S : (x_j)_j \mapsto S(x_j)_j = (s_j x_j)_j. \end{array} \tag{5}$$

Then S maps $\Phi(\mathcal{X})$ into a ball of finite radius R_S centered at the origin if and only if $(\sqrt{\lambda_j} s_j)_j \in \ell_2$.

Proof
($\Leftarrow$) Suppose $(s_j\sqrt{\lambda_j})_j \in \ell_2$ and let $R_S^2 := C_k^2 \|(s_j\sqrt{\lambda_j})_j\|_{\ell_2}^2 < \infty$. For any $\mathbf{x} \in \mathcal{X}$,

$$\|S\Phi(\mathbf{x})\|_{\ell_2}^2 = \sum_{j\in\mathbb{N}} s_j^2\lambda_j|\psi_j(\mathbf{x})|^2 \le \sum_{j\in\mathbb{N}} s_j^2\lambda_j C_k^2 = R_S^2. \tag{6}$$

Hence $S\Phi(\mathcal{X}) \subseteq \ell_2$.
($\Rightarrow$) Suppose $(s_j\sqrt{\lambda_j})_j$ is not in ℓ_2. Hence the sequence $(A_n)_n$ with $A_n := \sum_{j=1}^{n} s_j^2\lambda_j$ is unbounded. Now define

$$a_n(\mathbf{x}) := \sum_{j=1}^{n} s_j^2\lambda_j|\psi_j(\mathbf{x})|^2. \tag{7}$$

Then $\|a_n(\cdot)\|_{L_1(\mathcal{X})} = A_n$ due to the normalization condition on ψ_j. However, as $\mu(\mathcal{X}) < \infty$ there exists a set $\tilde{\mathcal{X}}$ of nonzero measure such that

$$a_n(\mathbf{x}) \ge \frac{A_n}{\mu(\mathcal{X})} \quad \text{for all } \mathbf{x} \in \tilde{\mathcal{X}}. \tag{8}$$

Combining the left side of (6) with (7) we obtain $\|S\Phi(\mathbf{x})\|_{\ell_2}^2 \ge a_n(\mathbf{x})$ for all $n \in \mathbb{N}$ and almost all $\mathbf{x}$. Since $a_n(\mathbf{x})$ is unbounded for a set $\tilde{\mathcal{X}}$ with nonzero measure in $\mathcal{X}$, we can see that $S\Phi(\mathcal{X}) \not\subset \ell_2$. ∎

The consequence of this result is that there exists no *axis parallel* ellipsoid $\mathcal{E}$ not completely containing the (also) axis parallel parallelepiped $\mathcal{B}$ of sidelength $(2C_k\sqrt{\lambda_j})_j$, such that $\mathcal{E}$ would contain $\Phi(\mathcal{X})$. More formally

$$\mathcal{B} \subset \mathcal{E} \text{ if and only if } \Phi(\mathcal{X}) \subset \mathcal{E}.$$

Hence $\Phi(\mathcal{X})$ contains a set of nonzero measure of elements near the corners of the parallelepiped.

Once we know that $\Phi(\mathcal{X})$ "fills" the parallelepiped described above we can use this result to construct an inverse mapping A from the unit ball in ℓ_2 to an ellipsoid $\mathcal{E}$ such that $\Phi(\mathcal{X}) \subset \mathcal{E}$ as in the following diagram.

$$\begin{tikzcd} \mathcal{X} \arrow[r, "\Phi"] & \Phi(\mathcal{X}) \arrow[r, "A^{-1}"] \arrow[d, phantom, "\cap"] & U_{\ell_2} \arrow[dl, "A"] \\ & \mathcal{E} & \end{tikzcd} \tag{9}$$

The operator A will be useful for computing the entropy numbers of concatenations of operators. (Knowing the inverse will allow us to compute the forward operator, and that can be used to bound the covering numbers of the class of functions, as shown in the next subsection.) We thus seek an operator $A : \ell_2 \to \ell_2$ such that

$$A(U_{\ell_2}) \subseteq \mathcal{E}. \tag{10}$$

We can ensure this by constructing A such that

$$A\colon (x_j)_j \mapsto (R_A a_j x_j)_j \tag{11}$$

with $R_A := C_k \|(\sqrt{\lambda_j}/a_j)_j\|_{\ell_2}$. From Proposition 3 it follows that all those operators A for which $R_A < \infty$ will satisfy (10). We call such scaling (inverse) operators *admissible*.

3.2 Entropy Numbers

The next step is to compute the entropy numbers of the operator A and use this to obtain bounds on the entropy numbers for kernel machines like SV machines. We will make use of the following theorem due to Gordon, König and Schütt [8, p. 226] (stated in the present form in [6, p. 17]).

Theorem 4 *Let $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_j \geq \cdots \geq 0$ be a non–increasing sequence of non–negative numbers and let*

$$D\mathbf{x} = (\sigma_1 x_1, \sigma_2 x_2, \ldots, \sigma_j x_j, \ldots) \tag{12}$$

for $\mathbf{x} = (x_1, x_2, \ldots, x_j, \ldots) \in \ell_p$ be the diagonal operator from ℓ_p into itself, generated by the sequence $(\sigma_j)_j$, where $1 \leq p \leq \infty$. Then for all $n \in \mathbb{N}$,

$$\sup_{j \in \mathbb{N}} n^{-\frac{1}{j}} (\sigma_1 \sigma_2 \cdots \sigma_j)^{\frac{1}{j}} \leq \varepsilon_n(D) \leq 6 \sup_{j \in \mathbb{N}} n^{-\frac{1}{j}} (\sigma_1 \sigma_2 \cdots \sigma_j)^{\frac{1}{j}}. \tag{13}$$

We can exploit the freedom in choosing A to minimize an entropy number as the following corollary shows. This will be a key ingredient of our calculation of the covering numbers for SV classes, as shown below.

Corollary 5 (Entropy numbers for $\Phi(\mathcal{X})$) *Let $k\colon \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ be a Mercer kernel and let A be defined by (11). Then*

$$\varepsilon_n(A\colon \ell_2 \to \ell_2) \leq \inf_{(a_s)_s : \left(\sqrt{\lambda_s}/a_s\right)_s \in \ell_2} \sup_{j \in \mathbb{N}} 6 C_k \left\| \left(\sqrt{\lambda_s}/a_s\right)_s \right\|_{\ell_2} n^{-\frac{1}{j}} (a_1 a_2 \cdots a_j)^{\frac{1}{j}}. \tag{14}$$

This result follows immediately by identifying D and A and exploiting the freedom that we still have in choosing a particular operator A among the class of admissible ones.

As already described in Section 1 the hypotheses that a SV machine generates can be expressed as $\langle \mathbf{w}, \tilde{\mathbf{x}} \rangle + b$ where both $\mathbf{w}$ and $\tilde{\mathbf{x}}$ are defined in the feature space $\mathcal{S} = \mathrm{span}(\Phi(\mathcal{X}))$ and $b \in \mathbb{R}$. The kernel trick as introduced by [1] was then successfully employed in [4] and [7] to extend the Optimal Margin Hyperplane classifier to what is now known as the SV machine. (The "$+b$" term is readily dealt with; we omit such considerations here though.) Consider the class

$$\mathcal{F}_{R_\mathbf{w}} := \{\langle \mathbf{w}, \tilde{\mathbf{x}} \rangle : \tilde{\mathbf{x}} \in \mathcal{S}, \|\mathbf{w}\| \leq R_\mathbf{w}\} \subseteq \mathbb{R}^{\mathcal{S}}.$$

Note that $\mathcal{F}_{R_\mathbf{w}}$ depends implicitly on k since $\mathcal{S}$ does.

What we seek are the ℓ_∞^m covering numbers for the class $\mathcal{F}_{R_\mathbf{w}}$ induced by the kernel in terms of the parameter $R_\mathbf{w}$ which is the inverse of the size of the margin in feature space, or equivalently, the size of the weight vector in feature space as defined by the dot product in $\mathcal{S}$ (see [14, 13] for details). In the following we will call such hypothesis classes with length constraint on the weight vectors in feature space *SV classes*. Let T be the operator $T = S_{\tilde{\mathbf{X}}^m} R_\mathbf{w}$ where $R_\mathbf{w} \in \mathbb{R}$ and the operator $S_{\tilde{\mathbf{X}}^m}$ is defined by

$$\begin{array}{rcl} S_{\tilde{\mathbf{X}}^m} : \ell_2 & \to & \ell_\infty^m \\ S_{\tilde{\mathbf{X}}^m} : \mathbf{w} & \mapsto & (\langle \tilde{\mathbf{x}}_1, \mathbf{w}\rangle, \ldots, \langle \tilde{\mathbf{x}}_m, \mathbf{w}\rangle). \end{array} \tag{15}$$

with $\tilde{\mathbf{x}}_j \in \Phi(\mathcal{X})$ for all j. The following theorem is useful when computing entropy numbers in terms of T and A. It is originally due to Maurey, and was extended by Carl [5]. See [16] for some extensions and historical remarks.

Theorem 6 (Carl and Stephani [6, p. 246]) *Let $S \in \mathfrak{L}(H, \ell_\infty^m)$ where H is a Hilbert space. Then there exists a constant $c > 0$ such that for all $m \in \mathbb{N}$, and $1 \le j \le m$*

$$e_n(S) \le c\|S\| \left(n^{-1} \log\left(1 + \frac{m}{n}\right)\right)^{1/2}.$$

An alternative proof of this result (given in [16]) provides a small explicit value for the constant: $c = 2(\frac{6}{2-\log 3})^{1/2} \le 5.3771$.

The restatement of Theorem 6 in terms of $\varepsilon_{2^{n-1}} = e_n$ will be useful in the following. Under the assumptions above we have

$$\varepsilon_n(S) \le c\|S\| \left((\log n + 1)^{-1} \log\left(1 + \frac{m}{\log n + 1}\right)\right)^{1/2}. \tag{16}$$

Now we can combine the bounds on entropy numbers of A and $S_{\mathbf{X}^m}$ to obtain bounds for SV classes. First we need the following lemma.

Lemma 7 (Carl and Stephani [6, p. 11]) *Let E, F, G be Banach spaces, $R \in \mathfrak{L}(F, G)$, and $S \in \mathfrak{L}(E, F)$. Then, for $n, t \in \mathbb{N}$,*

$$\varepsilon_{nt}(RS) \le \varepsilon_n(R)\varepsilon_t(S) \tag{17}$$

$$\varepsilon_n(RS) \le \varepsilon_n(R)\|S\| \tag{18}$$

$$\varepsilon_n(RS) \le \varepsilon_n(S)\|R\|. \tag{19}$$

Note that the latter two inequalities follow directly from the fact that $\varepsilon_1(R) = \|R\|$ for all $R \in \mathfrak{L}(F, G)$.

Theorem 8 (Bounds for SV classes) *Let k be a Mercer kernel, let Φ be induced via (4) and let $T := S_{\tilde{\mathbf{X}}^m} R_\mathbf{w}$ where $S_{\tilde{\mathbf{X}}^m}$ is given by (15) and $R_\mathbf{w} \in \mathbb{R}^+$. Let A be defined by (11)* *and suppose* $\tilde{\mathbf{x}}_j = \Phi(\mathbf{x}_j)$ *for* $j = 1, \ldots, m$. *Then the entropy numbers of* T *satisfy the following inequalities:*

$$\begin{aligned}
\varepsilon_n(T) &\leq c\|A\|R_{\mathbf{w}} \log^{-1/2} n \log^{-1/2}\left(1 + \tfrac{m}{\log n}\right) && (20)\\
\varepsilon_n(T) &\leq R_{\mathbf{w}} \varepsilon_n(A) && (21)\\
\varepsilon_{nt}(T) &\leq c R_{\mathbf{w}} \log^{-1/2} n \log^{-1/2}\left(1 + \tfrac{m}{\log n}\right) \varepsilon_t(A)
\end{aligned}$$

where c *is defined as in Lemma 6.*

This result gives several options for bounding $\varepsilon_n(T)$. The reason for using ε_n instead of e_n is that the index only may be integer in the former case (whereas it can be in $[1, \infty)$ in the latter), thus making it easier to obtain tighter bounds. We shall see in examples later that the best inequality to use depends on the rate of decay of the eigenvalues of k. The result gives effective bounds on $\mathcal{N}^m(\varepsilon, \mathcal{F}_{R_{\mathbf{w}}})$ since

$$\varepsilon_n(T : \ell_2 \to \ell_\infty^m) \leq \varepsilon_0 \ \Rightarrow \ \mathcal{N}^m(\varepsilon_0, \mathcal{F}_{R_{\mathbf{w}}}) \leq n.$$

Proof We will use the following factorization of T to upper bound $\varepsilon_n(T)$.

$$\begin{tikzcd}
U_{\ell_2} \arrow[r, "T"] \arrow[d, "R_{\mathbf{w}}"'] & \ell_\infty^m \\
R_{\mathbf{w}} U_{\ell_2} \arrow[r, "A"'] \arrow[ur, "S_{\tilde{\mathbf{X}}^m}"] & R_{\mathbf{w}} \mathcal{E} \arrow[u, "S_{(A^{-1}\tilde{\mathbf{X}}^m)}"']
\end{tikzcd} \tag{22}$$

The top left part of the diagram follows from the definition of T. The fact that remainder commutes stems from the fact that since A is diagonal, it is self-adjoint and so

$$\langle \mathbf{w}, \tilde{\mathbf{x}} \rangle = \langle \mathbf{w}, AA^{-1}\tilde{\mathbf{x}} \rangle = \langle A\mathbf{w}, A^{-1}\tilde{\mathbf{x}} \rangle. \tag{23}$$

Instead of computing the covering number of $T = S_{\tilde{\mathbf{X}}^m} R_{\mathbf{w}}$ directly, which is difficult or wasteful, as the the bound on $S_{\tilde{\mathbf{X}}^m}$ does not take into account that $\tilde{\mathbf{x}} \in \mathcal{E}$ but just makes the assumption of $\tilde{\mathbf{x}} \in \rho U_{\ell_2}$ for some $\rho > 0$, we will represent T as $S_{(A^{-1}\tilde{\mathbf{X}}^m)} A R_{\mathbf{w}}$. This is more efficient as we constructed A such that $A^{-1}\Phi(\mathcal{X}) \in U_{\ell_2}$ filling a larger proportion of it than just $\frac{1}{\rho}\Phi(\mathcal{X})$.

By construction of A and the Cauchy-Schwarz inequality we have $\|S_{A^{-1}\tilde{\mathbf{X}}^m}\| = 1$. Thus applying lemma 7 to the factorization of T and using Theorem 6 proves the theorem. ∎

One can give (see below) asymptotic rates of decay for $\varepsilon_n(A)$. (In fact we can determine non-asymptotic results with explicitly evaluable constants.) It is thus of some interest to give overall asymptotic rates of decay of $\varepsilon_n(T)$ in terms of the order of $\varepsilon_n(A)$.

Lemma 9 (Rate bounds on ε_n) *Let k be a Mercer kernel and suppose A is the scaling operator associated with it as defined by (11).*

1. *If* $\varepsilon_n(A) = O(\log^{-\alpha} n)$ *for some* $\alpha > 0$ *then* $\varepsilon_n(T) = O(\log^{-(\alpha+1/2)} n)$.
2. *If* $\log \varepsilon_n(A) = O(\log^{-\beta} n)$ *for some* $\beta > 0$ *then* $\log \varepsilon_n(T) = O(\log^{-\beta} n)$.

This Lemma (the proof of which is omitted) shows that in the first case, Maurey's result (theorem 6) allows an improvement in the exponent of the entropy number of T, whereas in the second, it affords none (since the entropy numbers decay so fast anyway). The Maurey result may still help in that case though for nonasymptotic n. In a nutshell we can always obtain rates of convergence better than those due to Maurey's theorem because we are not dealing with *arbitrary* mappings into infinite dimensional spaces. In fact, for logarithmic dependency of $\varepsilon_n(T)$ on n, the effect of the kernel is so strong that it completely dominates the $1/\varepsilon^2$ behaviour for arbitrary Hilbert spaces. An example of such a kernel is $k(x,y) = \exp(-(x-y)^2)$.

3.3 Empirical Bounds

Instead of theoretically determining the shape of $\Phi(\mathcal{X})$ *a priori* one could use the training and/or test data to empirically estimate its shape and use this quantity to compute an operator B_{emp} analogously to (9) which performs the mapping described above. In this subsection we will sketch a possible approach — the full development of these ideas would requires considerable further work and will be deferred to a subsequent paper. There are some remarks in the full version of this paper (NC2-TR-1998-019). Furthermore the statistical argument needed to exploit such techniques (bounding generalization error in terms of *empirical* covering numbers has now been developed — see [11].

4 Eigenvalue Decay Rates

The results presented above show that if one knows the eigenvalue sequence $(\lambda_i)_i$ of a compact operator, one can bound its entropy numbers. A commonly used kernel is $k(x,y) = e^{-(x-y)^2}$ which has noncompact support. The induced integral operator

$$(T_k f)(x) = \int_{-\infty}^{\infty} k(x,y) f(y) dy$$

then has a continuous spectrum and thus T_k is not compact [3, p.267]. The question arises: can we make use of such kernels in SV machines and still obtain generalization error bounds of the form developed above? This problem can be readily resolved by analysing the v-periodic extension of the kernel in question $k_v(x) := \sum_{j=-\infty}^{\infty} k(x - jv)$. A simple argument gives

Lemma 10 *Let $k: \mathbb{R} \to \mathbb{R}$ be a symmetric convolution kernel, let $K(\omega) = F[k(x)](\omega)$ denote the Fourier transform of $k(\cdot)$ and k_v denote the v–periodic kernel derived from k (also assume that k_v exists). Then k_v has a representation as a Fourier series with $\omega_0 := \frac{2\pi}{v}$ and*

$$k_v(x-y) = \sum_{j=-\infty}^{\infty} \frac{\sqrt{2\pi}}{v} K(j\omega_0) e^{ij\omega_0 x}$$

Moreover $\lambda_j = \sqrt{2\pi} K(j\omega_0)$ for $j \in \mathbb{Z}$ and $C_k = \sqrt{\frac{2}{v}}$.

We now consider how $\varepsilon_n(A: \ell_2 \to \ell_2)$ depends asymptotically on the eigenvalues of T_k.

Proposition 11 (Exponential–Polynomial decay) *Suppose k is a Mercer kernel with $\lambda_j = \beta^2 e^{-\alpha j^p}$ for some $\alpha, \beta, p > 0$. Then*

$$\ln \varepsilon_n^{-1}(A: \ell_2 \to \ell_2) = O(\ln^{\frac{p}{p+1}} n)$$

An example of such a kernel (for $p = 2$) is $k(x) = e^{-x^2}$. It can also be shown that the rate in the above proposition is asymptotically tight.

5 The Missing Pieces and Some Conclusions

In this short version we have omitted many details and extensions such as

Discretization How should one choose v in periodizing a non-compact kernel?

Higher Dimensions The results need to be extended to multi-dimensional kernels to be practically useful. Several additional technical complications arise in doing so.

Glueing it all Together We have given the ingredients but not baked the cake. Since the approach we have taken is new, and since there are a wide range of different uniform convergence results one may use we have refrained from putting it all together into "master generalization error theorem." It should be clear that it is *possible* to do so.

We have shown how to connect properties known about mappings into feature spaces with bounds on the covering numbers. Our reasoning relied on the fact that this mapping exhibits certain decay properties to ensure rapid convergence and a constraint on the size of the weight vector in feature space. This means that the corresponding algorithms have to restrict exactly this quantity to ensure good generalization performance. This is exactly what is done in Support Vector machines. The method used to obtain the results (reasoning via entropy numbers of operators) would seem to be a nice new viewpoint and valuable for other problems.

Acknowledgements This work was supported by the Australian Research Council and a grant of the DFG(# Ja 379/71).

References

[1] M. A. Aizerman, E. 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.

[2] N. Alon, S. Ben-David, N. Cesa-Bianchi, and D. Haussler. Scale–sensitive Dimensions, Uniform Convergence, and Learnability. *J. of the ACM*, 44(4):615–631, 1997.

[3] R. Ash. *Information Theory*. Interscience Publishers, New York, 1965.

[4] B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In D. Haussler, editor, *5th Annual ACM Workshop on COLT*, pages 144–152, Pittsburgh, PA, 1992. ACM Press.

[5] B. Carl. Inequalities of Bernstein-Jackson-type and the degree of compactness of operators in Banach spaces. *Annales de l'Institut Fourier*, 35(3):79–118, 1985.

[6] B. Carl and I. Stephani. *Entropy, compactness, and the approximation of operators*. Cambridge University Press, Cambridge, UK, 1990.

[7] C. Cortes and V. Vapnik. Support vector networks. *M. Learning*, 20:273 – 297, 1995.

[8] Y. Gordon, H. König, and C. Schütt. Geometric and probabilistic estimates for entropy and approximation numbers of operators. *Journal of Approximation Theory*, 49:219–239, 1987.

[9] H. König. *Eigenvalue Distribution of Compact Operators*. Birkhäuser, Basel, 1986.

[10] J. Shawe-Taylor, P.L. Bartlett, R.C. Williamson, and M. Anthony. A framework for structural risk minimisation. In *Proceedings of the 9th Annual Conference on Computational Learning Theory*, pages 68–76. ACM Press, 1996.

[11] J. Shawe-Taylor and R.C. Williamson. Generalization performance of classifiers in terms of observed covering numbers. In *Proceedings of EUROCOLT'99*, 1999. submitted.

[12] A.J. Smola, B. Schölkopf, and K.-R. Müller. The connection between regularization operators and support vector kernels. *Neural Networks*, 11:637–649, 1998.

[13] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer, N.Y., 1995.

[14] V. Vapnik and A. Chervonenkis. *Theory of Pattern Recognition [in Russian]*. Nauka, Moscow, 1974. (German Translation: W. Wapnik & A. Tscherwonenkis, *Theorie der Zeichenerkennung*, Akademie–Verlag, Berlin, 1979).

[15] V.N. Vapnik and A.Ya. Chervonenkis. Necessary and sufficient conditions for the uniform convergence of means to their expectations. *Theory of Probability and its Applications*, 26(3):532–553, 1981.

[16] R.C. Williamson, B. Schölkopf, and A.J. Smola. A Maximum Margin Miscellany. Typescript, March 1998.