

K -Dimensional Coding Schemes in Hilbert Spaces

Andreas Maurer¹ and Massimiliano Pontil²

¹ Adalbertstrasse 55
D-80799 München, Germany
am@andreas-maurer.eu
² Dept. of Computer Science
University College London
Malet Pl., WC1E, London, UK
m.pontil@cs.ucl.ac.uk

Abstract. This paper presents a general coding method where data in a Hilbert space are represented by finite dimensional coding vectors. The method is based on empirical risk minimization within a certain class of linear operators, which map the set of coding vectors to the Hilbert space. Two results bounding the expected reconstruction error of the method are derived, which highlight the role played by the codebook and the class of linear operators. The results are specialized to some cases of practical importance, including K -means clustering, nonnegative matrix factorization and other sparse coding methods.

Index Terms: Empirical risk minimization, estimation bounds, K -means clustering and vector quantization, statistical learning.

1 Introduction

We study a general class of K -dimensional coding methods for data drawn from a distribution μ on the unit ball of a Hilbert space H . These methods encode a data point $x \sim \mu$ as a vector $\hat{y} \in \mathbb{R}^K$, according to the formula

$$\hat{y} = \arg \min_{y \in Y} \|x - Ty\|^2,$$

where $Y \subseteq \mathbb{R}^K$ is a prescribed set of *codes* (called the *codebook*), which we can always assume to span \mathbb{R}^K , and $T : \mathbb{R}^K \rightarrow H$ is a linear map, which defines a particular *implementation* of the codebook. It embeds the codebook Y in H and yields the set $T(Y)$ of exactly codable patterns. If \hat{y} is the code found for x then $\hat{x} = T\hat{y}$ is the reconstructed data point. The quantity

$$f_T(x) = \min_{y \in Y} \|x - Ty\|^2$$

is called the reconstruction error.

Given a codebook Y and a finite number of independent observations $x_1, \dots, x_m \sim \mu$, a common sense approach searches for an implementation \hat{T} which is optimal on average over the observed points, that is

$$\hat{T} = \arg \min_{T \in \mathcal{T}} \frac{1}{m} \sum_{i=1}^m f_T(x_i), \quad (1)$$

where \mathcal{T} denotes some class of linear maps $T : \mathbb{R}^K \rightarrow H$. As we shall see, this framework is general enough to include principal component analysis, K -means clustering, non-negative matrix factorization [10] and the sparse coding method as proposed in [14].

Whenever the codebook Y is compact and \mathcal{T} is bounded in the operator norm this approach is justified by the following high-probability, uniform bound on the expected reconstruction error.

Theorem 1. *Suppose that Y is a closed subset of the unit ball of \mathbb{R}^K , that there is $c \geq 1$ such that $\|T\|_\infty \leq c$ for all $T \in \mathcal{T}$ and that $\delta \in (0, 1)$. Then with probability at least $1 - \delta$ in the observed data $x_1, \dots, x_m \sim \mu$ we have for every $T \in \mathcal{T}$ that*

$$\mathbb{E}_{x \sim \mu} f_T(x) - \frac{1}{m} \sum_{i=1}^m f_T(x_i) \leq 6c^2 K^2 \sqrt{\frac{\pi}{m}} + c^2 \sqrt{\frac{8 \ln 1/\delta}{m}}.$$

The bound is two-sided in the sense that also with probability at least $1 - \delta$ we have for every $T \in \mathcal{T}$ that

$$\frac{1}{m} \sum_{i=1}^m f_T(x_i) - \mathbb{E}_{x \sim \mu} f_T(x) \leq 6c^2 K^2 \sqrt{\frac{\pi}{m}} + c^2 \sqrt{\frac{8 \ln 1/\delta}{m}}.$$

Any compact subset of \mathbb{R}^K can of course be down-scaled to be contained in the unit ball, and the scaling factor can be absorbed in c , so that the above result is applicable to any compact codebook.

The theorem implies a bound on the excess risk: let $T_0 \in \mathcal{T}$ be a minimizer of the expected reconstruction error within the set \mathcal{T} . It follows from the definition of \hat{T} and the above result that the expected reconstruction error of \hat{T} is with high probability not more than $O(1/\sqrt{m})$ worse than that of T_0 .

This order in m is optimal, as we know from existing lower bounds for K -means clustering [3]. The above dependence on K is, however, generally not optimal, and can be considerably improved with a more careful analysis, if we are prepared to accept the slightly inferior rate of $\sqrt{\ln m/m}$ in the sample size. To state this improvement define

$$\|\mathcal{T}\|_Y = \sup_{T \in \mathcal{T}} \|T\|_Y = \sup_{T \in \mathcal{T}} \sup_{y \in Y} \|Ty\|.$$

We then have the following result.

Theorem 2. Assume that $\|\mathcal{T}\|_Y \geq 1$ and that the functions f_T for $T \in \mathcal{T}$, when restricted to the unit ball of H , have range contained in $[0, b]$. Fix $\delta > 0$.

Then with probability at least $1 - \delta$ in the observed data $x_1, \dots, x_m \sim \mu$ we have for every $T \in \mathcal{T}$ that

$$\mathbb{E}_{x \sim \mu} f_T(x) - \frac{1}{m} \sum_{i=1}^m f_T(x_i) \leq \frac{K}{\sqrt{m}} \left(14 \|\mathcal{T}\|_Y + \frac{b}{2} \sqrt{\ln(16m \|\mathcal{T}\|_Y^2)} \right) + b \sqrt{\frac{\ln 1/\delta}{2m}}.$$

The bound is two sided in the same sense as the previous result.

Both results immediately imply uniform convergence in probability. We are not aware of other results for nonnegative matrix factorization [10] or the sparse coding techniques as in [14].

Before proving our results, we will illustrate their implications in some cases of interest. It turns out that the dependence on K in Theorem 2 adapts to the specific situation under consideration.

A preliminary version of this paper appeared in the proceedings of the 2008 Algorithmic Learning Theory Conference [12]. The new version contains Theorem 1 and a simplified proof of Theorem 2 with improved constants.

2 Examples of coding schemes

Several coding schemes can be expressed in our framework. We describe some of these methods and how our result applies.

2.1 Principal component analysis

Principal component analysis (PCA) seeks a K -dimensional orthogonal projection which maximizes the projected variance and then uses this projection to encode future data. A projection P can be expressed as TT^* where T is an isometry which maps \mathbb{R}^K to the range of P . Since

$$\|Px\|^2 = \|x\|^2 - \|x - Px\|^2 = \|x\|^2 - \min_{y \in \mathbb{R}^K} \|x - Ty\|^2$$

finding P to maximize the true or empirical expectation of $\|Px\|^2$ is equivalent to finding T to minimize the corresponding expectation of $\min_{y \in \mathbb{R}^K} \|x - Ty\|^2$. We see that PCA is described by our framework upon the identifications $Y = \mathbb{R}^K$ and \mathcal{T} is restricted to the class of isometries $T : \mathbb{R}^K \rightarrow H$. Given $T \in \mathcal{T}$ and $x \in H$ the reconstruction error is

$$f_T(x) = \min_{y \in \mathbb{R}^K} \|x - Ty\|^2.$$

If the data are constrained to be in the unit ball of H , as we generally assume, then it is easily seen that we can take Y to be the unit ball of \mathbb{R}^K without changing any of the encodings. We can therefore apply Theorem 2 with $\|\mathcal{T}\|_Y = 1$

and $b = 1$. This is besides the point however, because in the simple case of PCA much better bounds are available (see [15], [19] and Lemma 6 below). In [19] local Rademacher averages are used to give faster rates under certain circumstances.

An objection to PCA is, that generic codes have K nonzero components, while for practical and theoretical reasons sparse codes with much less than K nonzero components may be preferable [14].

2.2 K -means clustering or vector quantization

Here $Y = \{e_1, \dots, e_K\}$, where the vectors e_k form an orthonormal basis of \mathbb{R}^K . An implementation T now defines a set of centers $\{Te_1, \dots, Te_K\}$, the reconstruction error is $\min_{k=1}^K \|x - Te_k\|^2$ and a data point x is coded by the e_k such that Te_k is nearest to x . The algorithm (1) becomes

$$\hat{T} = \arg \min_{T \in \mathcal{T}} \frac{1}{m} \sum_{i=1}^m \min_{k=1}^K \|x_i - Te_k\|^2.$$

It is clear that every center Te_k has at most unit norm, so that $\|\mathcal{T}\|_Y = 1$. Since all data points are in the unit ball we have $\|x - Te_k\|^2 \leq 4$ so we can set $b = 4$ and the bound in Theorem 2 becomes

$$\left(14 + 2\sqrt{\ln(16m)}\right) \frac{K}{\sqrt{m}} + \sqrt{\frac{8 \ln(1/\delta)}{m}}.$$

The order of this bound matches up to $\sqrt{\ln m}$ the order given in [4] or [16]. To illustrate our method we will also prove the bound

$$\sqrt{18\pi} \frac{K}{\sqrt{m}} + \sqrt{\frac{8 \ln(1/\delta)}{m}}$$

(Theorem 6), which is essentially the same as those in [4] or [16]. There is a lower bound of order $\sqrt{K/m}$ in [3], and it is unknown which of the two bounds (upper or lower) is tight.

In K -means clustering every code has only one nonzero component, so that sparsity is enforced in a maximal way. On the other hand this results in a weaker approximation capability of the coding scheme.

2.3 Nonnegative matrix factorization

Here Y is the positive orthant in \mathbb{R}^K , that is the cone

$$Y = \{y : y = (y_1, \dots, y_K), y_k \geq 0, 1 \leq k \leq K\}.$$

A chosen map T generates a cone $T(Y) \subset H$ onto which incoming data is projected. In the original formulation by Lee and Seung [10] it is postulated that both the data and the vectors Te_k be contained in the positive orthant

of some finite dimensional space, but we can drop most of these restrictions, keeping only the requirement that $\langle Te_k, Te_l \rangle \geq 0$ for $1 \leq k, l \leq K$.

No coding will change if we require that $\|Te_k\| = 1$ for all $1 \leq k \leq K$ by a suitable normalization. The set \mathcal{T} is then given by

$$\mathcal{T} = \{T : T \in \mathcal{L}(\mathbb{R}^K, H), \|Te_k\| = 1, \langle Te_k, Te_l \rangle \geq 0, 1 \leq k, l \leq K\}.$$

We can restrict Y to its intersection with the unit ball in \mathbb{R}^K (see Lemma 2 below). We obtain that $\|\mathcal{T}\|_Y = \sqrt{K}$. Hence, Theorem 2 yields the bound

$$\frac{K}{\sqrt{m}} \left(14\sqrt{K} + \frac{1}{2} \sqrt{\ln(16mK)} \right) + \sqrt{\frac{\ln(1/\delta)}{2m}}$$

on the estimation error. We do not know of any other generalization bounds for this coding scheme.

Nonnegative matrix factorization appears to encourage sparsity, but cases have been reported where sparsity was not observed [11]. In fact this undesirable behavior should be generic for exactly codable data. Various authors have therefore proposed additional constraints ([11], [7]). It is clear that additional constraints on \mathcal{T} can only improve estimation and that the passage from Y to a subset can only improve our bounds, because the quantity $\|\mathcal{T}\|_Y$ would decrease.

2.4 Sparse coding

Another method arises by choosing the ℓ_p -unit ball as a codebook. Let $Y = \{y : y \in \mathbb{R}^K, \|y\|_p \leq 1\}$ and $\mathcal{T} = \{T : \mathbb{R}^K \rightarrow H : \|Te_k\| \leq 1, 1 \leq k \leq K\}$. We have

$$\|Ty\| = \left\| \sum_{k=1}^K y_k Te_k \right\| \leq \sum_{k=1}^K |y_k| \|Te_k\| \leq \left(\sum_{k=1}^K \|Te_k\|^q \right)^{1/q} \leq K^{1/q} = K^{1-1/p}$$

implying that $\|\mathcal{T}\|_Y \leq K^{1-1/p}$.

By the same argument as above all the f_T have range contained in $[0, 1]$, so Theorem 2 can be applied with $b = 1$ to yield the bound

$$\frac{K}{\sqrt{m}} \left(14K^{1-1/p} + \frac{1}{2} \sqrt{\ln(16mK^{2-2/p})} \right) + \sqrt{\frac{\ln(1/\delta)}{2m}}$$

on the estimation error. The best bound is obtained when $p = 1$, and the order in K matches that of the bound for K -means clustering described earlier.

The method for $p = 1$ is similar to the sparse-coding method proposed by Olshausen and Field [14], with the difference that the term $\|y\|_1$ is used as a penalty term instead of the hard constraint $\|y\|_1 \leq 1$. The method of Olshausen and Field [14] approximates with a compromise of geometric proximity and sparsity and our result asserts that the observed value of this compromise generalizes to unseen data if enough data have been observed.

3 Proofs

We first introduce some notation, conventions and auxiliary results. Then we set about to prove Theorems 1 and 2.

3.1 Notation, definitions and auxiliary results

Throughout H denotes a Hilbert space. The term *norm* and the notation $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$ always refer to the Euclidean norm and inner product on \mathbb{R}^K or on H . Other norms are characterized by subscripts. If H_1 and H_2 are any Hilbert spaces $\mathcal{L}(H_1, H_2)$ denotes the vector space of bounded linear transformations from H_1 to H_2 . If $H_1 = H_2$ we just write $\mathcal{L}(H_1) = \mathcal{L}(H_1, H_1)$. With $\mathcal{U}(H_1, H_2)$ we denote the set of isometries in $\mathcal{L}(H_1, H_2)$, that is maps U satisfying $\|Ux\|_{H_2} = \|x\|_{H_1}$ for all $x \in H_1$.

We use $\mathcal{L}_2(H)$ for the set of Hilbert-Schmidt operators on H , which becomes itself a Hilbert space with the inner product $\langle T, S \rangle_2 = \text{tr}(T^*S)$ and the corresponding (Frobenius) norm $\|\cdot\|_2$.

For $x \in H$ the rank-one operator Q_x is defined by $Q_x z = \langle z, x \rangle x$. For any $T \in \mathcal{L}_2(H)$ the identity

$$\langle T^*T, Q_x \rangle_2 = \|Tx\|^2$$

is easily verified.

Suppose that $Y \subseteq \mathbb{R}^K$ spans \mathbb{R}^K . It is easily verified that the quantity

$$\|T\|_Y = \sup_{y \in Y} \|Ty\|$$

defines a norm on $\mathcal{L}(\mathbb{R}^K, H)$.

We use the following well known result on covering numbers (see, for example, Proposition 5 in [5]).

Proposition 1. *Let B be a ball of radius r in an N -dimensional Banach space and $\epsilon > 0$. There exists a subset $B_\epsilon \subset B$ such that $|B_\epsilon| \leq (4r/\epsilon)^N$ and $\forall z \in B, \exists z' \in B_\epsilon$ with $d(z, z') \leq \epsilon$, where d is the metric of the Banach space.*

The following concentration inequality, known as the bounded difference inequality [13], goes back to the work of Hoeffding [6].

Theorem 3. *Let μ_i be a probability measure on a space \mathcal{X}_i , for $i = 1, \dots, m$. Let $\mathcal{X} = \prod_{i=1}^m \mathcal{X}_i$ and $\mu = \otimes_{i=1}^m \mu_i$ be the product space and product measure respectively. Suppose the function $\Psi : \mathcal{X} \rightarrow \mathbb{R}$ satisfies*

$$|\Psi(\mathbf{x}) - \Psi(\mathbf{x}')| \leq c_i$$

whenever \mathbf{x} and $\mathbf{x}' \in \mathcal{X}$ differ only in the i -th coordinate, where c_1, \dots, c_m are some positive parameters. Then

$$\Pr_{\mathbf{x} \sim \mu} \{\Psi(\mathbf{x}) - \mathbb{E}_{\mathbf{x}' \sim \mu} \Psi(\mathbf{x}') \geq t\} \leq \exp\left(\frac{-2t^2}{\sum_{i=1}^m c_i^2}\right).$$

Throughout σ_i will denote a sequence of mutually independent random variables, uniformly distributed on $\{-1, 1\}$ and γ_i, γ_{ij} will be (multiple indexed) sequences of mutually independent Gaussian random variables, with zero mean and unit standard deviation.

If \mathcal{F} is a class of real-valued functions on a space \mathcal{X} and μ a probability measure on \mathcal{X} then for $m \in \mathbb{N}$ the Rademacher and Gaussian complexities of \mathcal{F} w.r.t. μ are defined ([9],[2]) as

$$\mathcal{R}_m(\mathcal{F}, \mu) = \frac{2}{m} \mathbb{E}_{\mathbf{x} \sim \mu^m} \mathbb{E}_{\sigma} \sup_{f \in \mathcal{F}} \sum_{i=1}^m \sigma_i f(x_i),$$

$$\Gamma_m(\mathcal{F}, \mu) = \frac{2}{m} \mathbb{E}_{\mathbf{x} \sim \mu^m} \mathbb{E}_{\gamma} \sup_{f \in \mathcal{F}} \sum_{i=1}^m \gamma_i f(x_i)$$

respectively.

Appropriately scaled Gaussian complexities can be substituted for Rademacher complexities, by virtue of the next Lemma. For a proof see, for example, [9, p. 97].

Lemma 1. *For $Y \subseteq \mathbb{R}^k$ we have $\mathcal{R}(Y) \leq \sqrt{\pi/2} \Gamma(Y)$.*

The next result is known as Slepian's lemma ([17], [9]).

Theorem 4. *Let Ω and Ξ be mean zero, separable Gaussian processes indexed by a common set \mathcal{S} , such that*

$$\mathbb{E}(\Omega_{s_1} - \Omega_{s_2})^2 \leq \mathbb{E}(\Xi_{s_1} - \Xi_{s_2})^2 \text{ for all } s_1, s_2 \in \mathcal{S}.$$

Then

$$\mathbb{E} \sup_{s \in \mathcal{S}} \Omega_s \leq \mathbb{E} \sup_{s \in \mathcal{S}} \Xi_s.$$

The following result, which generalizes Theorem 8 in [2], plays a central role in our proof.

Theorem 5. *Let $\{\mathcal{F}_n : 1 \leq n \leq N\}$ be a finite collection of $[0, b]$ -valued function classes on a space \mathcal{X} , and μ a probability measure on \mathcal{X} . Then $\forall \delta \in (0, 1)$ we have with probability at least $1 - \delta$ that*

$$\max_{n \leq N} \sup_{f \in \mathcal{F}_n} \left[\mathbb{E}_{x \sim \mu} f(x) - \frac{1}{m} \sum_{i=1}^m f(x_i) \right] \leq \max_{n \leq N} \mathcal{R}_m(\mathcal{F}_n, \mu) + b \sqrt{\frac{\ln N + \ln(1/\delta)}{2m}}.$$

Proof. Denote with Ψ_n the function on \mathcal{X}^m defined by

$$\Psi_n(\mathbf{x}) = \sup_{f \in \mathcal{F}_n} \left[\mathbb{E}_{x \sim \mu} f(x) - \frac{1}{m} \sum_{i=1}^m f(x_i) \right], \quad \mathbf{x} \in \mathcal{X}^m.$$

By standard symmetrization (see, for example, [18]) we have $\mathbb{E}_{\mathbf{x} \sim \mu^m} \Psi_n(\mathbf{x}) \leq \mathcal{R}_m(\mathcal{F}_n, \mu) \leq \max_{n \leq N} \mathcal{R}_m(\mathcal{F}_n, \mu)$. Modifying one of the x_i can change the value

of any $\Psi_n(\mathbf{x})$ by at most b/m , so that by a union bound and the bounded difference inequality (Theorem 3)

$$\Pr \left\{ \max_{n \leq N} \Psi_n > \max_{n \leq N} \mathcal{R}_m(\mathcal{F}_n, \mu) + t \right\} \leq \sum_n \Pr \{ \Psi_n > \mathbb{E} \Psi_n + t \} \leq N e^{-2m(t/b)^2}.$$

Solving $\delta = N e^{-2m(t/b)^2}$ for t gives the result. \square

Notice that replacing the functions $f \in \mathcal{F}_n$ by $b - f$ does not affect the Rademacher complexities, so the above result can be used in a two-sided way.

The following lemma was used in Section 2.3.

Lemma 2. *Suppose $\|x\| \leq 1$, $\|c_k\| = 1$, $\langle c_k, c_l \rangle \geq 0$, $y \in \mathbb{R}^K$, $y_i \geq 0$. If y minimizes*

$$h(y) = \left\| x - \sum_{k=1}^K y_k c_k \right\|^2,$$

then $\|y\| \leq 1$.

Proof. Assume that y is a minimizer of h and $\|y\| > 1$. Then

$$\left\| \sum_{k=1}^K y_k c_k \right\|^2 = \|y\|^2 + \sum_{k \neq l} y_k y_l \langle c_k, c_l \rangle > 1.$$

Let the real-valued function f be defined by $f(t) = h(ty)$. Then

$$\begin{aligned} f'(1) &= 2 \left(\left\| \sum_{k=1}^K y_k c_k \right\|^2 - \left\langle x, \sum_{k=1}^K y_k c_k \right\rangle \right) \\ &\geq 2 \left(\left\| \sum_{k=1}^K y_k c_k \right\|^2 - \left\| \sum_{k=1}^K y_k c_k \right\| \right) \\ &= 2 \left(\left\| \sum_{k=1}^K y_k c_k \right\| - 1 \right) \left\| \sum_{k=1}^K y_k c_k \right\| \\ &> 0. \end{aligned}$$

So f cannot have a minimum at 1, whence y cannot be a minimizer of h . \square

3.2 Proof of the main results

We now fix a spanning codebook $Y \subseteq \mathbb{R}^K$ and recall that, for $T \in \mathcal{L}(\mathbb{R}^K, H)$, we had introduced the notation

$$f_T(x) = \inf_{y \in Y} \|x - Ty\|^2, x \in H.$$

Our principal object of study is the function class

$$\mathcal{F} = \{f_T : T \in \mathcal{T}\},$$

where $\mathcal{T} \subset \mathcal{L}(\mathbb{R}^K, H)$ is some fixed set of candidate implementations of our coding scheme. We first address the rather general Theorem 1 which can be treated in parallel to the case of K -means clustering. We begin with a technical lemma.

Lemma 3. *Suppose that*

1. $(e_k : 1 \leq k \leq K)$ is an orthonormal basis of \mathbb{R}^K ;
2. \mathcal{T} is the class of linear operators $T : \mathbb{R}^K \rightarrow H$ with $\|Te_k\| \leq c$;
3. $(x_i : 1 \leq i \leq m)$ is a sequence $x_i \in H$, $\|x_i\| \leq 1$;
4. $(\gamma_{ik} : 1 \leq i \leq m, 1 \leq k \leq K)$ and $(\gamma_{ikl} : 1 \leq i \leq m, 1 \leq k, l \leq K)$ are orthogaussian sequences.

Then the following three inequalities hold

$$\begin{aligned} \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{i=1}^m \sum_{k=1}^K \gamma_{ik} \langle x_i, Te_k \rangle &\leq cK\sqrt{m} \\ \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{i=1}^m \sum_{k=1}^K \gamma_{ik} \|Te_k\|^2 &\leq c^2 K\sqrt{m} \\ \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{i=1}^m \sum_{k,l=1}^K \gamma_{ikl} \langle Te_k, Te_l \rangle &\leq c^2 K^2\sqrt{m}. \end{aligned}$$

Proof. Using Cauchy-Schwarz' and Jensen's inequalities and the orthogaussian properties of the γ_{ik} , we get

$$\mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{k=1}^K \sum_{i=1}^m \gamma_{ik} \langle x_i, Te_k \rangle \leq c \mathbb{E}_\gamma \sum_{k=1}^K \left\| \sum_{i=1}^m \gamma_{ik} x_i \right\| \leq cK\sqrt{m}$$

which is the first inequality. Similarly we obtain

$$\begin{aligned} \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{k=1}^K \sum_{i=1}^m \gamma_{ik} \|Te_k\|^2 &\leq c^2 \mathbb{E}_\gamma \sum_{k=1}^K \left| \sum_{i=1}^m \gamma_{ik} \right| \leq c^2 K\sqrt{m} \\ \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{k,l=1}^K \sum_{i=1}^m \gamma_{ikl} \langle Te_k, Te_l \rangle &\leq c^2 \mathbb{E}_\gamma \sum_{k,l=1}^K \left| \sum_{i=1}^m \gamma_{ikl} \right| \leq c^2 K^2\sqrt{m}. \end{aligned}$$

□

Proposition 2. *Suppose that the probability measure μ is supported on the unit ball of H , that $\{e_k : 1 \leq k \leq K\}$ is an orthonormal basis of \mathbb{R}^K and that \mathcal{T} is*

a class of linear operators $T : \mathbb{R}^K \rightarrow H$ with $\|Te_k\| \leq c$ for $1 \leq k \leq K$, with $c \geq 1$. Let Y be a nonempty closed subset of the unit ball in \mathbb{R}^K and

$$\mathcal{F}_Y = \left\{ x \in H \mapsto \min_{y \in Y} \|x - Ty\|^2 : T \in \mathcal{T} \right\}.$$

Then

$$\mathcal{R}(\mathcal{F}_Y, \mu) \leq 6c^2 K^2 \sqrt{\frac{\pi}{m}}.$$

and if $Y = \{e_k : 1 \leq k \leq K\}$ then the bound improves to

$$\mathcal{R}(\mathcal{F}_Y, \mu) \leq c^2 K \sqrt{\frac{18\pi}{m}}.$$

Proof. By Lemma 1 it suffices to bound the corresponding Gaussian averages, which we shall do using Slepian's Lemma (Theorem 4). First fix a sample \mathbf{x} and define Gaussian processes Ω and Ξ indexed by \mathcal{T}

$$\begin{aligned} \Omega_T &= \sum_i \gamma_i \min_y \|x_i - Ty\|^2 \quad \text{and} \\ \Xi_T &= \sqrt{8} \sum_{ik} \gamma_{ik} \langle x_i, Te_k \rangle + \sqrt{2} \sum_{ilk} \gamma_{ilk} \langle Te_l, Te_k \rangle. \end{aligned}$$

Suppose $T_1, T_2 \in \mathcal{T}$. For any $x \in H$ we have, using $(a+b)^2 \leq 2a^2 + 2b^2$ and Cauchy-Schwarz

$$\begin{aligned} & \left(\min_{y \in Y} \|x - T_1 y\|^2 - \min_y \|x - T_2 y\|^2 \right)^2 \\ & \leq \left(\max_{y \in Y} \|x - T_1 y\|^2 - \|x - T_2 y\|^2 \right)^2 \\ & \leq 8 \max_{y \in Y} \left(\sum_k y_k \langle x, (T_1 - T_2) e_k \rangle \right)^2 + 2 \max_{y \in Y} \left(\sum_{kl} y_k y_l \langle e_k, (T_1^* T_1 - T_2^* T_2) e_l \rangle \right)^2 \\ & \leq 8 \sum_k (\langle x, T_1 e_k \rangle - \langle x, T_2 e_k \rangle)^2 + 2 \sum_{kl} (\langle T_1 e_k, T_1 e_l \rangle - \langle T_2 e_k, T_2 e_l \rangle)^2. \end{aligned}$$

We therefore have

$$\begin{aligned} \mathbb{E}(\Omega_{T_1} - \Omega_{T_2})^2 &= \sum_i \left(\min_y \|x_i - T_1 y\|^2 - \min_y \|x_i - T_2 y\|^2 \right)^2 \\ &\leq 8 \sum_{ik} (\langle x_i, T_1 e_k \rangle - \langle x_i, T_2 e_k \rangle)^2 + 2 \sum_{ikl} (\langle T_1 e_k, T_1 e_l \rangle - \langle T_2 e_k, T_2 e_l \rangle)^2 \\ &= \mathbb{E}(\Xi_{T_1} - \Xi_{T_2})^2. \end{aligned}$$

So, by Slepian's Lemma and the first and last inequalities in Lemma 3

$$\begin{aligned}\mathbb{E} \sup_{T \in \mathcal{T}} \Omega_T &\leq \mathbb{E} \sup_{T \in \mathcal{T}} \Xi_T \\ &\leq \sqrt{8} \mathbb{E} \sup_{T \in \mathcal{T}} \sum_{ik} \gamma_{ik} \langle x_i, T e_k \rangle + \sqrt{2} \mathbb{E} \sup_{T \in \mathcal{T}} \sum_{ilk} \gamma_{ilk} \langle T e_l, T e_k \rangle \\ &\leq cK \sqrt{8m} + c^2 K^2 \sqrt{2m}.\end{aligned}$$

Multiply by $\sqrt{2\pi}/m$ to get a bound on the Rademacher complexity of

$$\mathcal{R}(\mathcal{F}_Y, \mu) \leq 4cK \sqrt{\frac{\pi}{m}} + 2c^2 K^2 \sqrt{\frac{\pi}{m}} \leq 6c^2 K^2 \sqrt{\frac{\pi}{m}}.$$

To obtain the second conclusion we improve the bound on the Gaussian average. With Ω_T as above we set

$$\Xi_T = \sum_{i=1}^m \sum_{k=1}^K \gamma_{ik} \|x_i - T e_k\|^2.$$

Now we have for $T_1, T_2 \in \mathcal{T}$ that

$$\begin{aligned}\mathbb{E} (\Omega_{T_1} - \Omega_{T_2})^2 &= \sum_{i=1}^m \left(\min_{k=1}^K \|x_i - T_1 e_k\|^2 - \min_{k=1}^K \|x_i - T_2 e_k\|^2 \right)^2 \\ &\leq \sum_{i=1}^m \max_{k=1}^K \left(\|x_i - T_1 e_k\|^2 - \|x_i - T_2 e_k\|^2 \right)^2 \\ &\leq \sum_{i=1}^m \sum_{k=1}^K \left(\|x_i - T_1 e_k\|^2 - \|x_i - T_2 e_k\|^2 \right)^2 \\ &= \mathbb{E} (\Xi_{T_1} - \Xi_{T_2})^2.\end{aligned}$$

Again with Slepian's Lemma and the triangle inequality

$$\begin{aligned}\mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \Omega_T &\leq \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \Xi_T = \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{i=1}^m \sum_{k=1}^K \gamma_{ik} \|x_i - T e_k\|^2 \\ &\leq 2 \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{i=1}^m \sum_{k=1}^K \gamma_{ik} \langle x_i, T e_k \rangle + \mathbb{E}_\gamma \sup_{T \in \mathcal{T}} \sum_{i=1}^m \sum_{k=1}^K \gamma_{ik} \|T e_k\|^2 \\ &\leq 3c^2 K \sqrt{m},\end{aligned}$$

where the last inequality follows from the first two inequalities in Lemma 3. Multiply by $\sqrt{2\pi}/m$ as above \square

Theorem 1 follows from observing that the functions in \mathcal{F} map to $[0, 4c^2]$ and combining the above bound on the Rademacher complexity with Theorem 5 with $N = 1$ and $b = 4$.

The second conclusion of the proposition yields a bound for K -means clustering, corresponding to the choices $Y = \{e_1, \dots, e_K\}$ and $\mathcal{T} = \{T : \|Te_k\| \leq 1, 1 \leq k \leq K\}$. As already noted in Section 2.2 the vectors Te_k define the cluster centers. With Theorem 5 we obtain

Theorem 6. *For every $\delta > 0$ with probability greater $1 - \delta$ in the sample $\mathbf{x} \sim \mu^m$ we have for all $T \in \mathcal{T}$*

$$\mathbb{E}_{x \sim \mu} \min_{k=1}^K \|x - Te_k\|^2 \leq \frac{1}{m} \sum_{i=1}^m \min_{k=1}^K \|x_i - Te_k\|^2 + K \sqrt{\frac{18\pi}{m}} + \sqrt{\frac{8 \ln(1/\delta)}{m}}.$$

To prove Theorem 2 a more subtle approach is necessary. The idea is the following: every implementing map $T \in \mathcal{T}$ can be factored as $T = US$, where S is a $K \times K$ matrix, $S \in \mathcal{L}(\mathbb{R}^K)$, and U is an isometry, $U \in \mathcal{U}(\mathbb{R}^K, H)$. Suitably bounded $K \times K$ matrices form a compact, finite dimensional set, the complexity of which can be controlled using covering numbers, while the complexity arising from the set of isometries can be controlled with Rademacher and Gaussian averages. Theorem 5 then combines these complexity estimates.

For fixed $S \in \mathcal{L}(\mathbb{R}^K)$ we denote

$$\mathcal{G}_S = \{f_{US} : U \in \mathcal{U}(\mathbb{R}^K, H)\}.$$

Recall the notation $\|\mathcal{T}\|_Y = \sup_{T \in \mathcal{T}} \|T\|_Y = \sup_{T \in \mathcal{T}} \sup_{y \in Y} \|Ty\|$. With \mathcal{S} we denote the set of $K \times K$ matrices

$$\mathcal{S} = \{S \in \mathcal{L}(\mathbb{R}^K) : \|S\|_Y \leq \|\mathcal{T}\|_Y\}.$$

Lemma 4. *Assume $\|\mathcal{T}\|_Y \geq 1$, that the functions in \mathcal{F} , when restricted to the unit ball of H , have range contained in $[0, b]$, and that the measure μ is supported on the unit ball of H . Then with probability at least $1 - \delta$ we have for all $T \in \mathcal{T}$ that*

$$\begin{aligned} & \mathbb{E}_{x \sim \mu} f_T(x) - \frac{1}{m} \sum_{i=1}^m f_T(x_i) \\ & \leq \sup_{S \in \mathcal{S}} \mathcal{R}_m(\mathcal{G}_S, \mu) + \frac{bK}{2} \sqrt{\frac{\ln(16m \|\mathcal{T}\|_Y^2)}{m}} + \frac{8\|\mathcal{T}\|_Y}{\sqrt{m}} + b \sqrt{\frac{\ln(1/\delta)}{2m}}. \end{aligned}$$

Proof. Fix $\epsilon > 0$. The set \mathcal{S} is the ball of radius $\|\mathcal{T}\|_Y$ in the K^2 -dimensional Banach space $(\mathcal{L}(\mathbb{R}^K), \|\cdot\|_Y)$ so by Proposition 1 we can find a subset $\mathcal{S}_\epsilon \subset \mathcal{S}$, of cardinality $|\mathcal{S}_\epsilon| \leq (4\|\mathcal{T}\|_Y/\epsilon)^{K^2}$ such that every member of \mathcal{S} can be approximated by a member of \mathcal{S}_ϵ up to distance ϵ in the norm $\|\cdot\|_Y$.

We claim that for all $T \in \mathcal{T}$ there exist $U \in \mathcal{U}(\mathbb{R}^K, H)$ and $S_\epsilon \in \mathcal{S}_\epsilon$ such that

$$|f_T(x) - f_{US_\epsilon}(x)| < 4\|\mathcal{T}\|_Y \epsilon,$$

for all x in the unit ball of H . To see this write $T = US$ with $U \in \mathcal{U}(\mathbb{R}^K, H)$ and $S \in \mathcal{L}(\mathbb{R}^K)$. Then, since U is an isometry, we have

$$\|S\|_Y = \sup_{y \in Y} \|Sy\| = \sup_{y \in Y} \|Ty\| = \|T\|_Y \leq \|\mathcal{T}\|_Y$$

so that $S \in \mathcal{S}$. We can therefore choose $S_\epsilon \in \mathcal{S}_\epsilon$ such that $\|S_\epsilon - S\|_Y < \epsilon$. Then for $x \in H$, with $\|x\| \leq 1$, we have

$$\begin{aligned} |f_T(x) - f_{US_\epsilon}(x)| &= \left| \inf_{y \in Y} (\|x - USy\|^2) - \inf_{y \in Y} (\|x - US_\epsilon y\|^2) \right| \\ &\leq \sup_{y \in Y} \left| (\|x - USy\|^2 - \|x - US_\epsilon y\|^2) \right| \\ &= \sup_{y \in Y} |(US_\epsilon y - USy, 2x - (USy + US_\epsilon y))| \\ &\leq (2 + 2\|\mathcal{T}\|_Y) \sup_{y \in Y} \|(S_\epsilon - S)y\| \leq 4\|\mathcal{T}\|_Y \epsilon. \end{aligned}$$

Apply Theorem 5 to the finite collection of function classes $\{\mathcal{G}_S : S \in \mathcal{S}_\epsilon\}$ to see that with probability at least $1 - \delta$

$$\begin{aligned} &\sup_{T \in \mathcal{T}} \mathbb{E}_{x \sim \mu} f_T(x) - \frac{1}{m} \sum_{i=1}^m f_T(x_i) \\ &\leq \max_{S \in \mathcal{S}_\epsilon} \sup_{U \in \mathcal{U}(\mathbb{R}^K, H)} \mathbb{E}_{x \sim \mu} f_{US}(x) - \frac{1}{m} \sum_{i=1}^m f_{US}(x_i) + 8\|\mathcal{T}\|_Y \epsilon \\ &\leq \max_{S \in \mathcal{S}_\epsilon} \mathcal{R}_m(\mathcal{G}_S, \mu) + b \sqrt{\frac{\ln |\mathcal{S}_\epsilon| + \ln(1/\delta)}{2m}} + 8\|\mathcal{T}\|_Y \epsilon \\ &\leq \sup_{S \in \mathcal{S}} \mathcal{R}_m(\mathcal{G}_S, \mu) + \frac{bK}{2} \sqrt{\frac{\ln(16m\|\mathcal{T}\|_Y^2)}{m}} + \frac{8\|\mathcal{T}\|_Y}{\sqrt{m}} + b \sqrt{\frac{\ln(1/\delta)}{2m}}, \end{aligned}$$

where the last line follows from the known bound on $|\mathcal{S}_\epsilon|$, subadditivity of the square root and the choice $\epsilon = 1/\sqrt{m}$. \square

Remark 1. If H is finite dimensional the above result may be improved to

$$\mathbb{E}f_T - \hat{\mathbb{E}}f_T \leq \frac{b}{2} \sqrt{\frac{dK \ln(16m\|\mathcal{T}\|_Y^2)}{m}} + \frac{8\|\mathcal{T}\|_Y}{\sqrt{m}} + b \sqrt{\frac{\ln(1/\delta)}{2m}}. \quad (2)$$

To see this, follow the same lines as in Lemma 4 to note that

$$\sup_{T \in \mathcal{T}} \mathbb{E}f_T - \hat{\mathbb{E}}f_T \leq \max_{T \in \mathcal{T}_\epsilon} \mathbb{E}f_T - \hat{\mathbb{E}}f_T + 8\|\mathcal{T}\|_Y \epsilon,$$

where \mathcal{T}_ϵ is a subset of \mathcal{T} such that every member of \mathcal{T} can be approximated by a member of \mathcal{T}_ϵ up to distance ϵ in the norm $\|\cdot\|_Y$.

By Proposition 1, $|\mathcal{T}_\epsilon| \leq (4\|\mathcal{T}\|_Y/\epsilon)^{dK}$. Inequality (2) now follows from Theorem 5 with $N = |\mathcal{T}_\epsilon|$ and $\epsilon = 1/\sqrt{m}$.

To complete the proof of Theorem 2 we now fix some $S \in \mathcal{S}$ and focus on the corresponding function class \mathcal{G}_S .

Lemma 5. *For any $S \in \mathcal{L}(\mathbb{R}^K)$ we have*

$$\mathcal{R}(\mathcal{G}_S, \mu) \leq 2\sqrt{2\pi} \|S\|_Y \frac{K}{\sqrt{m}}.$$

Proof. Let $\|x_i\| \leq 1$ and define Gaussian processes Ω_U and Ξ_U indexed by $\mathcal{U}(\mathbb{R}^K, H)$

$$\begin{aligned} \Omega_U &= \sum_{i=1}^m \gamma_i \inf_{y \in Y} \|x_i - USy\|^2 \\ \Xi_U &= 2 \|S\|_Y \sum_{k=1}^K \sum_{i=1}^m \gamma_{ik} \langle x_i, Ue_k \rangle, \end{aligned}$$

where the e_k are the canonical basis of \mathbb{R}^K . For $U_1, U_2 \in \mathcal{U}(\mathbb{R}^K, H)$ we have

$$\begin{aligned} \mathbb{E}(\Omega_{U_1} - \Omega_{U_2})^2 &\leq \sum_{i=1}^m \left(\sup_{y \in Y} \|x_i - U_1Sy\|^2 - \|x_i - U_2S\|^2 \right)^2 \\ &\leq \sum_{i=1}^m \sup_{y \in Y} 4 \langle x_i, (U_2 - U_1)Sy \rangle^2 \\ &\leq 4 \sum_{i=1}^m \sup_{y \in Y} \|U_2^*x_i - U_1^*x_i\|^2 \|Sy\|^2 \\ &= 4 \|S\|_Y^2 \sum_{i=1}^m \sum_{k=1}^K (\langle x_i, U_1e_k \rangle - \langle x_i, U_2e_k \rangle)^2 \\ &= \mathbb{E}(\Xi_{U_1} - \Xi_{U_2})^2. \end{aligned}$$

It follows from Lemma 1 and Slepian's lemma (Theorem 4) that

$$\mathcal{R}_m(\mathcal{G}_S, \mu) \leq \mathbb{E}_{\mathbf{x} \sim \mu^m} \frac{2}{m} \sqrt{\frac{\pi}{2}} \mathbb{E}_\gamma \sup_U \Xi_U,$$

so the result follows from the following inequalities, using Cauchy-Schwarz' and Jensen's inequality, the orthonormality of the γ_{ik} and the fact that $\|x_i\| \leq 1$ on the support of μ .

$$\begin{aligned} \mathbb{E}_\gamma \sup_U \Xi_U &= 2 \|S\|_Y \mathbb{E} \sup_U \left\langle \sum_{k=1}^K \left\langle \sum_{i=1}^m \gamma_{ik} x_i, Ue_k \right\rangle \right\rangle \\ &\leq 2 \|S\|_Y \sum_{k=1}^K \mathbb{E} \left\| \sum_{i=1}^m \gamma_{ik} x_i \right\| \\ &\leq 2 \|S\|_Y K \sqrt{m}. \end{aligned}$$

□

Substitution of the last result in Lemma 4 and noting that, for $K \geq 1$, $2\sqrt{2\pi}K + 8 \leq 14K$, gives Theorem 2.

Observe that when the set \mathcal{S} contains only the identity matrix, the function class $\mathcal{G}_{\mathcal{S}}$ is the class of reconstruction errors of PCA. In this case, the result can be improved as shown by the next lemma.

Lemma 6. $\mathcal{R}(\mathcal{D}, \mu) \leq 2\sqrt{K/m}$.

Proof. Recall, for every $z \in H$, that the outer product operator Q_z is defined by $Q_z x = \langle x, z \rangle z$. With $\langle \cdot, \cdot \rangle_2$ and $\|\cdot\|_2$ denoting the Hilbert-Schmidt inner product and norm respectively we have for $\|x_i\| \leq 1$

$$\begin{aligned} \mathbb{E}_{\sigma} \sup_{f \in \mathcal{D}} \sum_{i=1}^m \sigma_i f(x_i) &= \mathbb{E}_{\sigma} \sup_{U \in \mathcal{U}} \sum_{i=1}^m \sigma_i \left(\|x_i\|^2 - \|UU^* x_i\|^2 \right) \\ &= \mathbb{E}_{\sigma} \sup_{U \in \mathcal{U}} \left\langle \sum_{i=1}^m \sigma_i Q_{x_i}, UU^* \right\rangle_2 \\ &\leq \mathbb{E}_{\sigma} \left\| \sum_{i=1}^m \sigma_i Q_{x_i} \right\|_2 \sup_{U \in \mathcal{U}} \|UU^*\|_2 \\ &\leq \sqrt{mK}, \end{aligned}$$

since the Hilbert-Schmidt norm of a K -dimensional projection is \sqrt{K} . The result follows upon multiplication with $2/m$ and taking the expectation in μ^m . \square

An application of Theorem 5 with $N = 1$ and $b = 1$ also give a generalization bound for PCA of order $\sqrt{K/m}$.

4 Concluding remarks

We have analyzed a general method to encode random vectors in a Hilbert space H . The method searches for an operator $T : \mathbb{R}^K \rightarrow H$ which minimizes, within some prescribed class \mathcal{T} , the empirical average of the reconstruction error, which is defined as the minimum distance between a given point in H and an image of the operator T acting on a prescribed codebook Y .

We have presented two approaches to upper bound the estimation error of the method in terms of the parameter K , the sample size m and the properties of the sets \mathcal{T} and Y . The first approach is based on a direct bound for the Rademacher average of the loss class induced by the reconstruction error. The bound matches the best known bound for K -means clustering in a Hilbert space [4] but also applies to other interesting coding techniques such as sparse coding and non-negative matrix factorization. The second approach uses a decomposition of the function class as a union of function classes parameterized by K -dimensional isometries. The main idea is to approximate the union with a finite union via covering numbers and then bound the complexity of each class under the union with Rademacher averages. This second result is more complicated than the first

one, however it provides in certain cases a better dependency of the bound on the parameter K at the expense of an additional logarithmic factor in m .

We conclude with some open problems and possible extensions which are suggested by this study. Firstly, it would be valuable to investigate the possibility of removing the logarithmic term in m in the bound of Theorem 2. Secondly, it would be important to elucidate whether the dependency in K in the same bound is optimal. The latter problem is also mentioned in [4] in the case of K -means clustering. Finally, it would be interesting to study possible improvements of our results in the case that additional assumptions on the probability measure μ are introduced. For example, in the case of K -means clustering in a finite dimensional Hilbert space [1] shows that for certain classes of probability measures the rate of convergence can be improved to $O(\log(m)/m)$ and it may be possible to obtain similar improvements in our general framework.

Acknowledgments

This work was supported by EPSRC Grants GR/T18707/01 and EP/D071542/1.

References

1. A. Antos, L. Györfi, A. Györfi. Individual convergence rates in empirical vector quantizer design. *IEEE Transactions on Information Theory*, 51(11):4013–4022, 2005.
2. P. L. Bartlett and S. Mendelson. Rademacher and Gaussian Complexities: Risk Bounds and Structural Results. *Journal of Machine Learning Research*, 3: 463–482, 2002.
3. P. Bartlett, T. Linder, G. Lugosi. The minimax distortion redundancy in empirical quantizer design. *IEEE Transactions on Information Theory*, 44: 1802–1813, 1998.
4. G. Biau, L. Devroye, G. Lugosi. On the performance of clustering in Hilbert spaces. *IEEE Transactions on Information Theory*, 54:781–790, 2008.
5. F. Cucker and S. Smale. On the mathematical foundations of learning, *Bulletin of the American Mathematical Society*, 39 (1):1–49, 2001.
6. W. Hoeffding. Probability inequalities for sums of bounded random variables, *Journal of the American Statistical Association*, 58:13–30, 1963.
7. P. O. Hoyer. Non-negative matrix factorization with sparseness constraints. *Journal of Machine Learning Research*, 5:1457–1469, 2004.
8. V. Koltchinskii and D. Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers, *The Annals of Statistics*, 30(1): 1–50, 2002.
9. M. Ledoux, M. Talagrand. *Probability in Banach Spaces*, Springer, 1991.
10. D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature* 401, 788–791, 1999.
11. S. Z. Li, X. Hou, H. Zhang, and Q. Cheng. Learning spatially localized parts-based representations. *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Vol. I, pages 207–212, Hawaii, USA, 2001.
12. A. Maurer and M. Pontil. Generalization bounds for K -dimensional coding schemes in Hilbert spaces. Proceedings of the 19th international conference on Algorithmic Learning Theory, pages 91, 2008,

13. C. McDiarmid. *Concentration*, in *Probabilistic Methods of Algorithmic Discrete Mathematics*, p195-248, Springer, Berlin, 1998.
14. B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609, 1996.
15. J. Shawe-Taylor, C. K. I. Williams, N. Cristianini, J. S. Kandola. On the eigen-spectrum of the Gram matrix and the generalization error of kernel-PCA. *IEEE Transactions on Information Theory* 51(7): 2510–2522, 2005.
16. O. Wigelius, A. Ambroladze, J. Shawe-Taylor. Statistical analysis of clustering with applications. Preprint, 2007.
17. D. Slepian. The one-sided barrier problem for Gaussian noise. *Bell System Tech. J.*, 41: 463–501, 1962.
18. A.W. van der Vaart and J.A. Wallner. *Weak Convergence and Empirical Processes*, Springer Verlag, 1996.
19. L. Zwald, L., O. Bousquet, and G. Blanchart. Statistical properties of kernel principal component analysis. *Machine Learning* 66(2-3): 259–294, 2006.