# The Sample Complexity of Dictionary Learning 

Daniel Vainsencher<br>danielv@tx.technion.ac.il<br>Department of Electrical Engineering<br>Technion, Israel Institute of Technology<br>Haifa 32000, Israel

Shie Mannor<br>shie@ee.technion.ac.il<br>Department of Electrical Engineering<br>Technion, Israel Institute of Technology<br>Haifa 32000, Israel

Alfred M. Bruckstein<br>freddy@cs.technion.ac.il<br>Department of Computer Science<br>Technion, Israel Institute of Technology<br>Haifa 32000, Israel

November 25, 2010


#### Abstract

A large set of signals can sometimes be described sparsely using a dictionary, that is, every element can be represented as a linear combination of few elements from the dictionary. Algorithms for various signal processing applications, including classification, denoising and signal separation, learn a dictionary from a set of signals to be represented. Can we expect that the representation found by such a dictionary for a previously unseen example from the same source will have $L_{2}$ error of the same magnitude as those for the given examples? We assume signals are generated from a fixed distribution, and study this questions from a statistical learning theory perspective.

We develop generalization bounds on the quality of the learned dictionary for two types of constraints on the coefficient selection, as measured by the expected $L_{2}$ error in representation when the dictionary is used. For the case of $l_{1}$ regularized coefficient selection we provide a generalization bound of the order of $O(\sqrt{n p \log (m \lambda) / m})$, where $n$ is the dimension, $p$ is the number of elements in the dictionary, $\lambda$ is a bound on the $l_{1}$ norm of the coefficient vector and $m$ is the number of samples, which complements existing results. For the case of representing a new signal as a combination of at most $k$ dictionary elements, we provide a bound of the order $O(\sqrt{n p \log (m k) / m})$ under an assumption on the level of orthogonality of the dictionary (low Babel function). We further show that this assumption holds for most dictionaries in high dimensions in a strong probabilistic sense. Our results further yield fast rates of order $1 / m$ as opposed to $1 / \sqrt{m}$ using localized Rademacher complexity. We provide similar results in a general setting using kernels with weak smoothness requirements.


## 1 Introduction

In processing signals from $\mathcal{X}=\mathbb{R}^{n}$ it is now a common technique to use sparse representations; that is, to approximate each signal $x$ by a "small" linear combination $a$ of elements $d_{i}$ from a dictionary
$D \in \mathcal{X}^{p}$, so that $x \approx D a=\sum_{i=1}^{p} a_{i} d_{i}$. This has various uses detailed in Section 1.1 The smallness of $a$ is often measured using either $\|a\|_{1}$, or the number of non zero elements in $a$, often denoted $\|a\|_{0}$. The approximation error is measured here using a Euclidean norm appropriate to the vector space. We denote the approximation error of $x$ using dictionary $D$ and coefficients from $A$ as

$$
\begin{equation*}
h_{A, D}(x)=\min _{a \in A}\|D a-x\|, \tag{1.1}
\end{equation*}
$$

where $A$ is one of the following sets determining the sparsity required of the representation:

$$
H_{k}=\left\{a:\|a\|_{0} \leq k\right\}
$$

induced a "hard" sparsity constraint, which we also call $k$ sparse representation, while

$$
R_{\lambda}=\left\{a:\|a\|_{1} \leq \lambda\right\}
$$

induces a convex constraint that is a "relaxation" of the previous constraint.
The dictionary learning problem is to find a dictionary $D$ minimizing

$$
\begin{equation*}
E(D)=\mathbb{E}_{x \sim \nu} h_{A, D}(x), \tag{1.2}
\end{equation*}
$$

where $\nu$ is a distribution over signals that is known to us only through samples from it. The problem addressed in this paper is the "generalization" (in the statistical learning sense) of dictionary learning: to what extent does the performance of a dictionary chosen based on a finite set of samples indicate its expected error in (1.2)? This clearly depends on the number of samples and other parameters of the problem such as dictionary size. In particular, an obvious algorithm is to represent each sample using itself, if the dictionary is allowed to be as large as the sample, but the performance on unseen signals is likely to disappoint.

To state our goal more quantitatively, assume that an algorithm finds a dictionary $D$ suited to $k$ sparse representation, in the sense that the average representation error $E_{m}(D)$ on the $m$ examples it is given is low. Our goal is to bound the generalization error $\varepsilon$, which is the additional expected error that might be incurred:

$$
E(D) \leq(1+\eta) E_{m}(D)+\varepsilon,
$$

where $\eta \geq 0$ is sometimes zero, and the bound depends on the number of samples and problem parameters. Since algorithms that find the optimal dictionary for a given set of samples (also known as empirical risk minimization, or ERM, algorithms) are not known for dictionary learning, we prove uniform convergence bounds that apply simultaneously over all admissible dictionaries $D$, thus bounding from above the sample complexity of the dictionary learning problem.

Many analytic and algorithmic methods relying on the properties of finite dimensional Euclidean geometry can be applied in more general settings by applying kernel methods. These consist of treating objects that are not naturally represented in $\mathbb{R}^{n}$ as having their similarity described by an inner product in an abstract feature space that is Euclidean. This allows the application of algorithms depending on the data only through a computation of inner products to such diverse objects as graphs, DNA sequences and text documents, that are not naturally represented using vector spaces (Shawe-Taylor and Cristianini, 2004). Is it possible to extend the usefulness of dictionary learning techniques to this setting? We address sample complexity aspects of this question as well.

### 1.1 Background and related work

Sparse representations are a standard practice in diverse fields such as signal processing, natural language processing, etc. Typically, the dictionary is assumed to be known. The motivation for sparse representations is indicated by the following results, in which we assume the signals come from $\mathcal{X}=\mathbb{R}^{n}$, and the representation coefficients from $A=H_{k}$ where $k<n, p$ and typically $h_{A, D}(x) \ll 1$.

- Compression: If a signal $x$ has an approximate sparse representation in some commonly known dictionary $D$, then by definition, storing or transmitting the sparse representation will not cause large error.
- Representation: If a signal $x$ has an approximate sparse representation in a dictionary $D$ that fulfills certain geometric conditions, then its sparse representation is unique and can be found efficiently (Bruckstein et al., 2009).
- Denoising: If a signal $x$ has a sparse representation in some known dictionary $D$, and $\tilde{x}=$ $x+\nu$, where the random noise $\nu$ is Gaussian, then the sparse representation found for $\tilde{x}$ will likely be very close to $x$ (for example Chen et al., 2001).
- Compressed sensing: Assuming that a signal $x$ has a sparse representation in some known dictionary $D$ that fulfills certain geometric conditions, this representation can be approximately retrieved with high probability from a small number of random linear measurements of $x$. The number of measurements needed depends on the sparsity of $x$ in $D$ (Candes and Tao, 2006).

The implications of these results are significant when a dictionary $D$ is known that sparsely represents simultaneously many signals. In some applications the dictionary is chosen based on prior knowledge, but in many applications the dictionary is learned based on a finite set of examples. To motivate dictionary learning, consider an image representation used for compression or denoising. Different types of images may have different properties (MRI images are not similar to scenery images), so that learning a specific dictionary to each type of images may lead to improved performance. The benefits of dictionary learning have been demonstrated in many applications (Protter and Elad, 2007; Peyré, 2009; Yang et al., 2009).

Two extensively used techniques related to dictionary learning are Principal Component Analysis (PCA) and $k$ means clustering. The former finds a single subspace minimizing the sum of squared representation errors which is very similar to dictionary learning with $A=H_{k}$ and $p=k$. The latter finds a set of locations minimizing the sum of squared distances between each signal and the location closest to it which is very similar to dictionary learning with $A=H_{1}$ where $p$ is the number of locations. Thus we could see dictionary learning as PCA with multiple subspaces, or as clustering where multiple locations are used to represent each signal. The sample complexity of both algorithms are well studied (Bartlett et al., 1998; Biau et al., 2008; Shawe-Taylor et al., 2005; Blanchard et al., 2007).

This paper does not address questions of computational cost, though they are very relevant. Finding optimal coefficients for $k$ sparse representation (that is, minimizing (1.1) with $A=H_{k}$ ) is NP-hard in general (Davis et al., 1997). Dictionary learning as an optimization problem, that
of minimizing (1.2) is less well understood, even for empirical $\nu$ (consisting of a finite number of samples), despite over a decade of work on related algorithms with good empirical results (Olshausen and Fieldt, 1997; Lewicki et al., 1998; Kreutz-Delgado et al., 2003; Aharon et al., 2005; Lee et al., 2007; Krause and Cevher, 2010; Mairal et al., 2010).

The only prior work we are aware of that addresses generalization in dictionary learning, by Maurer and Pontil (2010), addresses the convex representation constraint $A=R_{\lambda}$; we discuss the relation of our work to theirs in Section 2] Another related work studies the identifiability of dictionaries, giving conditions under which a dictionary may be exactly recovered. A recent example giving somewhat similar requirements on the number of samples (though in a different setting, and to obtain a different kind of result) is by Gribonval and Schnass (2009), which also includes a review of identifiability results.

## 2 Results

Except where we state otherwise, we assume signals are generated in the unit sphere $\mathbb{S}^{n-1}$.
A new approach to dictionary learning generalization. Our first main contribution is an approach to generalization bounds in dictionary learning that is complementary to that used by Maurer and Pontil (2010). Assume that the columns of the dictionary $D \in \mathbb{R}^{n \times p}$ are of unit length, and that each signal $x \in \mathbb{S}^{n-1}$ is approximately represented in the form $D a$ where the coefficient vector $a$ is known to fulfill a constraint of form $\|a\|_{1} \leq \lambda$. We quantify the complexity of the associated error function class in terms of $\lambda$, so that standard methods of uniform convergence give generalization error bounds $\varepsilon$ of order $O(\sqrt{n p \log (m \lambda) / m})$ with $\eta=0$. The method by Maurer and Pontil (2010) results in Theorem 3 given below providing generalization error bounds of order

$$
O\left(\sqrt{p \min (p, n)(\lambda+\sqrt{\log (m \lambda)})^{2} / m}\right) .
$$

Thus the latter are applicable to the case $n \gg p$, while our approach is not. However in the case $n<p$, also known in the literature as the "over-complete" case (Olshausen and Fieldt, 1997; Lewicki et al., 1998), the important complexity parameter is $\lambda$, on which our bounds depend only logarithmically, instead of polynomially. One case where this is significant is where the representation is chosen by solving a minimization problem such as $\min _{a}\|D a-X\|+\gamma \cdot\|a\|_{1}$ in which $\lambda=O\left(\gamma^{-1}\right)$.

Fast rates. For the case $\eta>0$ our methods are compatible with general fast rate methods of Bartlett et al. (2005), for bounds of order $O(n p \log (\lambda m) / m)$. The main significance of this is not in the numerical results achieved, due to the large constants, but in that the general statistical behavior they imply occurs in dictionary learning. For example, generalization error has a "proportional" component which is reduced when the empirical error is low. Whether fast rates results can be proved under the infinite dimension regime is an interesting question we leave open. Note that due to lower bounds by Bartlett et al. (1998) of order $\sqrt{m^{-1}}$ on the $k$-means clustering problem, which corresponds to dictionary learning for 1 -sparse representation, fast rates may be expected only with $\eta>0$, as presented here.

We now describe the relevant function class and the bounds on its complexity, which are proved in Section 3, proving the following theorem The resulting generalization bounds are given explicitly at the end of this section.

Theorem 1. The function class $\mathcal{G}_{\lambda}=\left\{h_{R_{\lambda}, D}: \mathbb{S}^{n-1} \rightarrow \mathbb{R}: D \in \mathbb{R}^{n \times p},\left\|d_{i}\right\| \leq 1\right\}$, taken as a metric space with the metric induced by $\|\cdot\|_{\infty}$, has an $\varepsilon$ cover of cardinality at most $(4 \lambda / \varepsilon)^{n p}$.

Extension to $k$ sparse representation. Our second main contribution is to extend both our approach and that of Maurer and Pontil (2010) to provide generalization bounds for dictionaries for $k$ sparse representations, by using a bound $\lambda$ on the $l_{1}$ norm of the representation coefficients when the dictionaries are close to orthogonal. Distance from orthogonality is measured by the Babel function, defined below and discussed in more detail in Section 4

Definition 1 (Babel function, Tropp 2004). For any $k \in \mathbb{N}$, the Babel function $\mu_{k}: \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^{+}$ is defined by:

$$
\mu_{k}(D)=\max _{\Lambda \subset\{1, \ldots, p\} ;|\Lambda|=k} \max _{i \notin \Lambda} \sum_{\lambda \in \Lambda}\left|\left\langle d_{\lambda}, d_{i}\right\rangle\right| .
$$

The following proposition, which is proved in Section 3, bounds the 1-norm of the dictionary coefficients for a $k$ sparse representation and also follows from analysis previously done by Donoho and Elad (2003); Tropp (2004).

Proposition 1. Let $\left\|d_{i}\right\| \in[1, \gamma]$ and $\mu_{k-1}(D)<1$, then a coefficient vector $a \in \mathbb{R}^{p}$ minimizing the $k$-sparse representation error $h_{H_{k}, D}(x)$ exists which has $\|a\|_{1} \leq \gamma k /\left(1-\mu_{k-1}(D)\right)$.

We now consider the class of all $k$ sparse representation error functions. We prove in Section 3 the following bound on the complexity of this class.
Corollary 2. The function class $\mathcal{F}_{\delta, k}=\left\{h_{H_{k}, D}: \mathbb{S}^{n-1} \rightarrow \mathbb{R}: \mu_{k-1}(D)<\delta\right\}$, taken as a metric space with the metric induced by $\|\cdot\|_{\infty}$, has an $\varepsilon$ cover of cardinality at most $(4 k /(\varepsilon(1-\delta)))^{n p}$.

The dependence of the last two results on $\mu_{k-1}(D)$ means that the resulting bounds will be meaningful only for algorithms which explicitly or implicitly prefer near orthogonal dictionaries. Contrast this to Theorem 1 which has no significant conditions on the dictionary.

Asymptotically almost all dictionaries are near orthogonal. A question that arises is what values of $\mu_{k-1}$ can be expected for parameters $n, p, k$ ? We discuss this question and prove the following probabilistic result in Section 4

Theorem 2. Suppose that $D$ consist of $p$ vectors chosen uniformly and independently from $\mathbb{S}^{n-1}$. Then we have

$$
P\left(\mu_{k}>\frac{1}{2}\right) \leq \frac{1}{\left(e^{(n-2) /(10 k \log p)^{2}}-1\right)}
$$

Since low values of the Babel function have implications to representation finding algorithms, this result is of interest also outside the context of dictionary learning. Essentially it means that random dictionaries of size sub-exponential in $(n-2) / k^{2}$ have low Babel function.

New generalization bounds for $l_{1}$ case. The covering number bound of Theorem 1 implies several generalization bounds for the problem of dictionary learning for $l_{1}$ regularized representation which we give here. These differ from those by Maurer and Pontil (2010) in depending more strongly on the dimension of the space, but less strongly on the particular regularization term. We first give the relevant specialization of the result by Maurer and Pontil (2010) for comparison and for reference as we will later build on it. This result is independent of the dimension $n$ of the underlying space, thus the Euclidean unit ball $B$ may be that of a general Hilbert space, and the errors measured by $h_{A, D}$ are in the same norm.

Theorem 3 (Maurer and Pontil 2010). Let $\max _{a \in A}\|a\|_{1} \leq \lambda$, and $\nu$ be any distribution on the unit sphere $B$. Then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all dictionaries $D \subset B$ with cardinality $p$ :

$$
E h_{A, D}^{2} \leq E E_{m} h_{A, D}^{2}+\sqrt{\frac{p^{2}\left(14 \lambda+1 / 2 \sqrt{\ln \left(16 m \lambda^{2}\right)}\right)^{2}}{m}}+\sqrt{\frac{x}{2 m}} .
$$

Using the covering number bound of Theorem 1 and a bounded differences concentration inequality (see Lemma 9), we obtain the following result. The details are given in Section 3

Theorem 4. Let $\lambda>0$, with $\nu$ a distribution on $\mathbb{S}^{n-1}$. Then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all $D$ with unit length columns:

$$
E h_{R_{\lambda}, D} \leq E E_{m} h_{R_{\lambda}, D}+\sqrt{\frac{n p \ln (4 \sqrt{m} \lambda)}{2 m}}+\sqrt{\frac{x}{2 m}}+\sqrt{\frac{4}{m}} .
$$

Using the same covering number bound and localized Rademacher complexity (see Lemma 10), we obtain the following fast rates result.

Theorem 5. Let $\lambda>0, K>1, \alpha>0$, with $\nu$ a distribution on $\mathbb{S}^{n-1}$. Then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all $D$ with unit length column:

$$
\begin{aligned}
E h_{R_{\lambda}, D} & \leq \frac{K}{K-1} E_{m} h_{R_{\lambda}, D}+6 K \max \left\{\frac{8 \alpha \lambda^{2}}{m},(480)^{2} \frac{(n p+1) \log \left(\frac{m}{\alpha}\right)}{m}, \frac{20+22 \log (m)}{m}\right\} \\
& +\frac{11 x+5 K}{m} .
\end{aligned}
$$

In any particular case, $\alpha$ and then $K$ may be chosen so as to minimize the right hand side.
Generalization bounds for $k$ sparse representation. Proposition 1 and Corollary 2 imply certain generalization bounds for the problem of dictionary learning for $k$ sparse representation, which we give here.

A straight forward combination of Theorem 2 of Maurer and Pontil (2010) (given here as Theorem 3) and Proposition 1 results in the following theorem.

Theorem 6. Let $\delta<1$ with $\nu$ a distribution on $\mathbb{S}^{n-1}$. Then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all $D$ s.t. $\mu_{k-1}(D) \leq \delta$ :

$$
E h_{H_{k}, D}^{2} \leq E E_{m} h_{H_{k}, D}^{2}+\frac{p}{\sqrt{m}}\left(\frac{14 k}{1-\delta}+\frac{1}{2} \sqrt{\ln \left(16 m\left(\frac{k}{1-\delta}\right)^{2}\right)}\right)+\sqrt{\frac{x}{2 m}} .
$$

In the case of clustering we have $k=1$ and $\delta=0$ and this result approaches the rates of Biau et al. (2008).

The following theorems follow from standard results and the covering number bound of Corollary 2

Theorem 7. Let $\delta<1$ with $\nu$ a distribution on $\mathbb{S}^{n-1}$. Then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all $D$ s.t. $\mu_{k-1}(D) \leq \delta$ :

$$
E h_{H_{k}, D} \leq E E_{m} h_{H_{k}, D}+\sqrt{\frac{n p \ln \frac{4 \sqrt{m} k}{1-\delta}}{2 m}}+\sqrt{\frac{x}{2 m}}+\sqrt{\frac{4}{m}}
$$

Theorem 8. Let $\delta<1<K, \alpha>0$ with $\nu$ a distribution on $\mathbb{S}^{n-1}$. Then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all $D$ s.t. $\mu_{k-1}(D) \leq \delta$ :

$$
\begin{aligned}
E h_{H_{k}, D} & \leq \frac{K}{K-1} E_{m} h_{H_{k}, D}+6 K \max \left\{\frac{8 \alpha k^{2}}{m(1-\delta)^{2}},(480)^{2} \frac{(n p+1) \log \left(\frac{m}{\alpha}\right)}{m}, \frac{20+22 \log (m)}{m}\right\} \\
& +\frac{11 x+5 K}{m}
\end{aligned}
$$

In any particular case, $\alpha$ and then $K$ may be chosen so as to minimize the right hand side.
Generalization bounds for dictionary learning in feature spaces. We further consider applications of dictionary learning to signals that are not represented as elements in a vector space, or that have a very high (possibly infinite) dimension.

In addition to providing an approximate reconstruction of signals, sparse representation can also be considered as a form of analysis, if we treat the choice of non zero coefficients and their magnitude as features of the signal. In the domain of images, this has been used to perform classification (in particular, face recognition) by Wright et al. (2008). Such analysis does not require that the data itself be represented in $\mathbb{R}^{n}$ (or in any vector space); it is enough that the similarity between data elements is induced from an inner product in a feature space. This requirement is fulfilled by using an appropriate kernel function.

Definition 3. Let $\mathcal{R}$ be a set of data representations, and let the kernel function $\kappa: \mathcal{R}^{2} \rightarrow \mathbb{R}$ and the feature mapping $\phi: \mathcal{R} \rightarrow \mathcal{H}$ be such that:

$$
\kappa(x, y)=\langle\phi(x), \phi(y)\rangle
$$

where $\mathcal{H}$ is some Hilbert space.
As a concrete example, choose a sequence of $n$ words, and let $\phi$ map a document to the vector of counts of appearances of each word in it (also called bag of words). Treating $\kappa(a, b)=\langle\phi(a), \phi(b)\rangle$ as the similarity between documents $a$ and $b$, is the well known "bag of words" approach, applicable to many document related tasks (Shawe-Taylor and Cristianini, 2004). Then the statement $\phi(a)+\phi(b) \approx \phi(c)$ does not imply that $c$ can be reconstructed from $a$ and $b$, but we might consider it indicative of the content of $c$. The dictionary of elements used for representation could be decided via dictionary learning, and it is natural to choose the dictionary so that the bags of words of documents are approximated well by small linear combinations of those in the dictionary.

As the example above suggests, the kernel dictionary learning problem is to find a dictionary $D$ minimizing

$$
\mathbb{E}_{x \sim \nu} h_{\phi, A, D}(x)
$$

where we consider the representation error function

$$
h_{\phi, A, D}(x)=\min _{a \in A}\|(\Phi D) a-\phi(x)\|_{\mathcal{H}}
$$

in which $\Phi$ acts as $\phi$ on the elements of $D, A \in\left\{R_{\lambda}, H_{k}\right\}$, and the norm $\|\cdot\|_{\mathcal{H}}$ is that induced by the kernel on the feature space $\mathcal{H}$.

Analogues of all the generalization bounds mentioned so far can be replicated in the kernel setting. The dimension free results of Maurer and Pontil (2010) apply most naturally in this setting, and may be combined with our results to cover also dictionaries for $k$ sparse representation, under reasonable assumptions on the kernel.

Proposition 2. Let $\nu$ be any distribution on $\mathcal{R}$ such that when $x \sim \nu$ we have $\|\phi(x)\| \leq 1$ with probability 1. Then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all $D \subset \mathcal{R}$ with cardinality $p$ such that $\Phi D \subset B_{\mathcal{H}}$ and $\mu_{k-1}(\Phi D) \leq \delta<1$ :

$$
E h_{\phi, H_{k}, D}^{2} \leq E_{m} h_{\phi, H_{k}, D}^{2}+\sqrt{\frac{p^{2}\left(14 k /(1-\delta)+1 / 2 \sqrt{\left.\ln \left(16 m\left(\frac{k}{1-\delta}\right)^{2}\right)\right)^{2}}\right.}{m}}+\sqrt{\frac{x}{2 m}} .
$$

Note that the Babel function is defined in terms of inner products between elements of $D$, and can therefore be computed in $\mathcal{H}$ by applications of the kernel.

This result is proved in Section 5, as well as the cover number bounds (using some additional definitions and assumptions described there) that are used to prove the remaining generalization bounds, of which one is given below.

Theorem 9. Let $\mathcal{R}$ have $\varepsilon$ covers of order $(C / \varepsilon)^{n}$. Let $\kappa: \mathcal{R}^{2} \rightarrow \mathbb{R}^{+}$be a kernel function s.t. $\kappa(x, y)=\langle\phi(X), \phi(Y)\rangle$, for $\phi$ which is uniformly L-Hölder of order $\alpha>0$ over $\mathcal{R}$, and let $\gamma=\max _{x \in \mathcal{R}}\|\phi(x)\|_{\mathcal{H}}$. Let $\delta<1$, and $\nu$ any distribution on $\mathcal{R}$, then with probability at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ drawn according to $\nu$, for all dictionaries $D \subset \mathcal{R}$ of cardinality p s.t. $\mu_{k-1}(\Phi D) \leq \delta<1$ (where $\Phi$ acts like $\phi$ on columns):

$$
E h_{H_{k}, D} \leq E E_{m} h_{H_{k}, D}+\gamma\left(\sqrt{\frac{n p \ln \left(\sqrt{m} C^{\alpha} \frac{k \gamma^{2} L}{1-\delta}\right)}{2 \alpha m}}+\sqrt{\frac{x}{2 m}}\right)+\sqrt{\frac{4}{m}} .
$$

The covering number bounds needed to prove this theorem and analogs for the other generalization bounds are proved in Section 5.

## 3 Covering numbers of $\mathcal{G}_{\lambda}$ and $\mathcal{F}_{\delta, k}$

The main content of this section is the proof of Theorem 2 and Corollary 2 We also show that the restriction of near-orthogonality on the set of dictionaries, on which we rely in the proof for $k$ sparse representation, is necessary to achieve a bound on $\lambda$. Lastly, we recall known results from statistical learning theory that link covering numbers to generalization bounds.

We recall the definition of the covering numbers we wish to bound. Anthony and Bartlett (1999) give a textbook introduction to covering numbers and their application to generalization bounds.

Definition 4 (Covering number). Let $(M, d)$ be a metric space and $S \subset M$. Then the $\varepsilon$ covering number of $S$ defined as $N(\varepsilon, S, d)=\min \left\{|A| \mid A \subset M\right.$ and $\left.S \subset\left(\bigcup_{a \in A} B_{d}(a, \varepsilon)\right)\right\}$ is the size of the minimal $\varepsilon$ cover of $S$ using $d$.

To prove Theorem 1 and Corollary 2 we first note that the space of all possible dictionaries is a subset of a unit ball in a Banach space of dimension $n p$ (with a norm specified below). Thus by proposition 5 formalized by Cucker and Smale (2002) the space of dictionaries has an $\varepsilon$ cover of size $(4 / \varepsilon)^{n p}$. We also note that a uniformly $L$ Lipschitz mapping between metric spaces converts $\varepsilon / L$ covers into $\varepsilon$ covers. Then it is enough to show that $\Psi_{\lambda}$ defined as $D \mapsto h_{R_{\lambda}, D}$ and $\Phi_{k}$ defined as $D \mapsto h_{H_{k}, D}$ are uniformly Lipschitz (when $\Phi_{k}$ is restricted to the dictionaries with $\mu_{k-1}(D) \leq c<1$ ). The proof of these Lipschitz properties is our next goal, in the form of Lemmas 7 and 8

The first step is to be clear about the metrics we consider over the spaces of dictionaries and of error functions. We start by defining the following norm.

Definition 5. Let $D \in \mathbb{R}^{n \times p}$. We denote $\|D\|_{M E}=\max _{i}\left\|d_{i}\right\|$ the norm of its maximal column.
We will use the fact $\|\cdot\|_{M E}$ upper bounds a certain induced norm.
Definition 6 (Induced matrix norm). Let $p, q \in \mathbb{N}$, then a matrix $A \in \mathbb{R}^{n \times m}$ can be considered as an operator $A:\left(\mathbb{R}^{m},\|\cdot\|_{p}\right) \rightarrow\left(\mathbb{R}^{n},\|\cdot\|_{q}\right)$. Then the $p, q$ induced norm is defined as $\|A\|_{p, q} \triangleq$ $\sup _{x \in \mathbb{R}^{m}\|x\|_{p}=1}\|A x\|_{q}$.

Fact 1. $\|D\|_{1,2} \leq\|D\|_{M E}$
The geometric interpretation of this fact is that $D a /\|a\|_{1}$ is a convex combination of vectors each of length at most $\|D\|_{M E}$, then $\|D a\|_{2} \leq\|D\|_{M E}\|a\|_{1}$.

The images of $\Psi_{\lambda}$ and $\Phi_{k}$ are sets of representation error functions-each dictionary induces a set of precisely representable signals, and a representation error function is simply a map of distances from this set. Representation error functions are clearly continuous, 1-Lipschitz, and into $[0,1]$. In this setting, a natural norm over the images is the supremum norm $\|\cdot\|_{\infty}$.

Lemma 7. The function $\Psi_{\lambda}$ is $\lambda$-Lipschitz from $\left(\mathbb{R}^{n \times m},\|\cdot\|_{M E}\right)$ to $C\left(\mathbb{S}^{n-1}\right)$.
Proof. Let $D$ and $D^{\prime}$ be two normalized dictionaries whose corresponding elements are at most $\varepsilon>0$ far from one another. Let $x$ be a unit signal and $D a$ an optimal representation for it. Then $\left\|\left(D-D^{\prime}\right) a\right\| \leq\left\|D-D^{\prime}\right\|_{1,2}\|a\|_{1} \leq\left\|D-D^{\prime}\right\|_{M E}\|a\|_{1} \leq \varepsilon \lambda$. Then $g_{\lambda, D^{\prime}}(x) \leq g_{\lambda, D}(x)+\varepsilon \lambda$ and by symmetry we have $\left|\Psi_{\lambda}(D)(x)-\Psi_{\lambda}\left(D^{\prime}\right)(x)\right| \leq \lambda \varepsilon$. This holds for all unit signals, then $\left\|\Psi_{\lambda}(D)-\Psi_{\lambda}\left(D^{\prime}\right)\right\|_{\infty} \leq \lambda \varepsilon$.

We now provide a proof for Proposition 1 which is used in the corresponding treatment for covering numbers under $k$ sparsity.

Of Proposition 1 Assume that $\mu_{k-1}(D) \leq \delta<1 \leq \min _{i \leq p}\left\|d_{i}\right\|_{2} \leq \gamma$. Let $D^{k}$ be a set of $k$ elements from $D$ achieving the minimum on $h_{H_{k}, D}(x)$, with $x \in \mathbb{S}^{n-1}$. We now consider the Gram matrix $G=\left(D^{k}\right)^{\top} D^{k}$. The matrix $G$ is symmetric, therefore it scales each point in the unit sphere by a non-negative combination of its real eigenvalues. Also, the diagonal entries of $G$ are the norms of the elements of $D^{k}$, therefore at least 1. By the Gersgorin theorem (Horn and Johnson, 1990), the eigenvalues of the Gram matrix are lower bounded by $1-\delta>0$. Then in particular $G$ has a symmetric inverse, which scales each point by no more than $1 /(1-\delta)$. Then $\left\|G^{-1}\right\|_{1,1} \leq 1 /(1-\delta)$.

In particular, elements of $D^{k}$ are linearly independent, which implies that the unique optimal representation of $x$ as a linear combination of the columns of $D^{k}$ is $D^{k} a$ with

$$
a=\left(\left(D^{k}\right)^{\top} D^{k}\right)^{-1}\left(D^{k}\right)^{\top} x
$$

By the definition of induced matrix norms, we have $\|a\|_{1} \leq\left\|\left(\left(D^{k}\right)^{\top} D^{k}\right)^{-1}\right\|_{1,1}\left\|\left(D^{k}\right)^{\top} x\right\|_{1} \leq$ $\gamma k /(1-\delta)$, the last bound because $x$ is a unit vector, and $D^{k}$ has $k$ columns whose norm is bounded by $\gamma$.

Lemma 8. The function $\Phi_{k}$ is a $k /(1-\delta)$-Lipschitz mapping from the set of normalized dictionaries with $\mu_{k-1}(D)<\delta$ with the metric induced by $\|\cdot\|_{M E}$ to $C\left(\mathbb{S}^{n-1}\right)$.

The proof of this lemma is the same as that of Lemma 7 except that $a$ is taken to be an optimal representation that fulfills $\|a\|_{1} \leq \lambda=k /\left(1-\mu_{k-1}(D)\right)$, whose existence is guaranteed by Proposition 1

This concludes the proof of Theorem 1 and Corollary 2
The next theorem shows that unfortunately, $\Phi$ is not uniformly $L$-Lipschitz for any constant $L$, requiring its restriction to an appropriate subset of the dictionaries.

Theorem 10. For any $k, n, p$, there exists $c>0$ and $q$, such that for every $\varepsilon>0$, there exist $D, D^{\prime}$ such that $\left\|D-D^{\prime}\right\|_{M E}<\varepsilon$ but $\left|\left(h_{H_{k}, D}(q)-h_{H_{k}, D^{\prime}}(q)\right)\right|>c$.

Proof. First we show that there exists $c>0$ such that every dictionary will have $k$ sparse representation error of at least $c$ on some signal. Let $\nu_{S^{n-1}}$ be the uniform probability measure on the sphere, and $A_{c}$ the probability assigned by it to the set within $c$ of a $k$ dimensional subspace. As $c \searrow 0, A_{c}$ also tends to zero, then there exists $c>0$ s.t. $\binom{p}{k} A_{c}<1$. Then for that $c$ there exists a set of positive measure on which $h_{H_{k}, D}>c$, let $q$ be a point in this set.

To complete the proof we consider a dictionary $D$ whose first $k-1$ elements are the standard basis $\left\{e_{1}, \ldots, e_{k-1}\right\}$, its $k$ the element is $D_{k}=\sqrt{1-\varepsilon^{2} / 2} e_{1}+\varepsilon e_{k} / 2$, and the remaining elements are chosen arbitrarily. Now construct $D^{\prime}$ to be identical to $D$ except its $k$ th element is $v=\sqrt{1-\varepsilon^{2} / 2} e_{1}+l q$ choosing $l$ so that $\|v\|_{2}=1$ (which implies that $|l|<\varepsilon / 2$ ). Then $\left\|D-D^{\prime}\right\|_{M E}=\left\|\varepsilon e_{k} / 2+l q\right\|_{2} \leq \varepsilon$ and $h_{H_{k}, D^{\prime}}(q)=0$.

To conclude the generalization bounds of Theorems 4,5,7, 8 and 9 from the covering number bounds we have provided, we use the following two results. The first has a simple proof which we therefore give here. The second result is an adaptation of results by Bartlett et al. (2005), to our needs, and explained further in the appendix.

Lemma 9. Let $\mathcal{F}$ be a class of $[0, B]$ functions with covering number bound $(C / \varepsilon)^{d}>e / B^{2}$ under the supremum norm. Then for every $x>0$, with probability of at least $1-e^{-x}$ over the $m$ samples in $E_{m}$ chosen according to $\nu$, for all $f \in \mathcal{F}$ :

$$
E f \leq E_{m} f+B\left(\sqrt{\frac{d \log (C \sqrt{m})}{2 m}}+\sqrt{\frac{x}{2 m}}\right)+\sqrt{\frac{4}{m}} .
$$

Lemma 10. If $\mathcal{F}$ is a class of $[0,1]$ functions with $C>2$ and $d \in \mathbb{N}$ s.t. $N\left(\varepsilon, \mathcal{F}, L_{2}(\nu)\right) \leq\left(\frac{C}{\varepsilon}\right)^{d}$ for every probability measure $\nu$ and $\varepsilon>0$, then for all $K, \alpha, x>0, f \in \mathcal{F}$, with probability at least $1-e^{-x}$ over the $m$ samples used in $E_{m}$ and drawn from $\nu$ :
$E f \leq \frac{K}{K-1} E_{m} f+6 K \max \left\{\frac{\alpha C^{2}}{2 m},(480)^{2} \frac{(d+1) \log \left(\frac{m}{\alpha}\right)}{m}, \frac{20+22 \log (m)}{m}\right\}+\frac{11 x+5 K}{m}$.
Our fast rates results are simple applications of this lemma, noting that an $\varepsilon$ cover in $C\left(\mathbb{S}^{n-1}\right)$ is also an $\varepsilon$ cover under an $L_{2}$ metric induced by any measure.

Of Lemma 9 We wish to bound $\sup _{f \in \mathcal{F}} E f-E_{m} f$. Take $\mathcal{F}_{\varepsilon}$ to be a minimal $\varepsilon$ cover of $\mathcal{F}$, then for an arbitrary $f$, denoting $f_{\varepsilon}$ an $\varepsilon$ close member of $\mathcal{F}_{\varepsilon}, E f-E_{m} f \leq E f_{\varepsilon}-E_{m} f_{\varepsilon}+2 \varepsilon$. In particular, $\sup _{f \in \mathcal{F}} E f-E_{m} f \leq 2 \varepsilon+\sup _{f \in \mathcal{F}_{\varepsilon}} E f-E_{m} f$. To bound the supremum on the now finite class of functions, note that $E f-E_{m} f$ is a function of $m$ independent variables (the samples chosen according to $\nu$ ), which changes by at most $B / m$ when one of the variables is modified. Then by the bounded differences inequality, $P\left(E f-E_{m} f-\mathbb{E}\left(E f-E_{m} f\right)>t\right)=P\left(E f-\mathbb{E}_{m} f>t\right) \leq$ $\exp \left(-2 m B^{-2} t^{2}\right)$.

The probability that any of the $\left|\mathcal{F}_{\varepsilon}\right|$ differences under the supremum is larger than $t$ may be union bounded as $\left|\mathcal{F}_{\varepsilon}\right| \cdot \exp \left(-2 m B^{-2} t^{2}\right) \leq \exp \left(d \log (C / \varepsilon)-2 m B^{-2} t^{2}\right)$.

In order to control the probability with $x$ as in the statement of the lemma, we need to have $x=d \log (C / \varepsilon)-m B^{-2} t^{2}$ and thus we choose $t=\sqrt{B^{2} / 2 m} \sqrt{d \log (C / \varepsilon)+x}$. Then with high probability we bound the supremum of differences by $t$ which is upper bounded, using the assumption on the covering number bound, by $B(\sqrt{d \log (C / \varepsilon) / 2 m}+\sqrt{x / 2 m})$.

Then the proof is completed by substitution into the bound over the whole function class $\mathcal{F}$ and taking $\varepsilon=1 / \sqrt{m}$.

## 4 On the Babel function

The Babel function is one of several metrics defined in the sparse representations literature to quantify an "almost orthogonality" property that dictionaries may enjoy. Such properties have been shown to imply theoretical properties such as uniqueness of the optimal $k$ sparse representation. In the algorithmic context, Donoho and Elad (2003) and Tropp (2004) use the Babel function to show that particular tractable algorithms for finding sparse representations are indeed approximation algorithms when applied to such dictionaries. This reinforces the practical importance of the learnability of this class of dictionary. We proceed to discuss some elementary properties of the Babel function, and then state a bound on the proportion of dictionaries having sufficiently good Babel function.

Measures of orthogonality are typically defined in terms of inner products between the elements of the dictionary. Perhaps the simplest of these measures of orthogonality is the following special case of the Babel function.

Definition 11. The coherence of a dictionary $D$ is $\mu_{1}(D)=\max _{i, j}\left|\left\langle d_{i}, d_{j}\right\rangle\right|$.
The Babel function, in considering sums of $k$ inner products at a time, rather than the maximum over all inner products, is better adapted to quantify the effects of non orthogonality on representing
a signal with particular level $k+1$ of sparsity. The additional expressive power of $\mu_{k}$ over $\mu_{1}$ is illustrated by considering that ensuring that $\mu_{k}<1$ by restricting $\mu_{1}$ implies the constraint $\mu_{1}(D)<1 / k$, which for $k>1$ would exclude a dictionary in which pairs of elements have inner product 0 except for some disjoint pairs whose inner product equals to half, despite such a dictionary having $\mu_{k}=1 / 2$ for any $k$.

To better understand $\mu_{k}(D)$, we consider first its extreme values. When $\mu_{k}(D)=0$, for any $k>1$, this means that $D$ is an orthogonal set (therefore $p \leq n$ ). The maximal value of $\mu_{k}(D)$ is $k$, and occurs only if some dictionary element is repeated (up to sign) at least $k+1$ times.

A well known generic class of dictionaries with more elements than a basis is that of frames (see Duffin and Schaeer, 1952), which include many wavelet systems and filter banks. Some frames can be trivially seen to fulfill our condition on the Babel function.

Proposition 3. Let $D \in \mathbb{R}^{n \times p}$ be a frame of $\mathbb{R}^{n}$, so that for every $v \in \mathbb{S}^{n-1}$ we have that $A \leq$ $\sum_{i=1}^{n}\left|\left\langle v, d_{i}\right\rangle\right| \leq B$, with $\left\|d_{i}\right\|_{2}=1$ for all $i$, and $B<1+1 /(p-1)$. Then $\mu_{k-1}(D)<1$.

This may be easily verified using the relation between $\|\cdot\|_{1}$ and $\|\cdot\|_{2}$ in $\mathbb{R}^{p-1}$.

### 4.1 Proportion of dictionaries with $\mu_{k-1}(D)<\delta$

We return to the question of the prevalence of dictionaries from $D_{\delta}$. Are almost all dictionaries in $D_{\delta}$ ? If the answer is affirmative, it implies that Theorem 8 is quite strong, and representation finding algorithms such as basis pursuit are almost always exact, which might help prove properties of dictionary learning algorithms. If the opposite is true and few dictionaries are in $D_{\delta}$, the results of this paper are weak. While there might be better measures on the space of dictionaries, we consider one that seems natural: suppose that a dictionary $D$ is constructed by choosing $p$ unit vectors uniformly from $\mathbb{S}^{n-1}$; what is the probability that $\mu_{k-1}(D)<\delta$ ?

Theorem 2 gives us the following answer to this question. Under the assumption that the sparsity parameter $k$ grows slowly, if at all, as $n \nearrow \infty$ (specifically, that $k \log p=o(\sqrt{n})$ ), this theorem implies that asymptotically almost all dictionaries under the Lebesgue measure are learnable.

The remainder of this section is devoted to the proof of Theorem 2. This proof relies heavily on the Orlicz norms for random variables and their properties; Van der Vaart and Wellner (1996) give a detailed introduction. We recall a few of the definitions and facts presented there.

Definition 12. Let $\psi$ be a non-decreasing, convex function with $\psi(0)=0$, and let $X$ be a random variable. Then

$$
\|X\|_{\psi}=\inf \left\{C>0: \mathbb{E} \psi\left(\frac{|X|}{C}\right)<1\right\}
$$

is called an Orlicz norm.
As may be verified, these are indeed norms for appropriate $\psi$, such as $\psi_{2} \triangleq e^{x^{2}}-1$, which is the case that will interest us most.

By the Markov inequality we can obtain that variables with finite Orlicz norms have light tails.
Fact 2. We have $P(|X|>x) \leq\left(\psi_{2}\left(x /\|X\|_{\psi_{2}}\right)\right)^{-1}$.
The next fact is an almost converse to the last fact, stating that light tailed random variables have finite $\psi_{2}$ Orlicz norms.

Fact 3. Let $A, B>0$ and $P(|X| \geq x) \leq A e^{-B x^{2}}$ for all $x$, where $p \geq 1$, then $\|X\|_{\psi_{2}} \leq$ $((1+A) / B)^{1 / 2}$.

The following bound on the maximum of variables with light tails.
Fact 4. We have $\left\|\max _{1 \leq i \leq m} X_{i}\right\|_{\psi_{2}} \leq K \sqrt{\log m} \max _{i}\left\|X_{i}\right\|_{\psi_{2}}$.
The constant $K$ may be upper bounded by $\sqrt{2}$. Note that the independence of $X_{i}$ is not required. We use also one isoperimetric fact about the sphere in high dimension.

Definition 13. The $\varepsilon$ expansion of a set $D$ in a metric space $(X, d)$ is defined as

$$
D_{\varepsilon}=\{x \in X \mid d(x, D) \leq \varepsilon\},
$$

where $d(x, A)=\inf _{a \in A} d(x, a)$.
Fact 5 (Lévy's isoperimetric inequality 1952). Let $C$ be one half of $\mathbb{S}^{n-1}$, then $\mu\left(\left(\mathbb{S}^{n-1} \backslash C_{\varepsilon}\right)\right) \leq$ $\sqrt{\frac{\pi}{8}} \exp \left(-\frac{(n-2) \varepsilon^{2}}{2}\right)$.

Our goal in the reminder of this subsection is to obtain the following bound.
Lemma 14. Let $D$ be a dictionary chosen at random as described above, then

$$
\left\|\mu_{k}(D)\right\|_{\psi_{2}} \leq 5 k \log p / \sqrt{n-2}
$$

Our probabilistic bound on $\mu_{k-1}$ is a direct conclusion of Fact 3 and Lemma 14 which we now proceed to prove. The plan of our proof is to bound the $\psi_{2}$ metric of $\mu_{k}$ from the inside terms and outward using Fact 4 to overcome the maxima over possibly dependent random variables.

Lemma 15. Let $X_{1}, X_{2}$ be unit vectors chosen uniformly and independently from $\mathbb{S}^{n-1}$, then

$$
\left\|\left|\left\langle X_{1}, X_{2}\right\rangle\right|\right\|_{\psi_{2}} \leq \sqrt{6 /(n-2)} .
$$

We denote the bound on the right hand side $W$.
Proof. Taking $X$ to be uniformly chosen from $\mathbb{S}^{n-1}$, for any constant unit vector $x_{0}$ we have that $\left\langle X, x_{0}\right\rangle$ is a light tailed random variable by Fact [5 By Fact 3] we may bound $\left\|\left\langle X, x_{0}\right\rangle\right\|_{\psi_{2}}$. Replacing $x_{0}$ by a random unit vector is equivalent to applying to $X$ a uniformly chosen rotation, which does not change the analysis.

The next step is to bound the inner maximum appearing in the definition of $\mu_{k}$.
Lemma 16. Let $\left\{d_{i}\right\}_{i=1}^{p}$ be uniformly and independently chosen unit vectors then

$$
\left\|\max _{\Lambda \subset\{2, \ldots, p\} \wedge|\Lambda|=k} \sum_{\lambda \in \Lambda}\left|\left\langle d_{1}, d_{\lambda}\right\rangle\right|\right\|_{\psi_{2}} \leq k K W \sqrt{\log (p-1)} .
$$

Proof. Take $X_{\lambda}$ to be $\left\langle D_{1}, D_{\lambda}\right\rangle$. Then using Fact 4 and the previous lemma we find

$$
\left\|\max _{1 \leq \lambda \leq p \wedge \lambda \neq i}\left|X_{\lambda}\right|\right\|_{\psi_{2}} \leq K \sqrt{\log (p-1)} \max _{\lambda}\left\|\left|X_{\lambda}\right|\right\|_{\psi_{2}} \leq K W \sqrt{\log (p-1)} .
$$

Define the random permutation $\lambda_{j}$ s.t. $\left|X_{\lambda_{j}}\right|$ are non-increasing. In this notation, it is clear that $\max _{\Lambda \subset\{2 \ldots p\} \wedge|\Lambda|=k} \sum_{\lambda \in \Lambda}\left|X_{\lambda}\right|=\sum_{j=1}^{k}\left|X_{\lambda_{j}}\right|$. Note that $\left|X_{\lambda_{i}}\right| \leq\left|X_{\lambda_{1}}\right|$ then for every i , $\left\|\left|X_{\lambda_{i}}\right|\right\|_{\psi_{2}} \leq\left\|\left|X_{\lambda_{1}}\right|\right\|_{\psi_{2}} \leq K W \sqrt{\log (p-1)}$.

By the triangle inequality, $\left\|\sum_{j=1}^{m}\left|X_{\lambda_{j}}\right|\right\|_{\psi_{2}} \leq \sum_{j=1}^{m}\left\|\left|X_{\lambda_{j}}\right|\right\|_{\psi_{2}} \leq m K W \sqrt{\log (p-1)}$.
Remark 1. Two facts are relevant to the tightness of the approximations in the last proof. First, that $\left|X_{\lambda_{i}}\right|$ are variables with positive expectation bounded away from zero, thus the norm of their sum must scale at least linearly in the number of summands, so the triangle inequality is essentially tight. Second we consider the bound $\left\|\left|X_{\lambda_{i}}\right|\right\|_{\psi_{2}} \leq\left\|\left|X_{\lambda_{1}}\right|\right\|_{\psi_{2}}$, and note its looseness is strictly limited by the slow growth of $\sqrt{\log (\cdot)}$, and in any case is bounded by 2 .

To complete the proof of Lemma 14 we replace $D_{1}$ with the dictionary element maximizing the Orlicz norm, by another application of Fact 4 and to complete the proof of Theorem 2, apply Fact 3 to the estimated Orlicz norm.

## 5 Dictionary learning in feature spaces

We propose in Section 2 a scenario in which dictionary learning is performed in a feature space corresponding to a kernel function. Here we show how to adapt the different generalization bounds discussed in this paper for the particular case of $\mathbb{R}^{n}$ to more general feature spaces, and the dependence of the sample complexities on the properties of the kernel function or the corresponding feature mapping. We begin with the relevant specialization of the results of Maurer and Pontil (2010) which have the simplest dependence on the kernel, and then discuss the extensions to $k$ sparse representation and to the cover number techniques presented in the current work.

Theorem [3 applies as is to the feature space, under the simple assumption that the dictionary elements and signals are in its unit ball which is guaranteed by some kernels such as the Gaussian kernel. Then we take $\nu$ on the unit ball of $\mathcal{H}$ to be induced by some distribution $\nu^{\prime}$ on the domain of the kernel, and the theorem applies to any such $\nu^{\prime}$ on $\mathcal{R}$. Nothing more is required if the representation is chosen from $R_{\lambda}$. The corresponding generalization bound for $k$ sparse representations when the dictionary elements are near orthogonal in the feature space is given in Proposition 2

Of Proposition 2 Proposition 1 applies with the Euclidean norm of $\mathcal{H}$, and $\gamma=1$. We apply Theorem 3 with $\lambda=k /(1-\delta)$.

The results so far show that generalization in dictionary learning can occur despite the potentially infinite dimension of the feature space, without considering practical issues of representation and computation. We now make the domain and applications of the kernel explicit in order to address a basic computational question, and allow the use of cover number based generalization bounds to prove Theorem 9 We now consider signals represented in a metric space $(\mathcal{R}, d)$, in which similarity is measured by the kernel $\kappa$ corresponding to the feature map $\phi: \mathcal{R} \rightarrow \mathcal{H}$. The
elements of a dictionary $D$ are now from $\mathcal{R}$, and we denote $\Phi D$ their mapping by $\phi$ to $\mathcal{H}$. Then representation error function used is $h_{\phi, A, D}$.

We now show that the approximation error in feature space is a quadratic function of the coefficient vector, which may be found by applications of the kernel.
Proposition 4. Computing the representation error at a given $x, a, D$ requires $O\left(p^{2}\right)$ kernel applications in general, and only $O\left(k^{2}+p\right)$ when a is $k$ sparse.

Proof. Writing the error;

$$
\begin{aligned}
\|(\Phi D) a-\phi(x)\|^{2} & =\langle(\Phi D) a-\phi(x),(\Phi D) a-\phi(x)\rangle \\
& =\langle(\Phi D) a,(\Phi D) a\rangle+\langle\phi(x), \phi(x)\rangle-2\langle\phi(x),(\Phi D) a\rangle \\
& =\left\langle\sum_{i=1}^{p} \phi\left(d_{i}\right) a_{i}, \sum_{j=1}^{p} \phi\left(d_{j}\right) a_{j}\right\rangle+\langle\phi(x), \phi(x)\rangle-2\left\langle\phi(x), \sum_{i=1}^{p} \phi\left(d_{i}\right) a_{i}\right\rangle \\
& =\sum_{i=1}^{p} a_{i} \sum_{j=1}^{p} a_{j}\left\langle\phi\left(d_{i}\right), \phi\left(d_{j}\right)\right\rangle+\langle\phi(x), \phi(x)\rangle-2 \sum_{i=1}^{p} a_{i}\left\langle\phi(x), \phi\left(d_{i}\right)\right\rangle \\
& =\sum_{i=1}^{p} a_{i} \sum_{j=1}^{p} a_{j} \kappa\left(d_{i}, d_{j}\right)+\kappa(x, x)-2 \sum_{i=1}^{p} a_{i} \kappa\left(x, d_{i}\right)
\end{aligned}
$$

We note that the $k$ sparsity constraint on $a$ poses algorithmic difficulties beyond those addressed here. Some of the common approaches to these, such as Orthogonal Matching Pursuit (Chen et al., 1989), also depend on the data only through their inner products, and may therefore be adapted to the kernel setting.

The cover number bounds depend strongly on the dimension of the space of dictionary elements. Taking $\mathcal{H}$ as the space of dictionary elements is the simplest approach, but may lead to vacuous or weak bounds, for example in the case of the Gaussian kernel whose feature space is infinite dimensional. Instead we propose to use the space of data representations $\mathcal{R}$, whose dimensions are generally bounded by practical considerations. In addition, we will assume that the kernel is not "too wild" in the following sense.

Definition 17. Let $L, \alpha>0$, and let $\left(A, d^{\prime}\right)$ and $(B, d)$ are metric spaces. We say a mapping $f: A \rightarrow B$ is uniformly $L$ Hölder of order $\alpha$ on a set $S \subset A$ if $\forall x, y \in S$, the following bound holds:

$$
d(f(x), f(y)) \leq L \cdot d^{\prime}(x, y)^{\alpha} .
$$

The relevance of this smoothness condition is as follows:
Fact 6. A Hölder function maps an $\varepsilon$ cover of $S$ to an $L \varepsilon^{\alpha}$ cover of its image $f(S)$. Thus, to obtain an $\varepsilon$ cover of the image of $S$, it is enough to begin with an $(\varepsilon / L)^{1 / \alpha}$ cover of $S$.

A Hölder feature map $\phi$ allows us to bound the cover numbers of the dictionary elements in $\mathcal{H}$ using their cover number bounds in $\mathcal{R}$. Note that not every kernel corresponds to a Hölder feature map (the Dirac $\delta$ kernel is a counter example: any two distinct elements are mapped to elements at a mutual distance of 1 ), and not for every kernel the feature map is known. The following lemma bounds the geometry of the feature map using that of the kernel.

Lemma 18. Let $\kappa(x, y)=\langle\phi(x), \phi(y)\rangle$, and assume further that $\kappa$ fulfills a Hölder condition of order $\alpha$ uniformly in each parameter, that is, $|\kappa(x, y)-\kappa(x+h, y)| \leq L\|h\|^{\alpha}$. Then $\phi$ uniformly fulfills a Hölder condition of order $\alpha / 2$ with constant $\sqrt{2 L}$.

This result is not sharp. For example, for the Gaussian case, both kernel and the feature map are Hölder order 1.

Proof. Using the Hölder condition, we have that $\|\phi(x)-\phi(y)\|_{\mathcal{H}}^{2}=\kappa(x, x)-\kappa(x, y)+\kappa(y, y)-$ $\kappa(x, y) \leq 2 L\|x-y\|^{\alpha}$. All that remains is to take the square root of both sides.

For a given feature mapping $\phi$, set of representations $\mathcal{R}$, we define two families of function classes so:

$$
\begin{aligned}
\mathcal{W}_{\phi, \lambda} & =\left\{h_{\phi, R_{\lambda}, D}: D \in \mathcal{D}^{p}\right\} \text { and } \\
\mathcal{Q}_{\phi, k, \delta} & =\left\{h_{\phi, H_{k}, D}: D \in \mathcal{D}^{p} \wedge \mu_{k-1}(\Phi D) \leq \delta\right\} .
\end{aligned}
$$

The next proposition completes this section by giving the cover number bounds for the representation error function classes induced by appropriate kernels, from which various generalization bounds easily follow, such as Theorem 9

Proposition 5. Let $\mathcal{R}$ be a set of representations with a cover number bound of $(C / \varepsilon)^{n}$, and let either $\phi$ be uniformly L Hölder condition of order $\alpha$ on $\mathcal{R}$, or $\kappa$ be uniformly L Hölder of order $2 \alpha$ on $\mathcal{R}$ in each parameter, and let $\gamma=\sup _{d \in \mathcal{R}}\|\phi(d)\|_{\mathcal{H}}$. Then the function classes $\mathcal{W}_{\phi, \lambda}$ and $\mathcal{Q}_{\phi, k, \delta}$ taken as metric spaces with the supremum norm, have $\varepsilon$ covers of cardinalities at most $\left(C(\lambda \gamma L / \varepsilon)^{1 / \alpha}\right)^{n p}$ and $\left(C\left(k \gamma^{2} L /(\varepsilon(1-\delta))\right)^{1 / \alpha}\right)^{n p}$, respectively.

Proof. We first consider the simpler case of $l_{1}$ constrained coefficients. If $\|a\|_{1} \leq \lambda$ and also $\max _{d \in \mathcal{D}}\|\phi(d)\|_{\mathcal{H}} \leq \gamma$ then by the considerations applied in section 3, to obtain an $\varepsilon$ cover of the set $\left\{\min _{a}\|(\Phi D) a-\phi(x)\|_{\mathcal{H}}: D \in \mathcal{D}\right\}$, it is enough to obtain an $\varepsilon /(\lambda \gamma)$ cover of $\{\Phi D: D \in \mathcal{D}\}$. If also $\phi$ is uniformly $L$ Hölder of order $\alpha$ over $\mathcal{R}$ then an $(\lambda \gamma L / \varepsilon)^{-1 / \alpha}$ cover of the set of dictionaries is sufficient, which as we have seen requires at most $\left(C(\lambda \gamma L / \varepsilon)^{1 / \alpha}\right)^{n p}$ elements.

In the case of $l_{0}$ constrained representation, the bound on $\lambda$ due to Proposition $\square$ is $\gamma k(1-\delta)$, and the result follows from the above by substitution.

## 6 Conclusions

Our work has several implications on the design of dictionary learning algorithms as used in signal, image, and natural language processing. First, the fact that generalization is only logarithmically dependent on the $l_{1}$ norm of the coefficient vector widens the set of applicable approaches to penalization. Second, in the particular case of $k$ sparse representation, we have shown that the Babel function is a key property for the generalization of dictionaries. It might thus be useful to modify dictionary learning algorithms so that they obtain dictionaries with low Babel functions, possibly through regularization or through certain convex relaxations. Third, mistake bounds (e.g., Mairal et al. 2010) on the quality of the solution to the coefficient finding optimization problem
may lead to generalization bounds for practical algorithms, by tying such algorithms to $k$ sparse representation.

The upper bounds presented here invite complementary lower bounds. The existing lower bounds for $k=1$ (vector quantization) and for $k=p$ (representation using PCA directions) are applicable, but do not capture the geometry of general $k$ sparse representation, and in particular do not clarify the effective dimension of the unrestricted class of dictionaries for it. We have not excluded the possibility that the class of unrestricted dictionaries has the same dimension as that of those with small Babel function. The best upper bound we know for the larger class, being the trivial one of order $\left.O\binom{p}{k} n^{2} / m\right)$, leaves a significant gap for future exploration.

We mention also that the dependence on $\mu_{k-1}$ can also be viewed from an "algorithmic luckiness" perspective (Herbrich and Williamson, 2003): if the dictionary has favorable geometry in the sense that the Babel function is small the generalization bounds are encouraging.

## Acknowledgments

We thank Shahar Mendelson for helpful discussions. This research was partly supported by the European Communitys FP7-FET program, SMALL project, under grant agreement no. 225913.

## References

Michal Aharon, Michael Elad, and Alfred M. Bruckstein. K-SVD: Design of dictionaries for sparse representation. In PROCEEDINGS OF SPARS, pages 9-12, 2005.

Martin Anthony and Peter L. Bartlett. Neural network learning: Theoretical foundations. Cambridge Univ Pr, 1999.

Peter L. Bartlett, Tamás Linder, and Gábor Lugosi. The minimax distortion redundancy in empirical quantizer design. IEEE Transactions on Information theory, 44(5):1802-1813, 1998.

Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. Ann. Statist., 33:1497-1537, 2005.

Gérard Biau, Luc Devroye, and Gábor Lugosi. On the performance of clustering in Hilbert spaces. Information Theory, IEEE Transactions on, 54(2):781-790, 2008. ISSN 0018-9448.
G. Blanchard, O. Bousquet, and L. Zwald. Statistical properties of kernel principal component analysis. Machine Learning, 66(2):259-294, 2007. ISSN 0885-6125.

Alfred M. Bruckstein, David L. Donoho, and Michael Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. SIAM Review, 51(1):34-81, 2009.

Emmanuel J. Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? IEEE Transactions on Information Theory, 52(12):5406-5425, 2006.
S. Chen, SA Billings, and W. Luo. Orthogonal least squares methods and their application to nonlinear system identification. International Journal of Control, 50(5):1873-1896, 1989.

Scott S. Chen, David L. Donoho, and Michael A. Saunders. Atomic decomposition by basis pursuit. SIAM Review, 43(1):129-159, 2001.
F. Cucker and S. Smale. On the mathematical foundations of learning. BULLETIN-AMERICAN MATHEMATICAL SOCIETY, 39(1):1-50, 2002.

Geoff Davis, Stèphane Mallat, and Marco Avellaneda. Adaptive greedy approximations. Constructive approximation, 13(1):57-98, 1997.

David L. Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via $\ell_{1}$ minimization. Proceedings of the National Academy of Sciences, 100(5):2197, 2003.

Richard J. Duffin and Albert C. Schaeer. A class of nonharmonic Fourier series. Trans. Amer. Math. Soc, 72:341-366, 1952.

Rémi Gribonval and Karin Schnass. Dictionary identification - sparse matrix-factorisation via $\ell_{1}-$ minimisation. Available at: arXiv:0904.4774v2 [cs.IT], 2009.

Ralf Herbrich and Robert Williamson. Algorithmic luckiness. JMLR, 3:175-212, 2003. ISSN 1532-4435.

Roger A. Horn and Charles R. Johnson. Matrix analysis. Cambridge Univ Pr, 1990.
Sham Kakade and Ambuj Tewari. Learning theory lecture notes, 2008. http://ttic.uchicago.edu/~tewari/lectures/lecture15.pdf.

Andreas Krause and Volkan Cevher. Submodular dictionary selection for sparse representation. In ICML, 2010.

Kenneth Kreutz-Delgado, Joseph F. Murray, Bhaskar D. Rao, Kjersti Engan, Te-Won Lee, and Terrance J. Sejnowski. Dictionary learning algorithms for sparse representation. Neural computation, 15(2):349-396, 2003.

Honglak Lee, Alexis Battle, Rajat Raina, and Andrew Y. Ng. Efficient sparse coding algorithms. Advances in neural information processing systems, 19:801, 2007.
P. Lévy. Problemes concrets danalyse fonctionnelle. Bull. Amer. Math. Soc, 58:408-411, 1952. ISSN 1088-9485.

Michael S. Lewicki, Terrence J. Sejnowski, and Howard Hughes. Learning overcomplete representations. Neural Computation, 12:337-365, 1998.

Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online learning for matrix factorization and sparse coding. JMLR, 11:19-60, 2010.
A. Maurer and M. Pontil. K-Dimensional Coding Schemes in Hilbert Spaces. Available at: arXiv:1002.0832v1 [stat.ML], 2010.

Bruno A. Olshausen and David J. Fieldt. Sparse coding with an overcomplete basis set: a strategy employed by V1. Vision Research, 37:3311-3325, 1997.

Gabriel Peyré. Sparse modeling of textures. Journal of Mathematical Imaging and Vision, 34(1): 17-31, 2009.

Matan Protter and Michael Elad. Sparse and redundant representations and motion-estimation-free algorithm for video denoising. Wavelets XII. Proceedings of the SPIE, 6701:43, 2007.
J. Shawe-Taylor and N. Cristianini. Kernel methods for pattern analysis. Cambridge Univ Pr, 2004.
J. Shawe-Taylor, C.K.I. Williams, N. Cristianini, and J. Kandola. On the eigenspectrum of the Gram matrix and the generalization error of kernel-PCA. Information Theory, IEEE Transactions on, 51(7):2510-2522, 2005. ISSN 0018-9448.

Joel A. Tropp. Greed is good: Algorithmic results for sparse approximation. IEEE Trans. Inform. Theory, 50:2231-2242, 2004.

AW Van der Vaart and J.A. Wellner. Weak convergence and empirical processes. Springer Verlag, 1996.

John Wright, Allen Y. Yang, Arvind. Ganesh, S. Shankar Sastry, and Yi Ma. Robust face recognition via sparse representation. IEEE Transactions on Pattern Analysis and Machine Intelligence, pages 210-227, 2008.

Jianchao Yang, Kai Yu, Yihong Gong, and Thomas Huang. Linear spatial pyramid matching using sparse coding for image classification. In IEEE Conference on Computer Vision and Pattern Recognition, 2009.

## Appendix A: generalization with fast rates

In this appendix we justify some adaptations in Lemma 10 relative to its origins in Bartlett et al. (2005). Specifically, we assume only growth rates instead of combinatorial dimensions and give explicit constants for this case (no particular effort was made to make the constants tight).

Following are some concepts and general results needed to prove Lemma 10 beyond those introduced in the main body of paper.

Definition 19. Let $F$ be a subset of a vector space $X, x \in X$. The star shaped closure of $F$ around $x$ is

$$
\star(F, x)=\{\lambda f+(1-\lambda) x: f \in F \wedge \lambda \in[0,1]\} .
$$

Definition 20. A function $f: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$is called sub-root if it is non negative, non decreasing and if $r \mapsto f(r) / \sqrt{( } r)$ is non increasing for $r>0$.

Definition 21. Let $\left\{Z_{i}\right\}_{i=1}^{m} \cup\left\{\varepsilon_{i}\right\}_{i=1}^{m}$ be independent variables, where $\varepsilon_{i}$ are uniform over $\{-1,1\}$, and $Z_{i}$ are i.i.d. The empirical Rademacher average of $\mathcal{F}$ is

$$
\hat{R}_{m}(\mathcal{F})=\mathbb{E}\left[\left.\sup _{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^{n} \varepsilon_{i} f_{i}\left(Z_{i}\right) \right\rvert\, Z_{1}, \ldots, Z_{m}\right]
$$

Lemma 22. Let $\hat{R}_{m}(\mathcal{F})$ be the empirical Rademacher averages of $\mathcal{F}$ on a sample $\left\{Z_{i}\right\}_{i=1}^{m}$. We have

$$
\hat{R}_{m}(\mathcal{F}) \leq 12 \int_{0}^{\infty} \sqrt{\frac{\log N\left(\varepsilon, \mathcal{F}, L_{2}\left(\nu_{m}\right)\right)}{m}} d \varepsilon
$$

where $\nu_{m}=m^{-1} \sum_{i=1}^{m} \delta_{Z_{i}}$.
See Kakade and Tewari (2008) for a proof.
Lemma 23. For any $\gamma \geq e^{\frac{1}{2}}$ and $x \in[0,1]$, we have $\int_{0}^{x} \sqrt{\log (\gamma / \varepsilon)} d \varepsilon \leq 2 x \sqrt{\log (\gamma / x)}$.
Proof. The indefinite integral $\int_{0}^{x} \sqrt{\log \frac{\gamma}{\varepsilon}} d \varepsilon$ is $x \sqrt{\log \frac{\gamma}{x}}-\frac{\sqrt{\pi}}{2} \gamma \cdot \operatorname{erf}\left(\sqrt{\log \frac{\gamma}{x}}\right)$, where $\operatorname{erf}(t)=$ $2 / \sqrt{\pi} \int_{0}^{t} e^{-u^{2}} d u$. Then

$$
\begin{aligned}
\int_{0}^{x} \sqrt{\log \frac{\gamma}{\varepsilon}} d \varepsilon & =\left[x \sqrt{\log \frac{\gamma}{x}}-\frac{\sqrt{\pi}}{2} \gamma \operatorname{erf}\left(\sqrt{\log \frac{\gamma}{x}}\right)\right]_{0}^{x} \\
& =x \sqrt{\log \frac{\gamma}{x}}-\frac{\sqrt{\pi}}{2} \gamma \operatorname{erf}\left(\sqrt{\log \frac{\gamma}{x}}\right)-\lim _{x \rightarrow 0}\left(x \sqrt{\log \frac{\gamma}{x}}-\frac{\sqrt{\pi}}{2} \gamma \operatorname{erf}\left(\sqrt{\log \frac{\gamma}{x}}\right)\right) \\
& =x \sqrt{\log \frac{\gamma}{x}}-\frac{\sqrt{\pi}}{2} \gamma\left(\operatorname{erf}\left(\sqrt{\log \frac{\gamma}{x}}\right)-\operatorname{erf}(\infty)\right)
\end{aligned}
$$

The error function erf is related to the tail probability of a normal variable, also known as the $Q$ function. In particular, $Q(x)=\frac{1}{2}\left(1-\operatorname{erf}\left(\frac{x}{\sqrt{2}}\right)\right) \Longleftrightarrow 2 Q(x)=1-\operatorname{erf}\left(\frac{x}{\sqrt{2}}\right) \Longleftrightarrow$ $2 Q(\sqrt{2} x)=1-\operatorname{erf}(x)$. We thus substitute and then use the bound $Q(x)<e^{-x^{2} / 2} /(x \sqrt{2 \pi})$ :

$$
\begin{aligned}
x \sqrt{\log \frac{\gamma}{x}}-\frac{\sqrt{\pi}}{2} \gamma\left(\operatorname{erf}\left(\sqrt{\log \frac{\gamma}{x}}\right)-1\right) & =x \sqrt{\log \frac{\gamma}{x}}-\frac{\sqrt{\pi}}{2} \gamma\left(-2 Q\left(\sqrt{2 \log \frac{\gamma}{x}}\right)\right) \\
& =x \sqrt{\log \frac{\gamma}{x}}+\sqrt{\pi} \gamma Q\left(\sqrt{2 \log \frac{\gamma}{x}}\right)
\end{aligned}
$$

(Bound on $Q$ )

$$
\begin{aligned}
& <x \sqrt{\log \frac{\gamma}{x}}+\sqrt{\pi} \frac{1}{\sqrt{2 \log \frac{\gamma}{x}} \sqrt{2 \pi}} e^{-\frac{\left(\sqrt{2 \log \frac{\gamma}{x}}\right)^{2}}{2}} \\
& =x \sqrt{\log \frac{\gamma}{x}}+x \frac{1}{2 \sqrt{\log \frac{\gamma}{x}}} \\
& =x\left(\sqrt{\log \frac{\gamma}{x}}+\frac{1}{2 \sqrt{\log \frac{\gamma}{x}}}\right)
\end{aligned}
$$

By our assumptions, $\frac{\gamma}{x} \geq \gamma \geq e^{1 / 2} \Longleftrightarrow \sqrt{\log \frac{\gamma}{x}}>\sqrt{\frac{1}{2}} \Longleftrightarrow \frac{1}{2 \sqrt{\log \frac{\gamma}{x}}}<\sqrt{\frac{1}{2}}$, then $x\left(\sqrt{\log \frac{\gamma}{x}}+\frac{1}{2 \sqrt{\log \frac{\gamma}{x}}}\right) \leq 2 x \sqrt{\log \frac{\gamma}{x}}$, completing the proof.

We return to prove Lemma 10

Proof. The core of the proof is to define particular sub-root function, and show its fixed point decays as $1 / m$. We then apply Theorem 3.3 of Bartlett, Bousquet and Mendelson Bartlett et al. (2005) to this sub-root function to complete the proof.

We define the function

$$
\psi(r)=10 \mathbb{E} R_{m}\left\{f \in \star(\mathcal{F}, 0) \mid E f^{2} \leq r\right\}+\frac{11 \log m}{m}
$$

By Lemma 3.4 of Bartlett et al. (2005) (with $T f=E f^{2}$ and $\hat{f}=0$ ) and Lemma 3.2 Bartlett et al. (2005), this function is sub-root and thus has a unique fixed point, which we denote $r^{*}$, and $r<$ $r^{*} \Longleftrightarrow r<\psi(r)$.

To upper bound $r^{*}$ we first construct an upper bound on $\psi$, in which $E f^{2}$ is replaced by $E_{m} f^{2}$, valid for $r \geq r^{*}$. The expectation of this upper bound is controlled using an entropy integral.

We make two observations.

1. By Corollary 2.2 of Bartlett et al. (2005), with $b=1$, for $r>\psi(r)$ with probability at least $1-1 / m$,

$$
\left\{f \in \star(\mathcal{F}, 0): E f^{2} \leq r\right\} \subset\left\{f \in \star(F, 0): E_{m} f^{2} \leq 2 r\right\} .
$$

2. By assumption $(\forall f \in \mathcal{F})\|f\|_{L_{\infty}} \leq 1$ and this implies $R_{m}\left\{f \in \star(\mathcal{F}, 0): E f^{2} \leq r\right\} \leq 1$.

Combining the observations, we can bound

$$
\mathbb{E} R_{m}\left\{f \in \star(\mathcal{F}, 0): E f^{2} \leq r\right\} \leq \frac{1}{m}+\mathbb{E} R_{m}\left\{f \in \star(F, 0): E_{m} f^{2} \leq 2 r\right\}
$$

Then $\psi(r) \leq 10\left(\frac{1}{m}+\mathbb{E} R_{m}\left\{f \in \star(\mathcal{F}, 0): E_{m} f^{2} \leq 2 r\right\}\right)+\frac{11 \log (m)}{m}$, and in particular

$$
r^{*}=\psi\left(r^{*}\right) \leq 10\left(\frac{1}{m}+\mathbb{E} R_{m}\left\{f \in \star(\mathcal{F}, 0): E_{m} f^{2} \leq 2 r^{*}\right\}\right)+\frac{11 \log (m)}{m}
$$

We denote $\nu_{m}$ the empirical measure induced by the $m$ samples (whose expectation is $E_{m}$ ). Under the metric $L_{2}\left(\nu_{m}\right)$, the set $\left\{f \in \star(\mathcal{F}, 0): E_{m} f^{2} \leq 2 r\right\}$ is covered by a single ball of radius $\sqrt{2 r}$ around the zero function. Applying Lemma [22] we have

$$
\begin{aligned}
\hat{R}_{m}\left\{f \in \star(\mathcal{F}, 0): E_{m} f^{2} \leq 2 r\right\} & \leq 12 \int_{0}^{\infty} \sqrt{\frac{\log N\left(\varepsilon, \star(\mathcal{F}, 0), L_{2}\left(\nu_{m}\right)\right)}{m}} d \varepsilon \\
& =12 \int_{0}^{\sqrt{2 r}} \sqrt{\frac{\log N\left(\varepsilon, \star(\mathcal{F}, 0), L_{2}\left(\nu_{m}\right)\right)}{m}} d \varepsilon
\end{aligned}
$$

Since $(\forall f \in \mathcal{F})\|f\|_{L_{2}\left(\nu_{m}\right)} \leq 1$, an $\varepsilon$ cover of $\mathcal{F}$ can be converted into an $\varepsilon$ cover of $\star(\mathcal{F}, 0)$ by replacing each element by $1 / \varepsilon+1$ balls on the segment from it to 0 . Then $N\left(\varepsilon, \star(\mathcal{F}, 0), L_{2}\left(\nu_{m}\right)\right) \leq$
$2 N\left(\varepsilon, \mathcal{F}, L_{2}\left(\nu_{m}\right)\right) / \varepsilon(a)$. Using also the assumption on $C(b)$ and Lemma $23(c)$, we find

$$
\begin{aligned}
12 \int_{0}^{\sqrt{2 r}} \sqrt{\frac{\log N\left(\varepsilon, \star(\mathcal{F}, 0), L_{2}\left(\nu_{m}\right)\right)}{m}} d \varepsilon & \stackrel{(a)}{\leq} 12 \int_{0}^{\sqrt{2 r}} \sqrt{\frac{\log \left(\left(\frac{C}{\varepsilon}\right)^{d} \frac{2}{\varepsilon}\right)}{m}} d \varepsilon \\
& \stackrel{(b)}{\leq} 12 \sqrt{\frac{d+1}{m}} \int_{0}^{\sqrt{2 r}} \sqrt{\log \left(\frac{C}{\varepsilon}\right)} d \varepsilon \\
& \stackrel{(c)}{\leq} 24 \sqrt{\frac{2 r(d+1) \log \left(\frac{C}{\sqrt{2 r}}\right)}{m}} .
\end{aligned}
$$

Substituting into $\psi\left(r^{*}\right)$, we obtain that

$$
\begin{aligned}
r^{*} & \leq 10\left(\frac{1}{m}+24 \sqrt{\frac{2 r^{*}(d+1) \log \left(\frac{C}{\sqrt{2 r^{*}}}\right)}{m}}\right)+\frac{11 \log (m)}{m} \\
& =240 \sqrt{\frac{2 r^{*}(d+1) \log \left(\frac{C}{\left.\sqrt{2 r^{*}}\right)}\right.}{m}+\frac{10+11 \log (m)}{m}} .
\end{aligned}
$$

Let $\alpha>0$ be fixed. If $r^{*} \leq \alpha C^{2} / 2 m$, our first step is complete. If not, then

$$
r^{*}>\alpha C^{2} / 2 m \Longleftrightarrow \sqrt{m / \alpha}>C / \sqrt{2 r^{*}},
$$

and then

$$
\begin{aligned}
r^{*} & \leq 240 \sqrt{\frac{2 r^{*}(d+1) \log \left(\sqrt{\frac{m}{\alpha}}\right)}{m}+\frac{10+11 \log (m)}{m}} \\
& =240 \sqrt{\frac{r^{*}(d+1) \log \left(\frac{m}{\alpha}\right)}{m}}+\frac{10+11 \log (m)}{m} \\
& \leq 2 \max \left\{240 \sqrt{\frac{r^{*}(d+1) \log \left(\frac{m}{\alpha}\right)}{m}}, \frac{10+11 \log (m)}{m}\right\} .
\end{aligned}
$$

Then either $r^{*} \leq(20+22 \log (m)) / m$ (and the first step is complete), or

$$
r^{*} \leq 480 \sqrt{\left(r^{*}(d+1) \log (m / \alpha)\right) / m} \Longleftrightarrow r^{*} \leq(480)^{2}((d+1) \log (m / \alpha)) / m
$$

and again we are done. We conclude that

$$
r^{*} \leq \max \left\{\frac{\alpha C^{2}}{2 m},(480)^{2} \frac{(d+1) \log \left(\frac{m}{\alpha}\right)}{m}, \frac{20+22 \log (m)}{m}\right\}
$$

Having proved $r^{*}$ decays approximately as $1 / m$, we apply Theorem 3.3 of Bartlett et al. (2005), with $a=0 ; b=1 ; B=1 ; T f=E f^{2}$. By definition of $T$, it is clear that

$$
\begin{aligned}
\psi(r) & =10 \mathbb{E} R_{m}\left\{f \in \star(\mathcal{F}, 0) \mid E f^{2} \leq r\right\}+\frac{11 \log m}{m} \\
& \geq \mathbb{E} R_{m}\left\{f \in \star(\mathcal{F}, 0) \mid E f^{2} \leq r\right\} \\
& =\mathbb{E} R_{m}\{f \in \star(\mathcal{F}, 0) \mid T f \leq r\}
\end{aligned}
$$

holds, then we can use part 2 of Theorem 3.3 of Bartlett et al. (2005), which allows the conclusion that for all $f \in \mathcal{F}, E f \leq \frac{K}{K-1} E_{m} f+6 K r^{*}+\frac{11 x+5 K}{m}$.

