

Compressed sensing and best k -term approximation ^{*}

Albert Cohen, Wolfgang Dahmen, and Ronald DeVore

March 3, 2008

Abstract

Compressed sensing is a new concept in signal processing where one seeks to minimize the number of measurements to be taken from signals while still retaining the information necessary to approximate them well. The ideas have their origins in certain abstract results from functional analysis and approximation theory by Kashin [15] but were recently brought into the forefront by the work of Candès, Romberg and Tao [5, 3, 4] and Donoho [9] who constructed concrete algorithms and showed their promise in application. There remain several fundamental questions on both the theoretical and practical side of compressed sensing. This paper is primarily concerned about one of these theoretical issues revolving around just how well compressed sensing can approximate a given signal from a given budget of fixed linear measurements, as compared to adaptive linear measurements. More precisely, we consider discrete signals $x \in \mathbb{R}^N$, allocate $n < N$ linear measurements of x , and we describe the range of k for which these measurements encode enough information to recover x in the sense of ℓ_p to the accuracy of best k -term approximation. We also consider the problem of having such accuracy only with high probability.

1 Introduction

The typical paradigm for obtaining a compressed version of a discrete signal represented by a vector $x \in \mathbb{R}^N$ is to choose an appropriate basis, compute the coefficients of x in this basis, and then retain only the k largest of these with $k < N$. If we are interested in a bit stream representation, we also need in addition to quantize these k coefficients.

Assuming, without loss of generality, that x already represents the coefficients of the signal in the appropriate basis, this means that we pick an approximation to x in the set Σ_k of k -sparse vectors

$$\Sigma_k := \{x \in \mathbb{R}^N : \#\text{supp}(x) \leq k\}, \quad (1.1)$$

^{*}This research was supported by the Office of Naval Research Contracts ONR-N00014-03-1-0051, ONR/DEPSCoR N00014-03-1-0675 and ONR/DEPSCoR N00014-00-1-0470; DARPA Grant N66001-06-1-2001; the Army Research Office Contract DAAD 19-02-1-0028; the AFOSR Contract UF/USAF F49620-03-1-0381; the NSF contracts DMS-0221642 and DMS-0200187; the French-German PROCOPE contract 11418YB; and by the European Community's Human Potential Programme under contract HPRN-CT-202-00286, (BREAKING COMPLEXITY)

where $\text{supp}(x)$ is the support of x , i.e., the set of i for which $x_i \neq 0$, and $\#A$ is the number of elements in the set A . The best performance that we can achieve by such an approximation process in some given norm $\|\cdot\|_X$ of interest is described by the *best k -term approximation error*

$$\sigma_k(x)_X := \inf_{z \in \Sigma_k} \|x - z\|_X. \quad (1.2)$$

This approximation process should be considered as *adaptive* since the indices of those coefficients which are retained vary from one signal to another. On the other hand, this procedure is stressed on the front end by the need to first compute all of the basis coefficients. The view expressed by Candès, Romberg, and Tao [5, 3, 4] and Donoho [9] is that since we retain only a few of these coefficients in the end, perhaps it is possible to actually compute only a few *non-adaptive* linear measurements in the first place and still retain the necessary information about x in order to build a compressed representation.

These ideas have given rise to a very lively area of research called *compressed sensing* which poses many intriguing questions, of both a theoretical and practical flavor. The present paper is an excursion into this area, focusing our interest on the question of just how well compressed sensing can perform in comparison to best k -term approximation.

To formulate the problem, we are given a budget of n questions we can ask about x . These questions are required to take the form of asking for the values $\lambda_1(x), \dots, \lambda_n(x)$ where the λ_j are fixed linear functionals. The information we gather about x can therefore be described by

$$y = \Phi x, \quad (1.3)$$

where Φ is an $n \times N$ matrix called the *encoder* and $y \in \mathbb{R}^n$ is the *information vector*. The rows of Φ are representations of the linear functionals λ_j , $j = 1, \dots, n$.

To extract the information that y holds about x , we use a *decoder* Δ which is a mapping from $\mathbb{R}^n \rightarrow \mathbb{R}^N$. We emphasize that Δ is not required to be linear. Thus, $\Delta(y) = \Delta(\Phi x)$ is our approximation to x from the information we have retained. We shall denote by $\mathcal{A}_{n,N}$ the set of all encoder-decoder pairs (Φ, Δ) with Φ an $n \times N$ matrix.

There are two common ways to evaluate the performance of an encoding-decoding pair $(\Phi, \Delta) \in \mathcal{A}_{n,N}$. The first is to ask for the largest value of k such that the encoding-decoding is exact for all k -sparse vectors, i.e.

$$x \in \Sigma_k \Rightarrow \Delta(\Phi x) = x. \quad (1.4)$$

It is easy to see (see §2) that given n, N , there are $(\Delta, \Phi) \in \mathcal{A}_{n,N}$ such that (1.4) holds for all $k \leq n/2$. Or put in another way, given k , we can achieve exact recovery on Σ_k whenever $n \geq 2k$. Unfortunately such encoder/decoder pairs are not numerically friendly as is explained in §2.

Generally speaking, our signal will not be in Σ_k with k small but may be approximated well by the elements in Σ_k . Therefore, we would like our algorithms to perform well in this case as well. One way of comparing compressed sensing with best k -term approximation is to consider their respective performance on a specific class of vectors $K \subset \mathbb{R}^N$. For such a class we can define on the one hand

$$\sigma_k(K)_X := \sup_{x \in K} \sigma_k(x)_X, \quad (1.5)$$

and

$$E_n(K)_X := \inf_{(\Phi, \Delta) \in \mathcal{A}_{n,N}} \sup_{x \in K} \|x - \Delta(\Phi x)\|_X \quad (1.6)$$

which describe respectively the performance of the two methods over this class. We are now interested in the largest value of k such that $E_n(K)_X \leq C_0 \sigma_k(K)_X$ for a constant C_0 independent of the parameters k, n, N . Results of this type were established already in the 1970's under the umbrella of what is called n -widths. The deepest results of this type were given by Kashin [15] with later improvements by Garneev and Gluskin [10, 14]. We recall this well-known story briefly in §2.

The results on n -widths referred to above give matching upper and lower estimates for $E_n(K)_X$ in the case that K is a typical sparsity class such as a ball in ℓ_p^N where

$$\|x\|_{\ell_p} := \|x\|_{\ell_p^N} := \begin{cases} \left(\sum_{i=1}^N |x_i|^p \right)^{1/p}, & 0 < p < \infty, \\ \max_{j=1, \dots, N} |x_j|, & p = \infty. \end{cases} \quad (1.7)$$

This in turn determines the largest range of k for which we can obtain comparisons of the form $E_n(K)_X \leq C_0 \sigma_k(K)_X$. One such result is the following: for $K = U(\ell_1^N)$, $X = \ell_2^N$, one has

$$E_k(U(\ell_1^N))_{\ell_2^N} \leq C_0 \sigma_k(U(\ell_1^N))_{\ell_2^N} \quad (1.8)$$

whenever

$$k \leq c_0 n / \log(N/n) \quad (1.9)$$

with absolute constants C_0, c_0 .

The decoders used in proving these theoretical bounds are far from being practical or numerically implementable. One of the remarkable achievements of the recent work of Candès, Romberg and Tao [3] and Donoho [9] is to give probabilistic constructions of matrices Φ which provide these bounds where the decoding can be done by solving the ℓ_1 minimization problem

$$\Delta(y) := \underset{\Phi z = y}{\operatorname{Argmin}} \|z\|_{\ell_1}. \quad (1.10)$$

The above results on approximation of classes is governed by the worst elements in the class. It is a more subtle problem to obtain estimates that depend on the individual characteristics of the target vector x . The main contribution of the present paper is to study a stronger way to compare the performance of k -term approximation in a compressed sensing algorithm. Namely, we address the following question:

For a given norm $\|\cdot\|_X$ and $k < N$, what is the minimal value of n , for which there exists a pair $(\Phi, \Delta) \in \mathcal{A}_{n,N}$ such that

$$\|x - \Delta(\Phi x)\|_X \leq C_0 \sigma_k(x)_X, \quad (1.11)$$

for all $x \in \mathbb{R}^N$, with C_0 a constant independent of k and N ?

If a result of the form (1.11) has been established then one can derive a result for a class K by simply taking the supremum over all $x \in K$. However, results on classes are less precise and informative than (1.11).

We shall say a pair $(\Phi, \Delta) \in \mathcal{A}_{n,N}$ satisfying (1.11) is *instance optimal* of order k with constant C_0 for the space X . In particular, we want to understand under what circumstances the minimal value of n is roughly of the same order as k , similar to (1.9). We shall see that the answer to this question strongly depends on the norm X under consideration.

The approximation accuracy of a compressed sensing matrix is determined by the null space

$$\mathcal{N} = \mathcal{N}(\Phi) := \{x \in \mathbb{R}^N : \Phi x = 0\}. \quad (1.12)$$

The importance of \mathcal{N} is that if we observe $y = \Phi x$ without any a-priori information on x , the set of z such that $\Phi z = y$ is given by the affine space

$$\mathcal{F}(y) := x + \mathcal{N}. \quad (1.13)$$

We bring out the importance of the null space in §3 where we formulate a property of the null space which is necessary and sufficient for Φ to have a decoder Δ for which the instance optimality (1.11) holds.

We apply this property in §4 to the case $X = \ell_1$. In this case, we show the minimal number of measurements n which ensures (1.11) is of the same order as k up to a logarithmic factor. In that sense, compressed sensing performs almost as good as best k -term approximation. We also show that, similar to the work of Candès, Romberg, and Tao this is achieved with the decoder Δ defined by ℓ_1 minimization. We should mention that our results in this section are essentially contained in the work of Candès, Romberg, and Tao [5, 6, 4] and we build on their ideas.

We next treat the case $X = \ell_2$ in §5. In this case, the situation is much less in favor of compressed sensing, since the minimal number of measurements n which ensures (1.11) is now of the same order as N .

In §6, we consider an important variant of the ℓ_2 case where we ask for ℓ_2 instance optimality in the sense of probability. Here, rather than requiring that (1.11) holds for all $x \in \mathbb{R}^N$, we ask only that for each given x it holds with high probability. We shall see that in the case $X = \ell_2$ the minimal number of measurements n for such results is dramatically reduced, down to the the order given by condition (1.9). Moreover, we show that standard constructions of random matrices such as Gaussian and Bernouli ensembles achieve this performance.

The striking contrast between the results of §5 and §6 show that the probabilistic setting plays a crucial role in ℓ^2 instance optimality. Similar results in the sense of probability have been obtained earlier in a series of paper [8, 11, 12, 13] that reflect the theoretical computer science approach to compressed sensing, also known as data sketching. A comparison with our results is in order.

First, the instance optimality bounds obtained in these papers are quantitatively more precise than ours, since they have the general form

$$\|x - \Delta(\Phi x)\|_{\ell^2} \leq (1 + \epsilon)\sigma_k(x)_{\ell^2}, \quad (1.14)$$

where $\epsilon > 0$ can be made arbitrarily small, at the expense of raising n , while in most of our results the constant C_0 in (1.11) cannot get arbitrarily close to 1. On the other hand, for a fixed $\epsilon > 0$, the ratio between n and k is generally not as good as in (1.9): for

instance the decoders proposed in [8] and [12] respectively use $n = \mathcal{O}(\frac{k}{\epsilon} \log(N)^{5/2})$ and $n = \mathcal{O}(\frac{k}{\epsilon^3} \log(N))$ samples in order to achieve (1.14).

Secondly, the type of encoding matrices which are proposed in these papers are of fairly different nature than those which are considered in §6, and our analysis actually *does not apply* to these matrices. Let us mention that one specific interest of the Gaussian matrices which are considered in the present paper is that they give rise to an encoding which is “robust” with respect to a change of the basis in which the signal is sparse, since the product of such a Φ by an $N \times N$ unitary matrix U results in a $n \times N$ matrix $\tilde{\Phi}$ with the same probability law as Φ .

Finally, one of the significant achievements in [8, 11, 12, 13] is the derivation of practical decoding algorithms of polynomial complexity in k (up to logarithmic factors), therefore typically faster than solving the ℓ^1 minimization problem, while we do not propose any such algorithm in the present paper.

Generally speaking, an important issue in compressed sensing is the practical implementation of the decoder Δ by a fast algorithm. While being aware of this fact, the main goal of the present paper is to understand the theoretical limits of compressed sensing in comparison to nonlinear approximation. Therefore the main question that we address is “*how many measurements do we need so that some decoder recovers x up to some prescribed tolerance*”, rather than “*what is the fastest algorithm which allows to recover x from these measurements up to the same tolerance*”.

The last sections of the paper are devoted to additional results which complete the theory. In order to limit the size of the paper we only give a sketch of the proofs in those sections. The case $X = \ell_p$ for $1 < p < 2$ is treated in §7, and in §8 we discuss another type of estimate that we refer to as *mixed-norm instance optimality*. Here the estimate (1.11) is replaced by an estimate of the type

$$\|x - \Delta(\Phi x)\|_X \leq C_0 k^{-s} \sigma_k(x)_Y, \quad (1.15)$$

where Y differs from X and $s > 0$ is some relevant exponent. This type of estimate was introduced in [4] in the particular case $X = \ell_2$ and $Y = \ell_1$. We give examples in the case $X = \ell_p$ and $Y = \ell_q$ in which mixed-norm estimates allow us to recover better approximation estimates for compressed sensing than (1.11).

2 Performance over classes

We begin by recalling some well-known results concerning best k -term approximation which we shall use in the course of this paper. Given a sequence norm $\|\cdot\|_X$ on \mathbb{R}^N and a positive integer $r > 0$, we define the approximation class \mathcal{A}^r by means of

$$\|x\|_{\mathcal{A}^r(X)} := \max_{1 \leq k \leq N} k^r \sigma_k(x)_X. \quad (2.1)$$

Notice that since we are in a finite dimensional space \mathbb{R}^N , this (quasi-)norm will be finite for all $x \in \mathbb{R}^N$.

A simple, yet fundamental chapter in k -term approximation is to connect the approximation norm in (2.1) with traditional sequence norms. For this, we define for any

$0 < q < \infty$, the weak- ℓ_q norm as

$$\|x\|_{w\ell_q}^q := \sup_{\epsilon > 0} \epsilon^q \#\{i ; |x_i| > \epsilon\}. \quad (2.2)$$

Again, for any $x \in \mathbb{R}^N$ all of these norms are finite.

If we fix the ℓ_p norm in which approximation error is to be measured, then for any $x \in \mathbb{R}^N$, we have for $q := (r + 1/p)^{-1}$,

$$B_0 \|x\|_{w\ell_q} \leq \|x\|_{\mathcal{A}^r} \leq B_1 r^{-1/p} \|x\|_{w\ell_q}, \quad x \in \mathbb{R}^N, \quad (2.3)$$

for two absolute constants $B_0, B_1 > 0$. Notice that the constants in these inequalities do not depend on N . Therefore, $x \in \mathcal{A}^r$ is equivalent to $x \in w\ell_q$ with equivalent norms.

Since the ℓ_q norm is larger than the weak ℓ_q norm, we can replace the weak ℓ_q norm by the ℓ_q norm in the right inequality of (2.3). However, the constant can be improved via a direct argument. Namely, if $1/q = r + 1/p$, then for any $x \in \mathbb{R}^N$,

$$\sigma_k(x)_{\ell_p} \leq \|x\|_{\ell_q} k^{-r}, \quad k = 1, 2, \dots, N. \quad (2.4)$$

To prove this, take Λ as the set of indices corresponding to the k largest entries in x . If ϵ is the size of the smallest entry in Λ , then $\epsilon \leq \|x\|_{w\ell_q} k^{-1/q} \leq \|x\|_{\ell_q} k^{-1/q}$ and therefore

$$\sigma_k(x)_{\ell_p}^p = \sum_{i \notin \Lambda} |x_i|^p \leq \epsilon^{p-q} \sum_{i \notin \Lambda} |x_i|^q \leq k^{-\frac{p-q}{q}} \|x\|_{\ell_q}^{p-q} \|x\|_{\ell_q}^q, \quad (2.5)$$

so that (2.4) follows.

From this, we see that if we consider the class $K = U(\ell_q^N)$, we have

$$\sigma_k(K)_{\ell_p} \leq k^{-r}, \quad (2.6)$$

with $r = 1/q - 1/p$. On the other hand, taking $x \in K$ such that $x_i = (2k)^{-1/q}$ for $2k$ indices and 0 otherwise, we find that

$$\sigma_k(x)_{\ell_p} = [k(2k)^{-p/q}]^{1/p} = 2^{-1/q} k^{-r}, \quad (2.7)$$

so that $\sigma_k(K)_X$ can be framed by

$$2^{-1/q} k^{-r} \leq \sigma_k(K)_{\ell_p} \leq k^{-r}. \quad (2.8)$$

We next turn to the performance of compressed sensing over *classes of vectors*, by studying the quantity $E_n(K)_X$ defined by (1.6). As we have mentioned, the optimal performance of sensing algorithms is closely connected to the concept of Gelfand width which are in some sense dual to the perhaps better known Kolmogorov widths. If K is a compact set in X , and n is a positive integer, then the Gelfand width of K and of order n is by definition

$$d^n(K)_X := \inf_Y \sup\{\|x\|_X ; x \in K \cap Y\} \quad (2.9)$$

where the infimum is taken over all subspaces Y of X of codimension less or equal to n . This quantity is equivalent to $E_n(K)_X$ according to the following well known result.

Lemma 2.1 *Let $K \subset \mathbb{R}^N$ be any set for which $K = -K$ and for which there is a $C_0 > 0$ such that $K + K \subset C_0K$. If $X \subset \mathbb{R}^N$ is any normed space, then*

$$d^n(K)_X \leq E_n(K)_X \leq C_0 d^n(K)_X, \quad 1 \leq n \leq N. \quad (2.10)$$

Proof: We give a proof for completeness of this paper. We first remark that the null space $Y = \mathcal{N}$ of Φ is of codimension less or equal to n . Conversely, given any space $Y \subset \mathbb{R}^N$ of codimension n , we can associate its orthogonal complement Y^\perp which is of dimension n and the $n \times N$ matrix Φ whose rows are formed by any basis for Y^\perp . Through this identification, we see that

$$d^n(K)_X = \inf_{\Phi} \sup \{ \|\eta\|_X : \eta \in \mathcal{N} \cap K \}, \quad (2.11)$$

where the infimum is taken over all $n \times N$ matrices Φ .

Now, if (Φ, Δ) is any encoder-decoder pair and $z = \Delta(0)$, then for any $\eta \in \mathcal{N}$, we also have $-\eta \in \mathcal{N}$. It follows that either $\|\eta - z\|_X \geq \|\eta\|_X$ or $\|-\eta - z\|_X \geq \|\eta\|_X$. Since $K = -K$ we conclude that

$$d^n(K)_X \leq \sup_{\eta \in \mathcal{N} \cap K} \|\eta - \Delta(\Phi\eta)\|_X. \quad (2.12)$$

Taking an infimum over all encoder-decoder pairs in $\mathcal{A}_{n,N}$, we obtain the left inequality in (2.10).

To prove the right inequality, we choose an optimal Y for $d^n(K)_X$ and use the matrix Φ associated to Y (i.e., the rows of Φ are a basis for Y^\perp). We define a decoder Δ for Φ as follows. Given y in the range of Φ , we recall that $\mathcal{F}(y)$ is the set of x such that $\Phi x = y$. If $\mathcal{F}(y) \cap K \neq \emptyset$, we take any $\bar{x}(y) \in \mathcal{F}(y) \cap K$ and define $\Delta(y) := \bar{x}(y)$. When $\mathcal{F}(y) \cap K = \emptyset$, we define $\Delta(y)$ as any element from $\mathcal{F}(y)$. This gives

$$E_n(K)_X \leq \sup_{x, x' \in \mathcal{F}(y) \cap K} \|x - x'\|_X \leq \sup_{\eta \in C_0[K \cap \mathcal{N}]} \|\eta\|_X \leq C_0 d^n(K)_X, \quad (2.13)$$

where we have used the fact that $x - x' \in \mathcal{N}$ and $x - x' \in C_0K$ by our assumptions on K . This proves the right inequality in (2.10). \square

The orders of the Gelfand widths of ℓ_q balls in ℓ_p are known except perhaps for the case $q = 1, p = \infty$. For the range of p, q that is relevant here even the constants are known. We recall the following results of Gluskin, Garneev and Kashin which can be found in [14, 10, 15], see also [16]. For $K = U(\ell_q^N)$, we have

$$C_1 \Psi(n, N, q, p) \leq d^n(K)_{\ell_p} \leq C_2 \Psi(n, N, q, p), \quad (2.14)$$

where C_1, C_2 only depend on p and q , and where

$$\Psi(n, N, q, p) := [\min(1, N^{1-1/q} n^{-1/2})]^{1/q-1/2}, \quad 1 \leq n \leq N, \quad 1 < q < p \leq 2, \quad (2.15)$$

and

$$\Psi(n, N, 1, 2) := \min \left\{ 1, \sqrt{\frac{\log(N/n)}{n}} \right\}. \quad (2.16)$$

Since $K = U(\ell_q^N)$ obviously satisfies the assumptions of Lemma 2.1 with $C_0 = 2$, we also have

$$C_1\Psi(n, N, q, p) \leq E_n(K)_{\ell_p} \leq 2C_2\Psi(n, N, q, p). \quad (2.17)$$

From (2.14), (2.16) and (2.10) we deduce indeed the announced fact that $E_n(U(\ell_1^N))_{\ell_2} \leq C_0\sigma_k(U(\ell_1^N))_{\ell_2}$ can only hold when k and the necessary number of measurements n are interrelated by (1.9). The possible range of k for which even instance optimality could hold is therefore also limited by (1.9), a relation that will turn up frequently in what follows.

3 Instance optimality and the null space of Φ

We now turn to the main question addressed in this paper, namely the study of instance optimality as expressed by (1.11). In this section, we shall see that (1.11) can be reformulated as a property of the null space \mathcal{N} of Φ . As it was already remarked in the proof of Lemma 2.1, this null space has codimension not larger than n .

We shall also need to consider sections of Φ obtained by keeping some of its columns: for $T \subset \{1, \dots, N\}$, we denote by Φ_T the $n \times \#T$ matrix formed from the columns of Φ with indices in T . Similarly we shall have to deal with restrictions x_T of vectors $x \in \mathbb{R}^N$ to sets T . However, it will be convenient to view such restrictions still as elements of \mathbb{R}^N , i.e. x_T agrees with x on T and has all components equal to zero whose indices do not belong to T .

We begin by studying under what circumstances the measurements $y = \Phi x$ uniquely determines each k -sparse vector $x \in \Sigma_k$. This is expressed by the following trivial lemma.

Lemma 3.1 *If Φ is any $n \times N$ matrix and $2k \leq n$, then the following are equivalent:*

- (i) *There is a decoder Δ such that $\Delta(\Phi x) = x$, for all $x \in \Sigma_k$,*
- (ii) $\Sigma_{2k} \cap \mathcal{N} = \{0\}$,
- (iii) *For any set T with $\#T = 2k$, the matrix Φ_T has rank $2k$.*
- (iv) *The symmetric non-negative matrix $\Phi_T^t \Phi_T$ is invertible, i.e. positive definite.*

Proof: The equivalence of (ii), (iii), (iv) is linear algebra.

(i) \Rightarrow (ii): Suppose (i) holds and $x \in \Sigma_{2k} \cap \mathcal{N}$. We can write $x = x_0 - x_1$ where both $x_0, x_1 \in \Sigma_k$. Since $\Phi x_0 = \Phi x_1$, we have, by (i), that $x_0 = x_1$ and hence $x = x_0 - x_1 = 0$.

(ii) \Rightarrow (i): Given any $y \in \mathbb{R}^n$, we define $\Delta(y)$ to be any element in $\mathcal{F}(y)$ with smallest support. Now, if $x_1, x_2 \in \Sigma_k$ with $\Phi x_1 = \Phi x_2$, then $x_1 - x_2 \in \mathcal{N} \cap \Sigma_{2k}$. From (ii), this means that $x_1 = x_2$. Hence, if $x \in \Sigma_k$ then $\Delta(\Phi x) = x$ as desired. \square

The properties discussed in Lemma 3.1 are algebraic properties of Φ . If N, k are fixed, the question arises as to how large do we need to make n so that there is a matrix Φ having the properties of the Lemma. It is easy to see that we can take $n = 2k$. Indeed, for any k and $N \geq 2k$, we can find a set Λ_N of N vectors in \mathbb{R}^{2k} such that any $2k$ of them are linearly independent. For example if $0 < x_1 < x_2 < \dots < x_N$ then the matrix whose (i, j) entry is x_j^{i-1} has the properties of Lemma 3.1. Its $2k \times 2k$ minors are Vandermonde matrices which are well known to be non-singular. Unfortunately, such

matrices are poorly conditioned when N is large and the process of recovering $x \in \Sigma_k$ from $y = \Phi x$ is therefore numerically unstable.

Stable recovery procedures have been proposed by Candès-Romberg-Tao and Donoho under stronger conditions on Φ . We shall make heavy use in this paper of the following property introduced by Candès and Tao. We say that Φ satisfies the *restricted isometry property* (RIP) of order k if there is a $0 < \delta_k < 1$ such that

$$(1 - \delta_k)\|z\|_{\ell_2} \leq \|\Phi_T z\|_{\ell_2} \leq (1 + \delta_k)\|z\|_{\ell_2}, \quad z \in \mathbb{R}^N, \quad (3.1)$$

holds for all T of cardinality k^1 . The RIP condition is equivalent to saying that the symmetric matrix $\Phi_T^t \Phi_T$ is positive definite with eigenvalues in $[(1 - \delta_k)^2, (1 + \delta_k)^2]$. Note that RIP of order k always implies RIP of order $l \leq k$. Note also that RIP of order $2k$ guarantees that the properties of Lemma 3.1 hold.

Candès and Tao have shown that any matrix Φ which satisfies the RIP property for k and sufficiently small δ_k will extract enough information about x to approximate it well and moreover the decoding can be done by ℓ_1 minimization. The key question then is given a fixed n, N , how large can we take k and still have matrices which satisfy RIP for k . It was shown by Candès and Tao [5], as well as Donoho [9], that certain families of random matrices will, with high probability, satisfy RIP of order k with $\delta_k \leq \delta < 1$ for some prescribed δ independent of N provided $k \leq c_0 n / \log(N/k)$. Here c_0 is a constant which when made small will make δ_k small as well. It should be stressed that all available constructions of such matrices (so far) involve random variables. For instance, as we shall recall in more detail in §6, the entries of Φ can be picked as i.i.d. Gaussian or Bernoulli variables with proper normalization.

We turn to the question of whether y contains enough information to approximate x to accuracy $\sigma_k(x)$ as expressed by (1.11). The following theorem shows that this can be understood through the study of the null space \mathcal{N} of Φ .

Theorem 3.2 *Given an $n \times N$ matrix Φ , a norm $\|\cdot\|_X$ and a value of k , then a sufficient condition that there exists a decoder Δ such that (1.11) holds with constant C_0 is that*

$$\|\eta\|_X \leq \frac{C_0}{2} \sigma_{2k}(\eta)_X, \quad \eta \in \mathcal{N}. \quad (3.2)$$

A necessary condition is that

$$\|\eta\|_X \leq C_0 \sigma_{2k}(\eta)_X, \quad \eta \in \mathcal{N}. \quad (3.3)$$

Proof: To prove the sufficiency of (3.2), we will define a decoder Δ for Φ as follows. Given any $y \in \mathbb{R}^N$, we consider the set $\mathcal{F}(y)$ and choose

$$\Delta(y) := \underset{z \in \mathcal{F}(y)}{\text{Argmin}} \sigma_k(z)_X. \quad (3.4)$$

¹The RIP condition could be replaced by the assumption that $C_0\|z\|_{\ell_2} \leq \|\Phi_T z\|_{\ell_2} \leq C_1\|z\|_{\ell_2}$ holds for all $\#(T) = k$, with absolute constants C_0, C_1 in all that follows. However, this latter condition is equivalent to having a rescaled matrix $\alpha\Phi$ satisfy RIP for some α and the rescaled matrix extracts exactly the same information from a vector x as Φ does.

We shall prove that for all $x \in \mathbb{R}^N$

$$\|x - \Delta(\Phi x)\|_X \leq C_0 \sigma_k(x)_X. \quad (3.5)$$

Indeed, $\eta := x - \Delta(\Phi x)$ is in \mathcal{N} and hence by (3.2), we have

$$\begin{aligned} \|x - \Delta(\Phi x)\|_X &\leq (C_0/2) \sigma_{2k}(x - \Delta(\Phi x))_X \\ &\leq (C_0/2) (\sigma_k(x)_X + \sigma_k(\Delta(\Phi x))_X) \\ &\leq C_0 \sigma_k(x)_X, \end{aligned}$$

where the second inequality uses the fact that $\sigma_{2k}(x+z)_X \leq \sigma_k(x)_X + \sigma_k(z)_X$ and the last inequality uses the fact that $\Delta(\Phi x)$ minimizes $\sigma_k(z)$ over $\mathcal{F}(y)$.

To prove the necessity of (3.3), let Δ be any decoder for which (1.11) holds. Let η be any element in $\mathcal{N} = \mathcal{N}(\Phi)$ and let η_0 be the best $2k$ -term approximation of η in X . Let $\eta_0 = \eta_1 + \eta_2$ be any splitting of η_0 into two vectors of support size k , we can write

$$\eta = \eta_1 + \eta_2 + \eta_3, \quad (3.6)$$

with $\eta_3 = \eta - \eta_0$. Since $-\eta_1 \in \Sigma_k$ we have by (1.11) that $-\eta_1 = \Delta(\Phi(-\eta_1))$, but since $\eta \in \mathcal{N}$, we also have $-\Phi\eta_1 = \Phi(\eta_2 + \eta_3)$ so that $-\eta_1 = \Delta(\Phi(\eta_2 + \eta_3))$. Using again (1.11) we derive

$$\begin{aligned} \|\eta\|_X &= \|\eta_2 + \eta_3 - \Delta(\Phi(\eta_2 + \eta_3))\|_X \leq C_0 \sigma_k(\eta_2 + \eta_3) \\ &\leq C_0 \|\eta_3\|_X = C_0 \sigma_{2k}(\eta), \end{aligned}$$

which is (3.3). □

When X is an ℓ_p space, the best k term approximation is obtained by leaving the k largest components of x unchanged and setting all the others to 0. Therefore the property

$$\|\eta\|_X \leq C \sigma_k(\eta)_X, \quad (3.7)$$

can be reformulated by saying that

$$\|\eta\|_X \leq C \|\eta_{T^c}\|_X, \quad (3.8)$$

holds for all $T \subset \{1, \dots, N\}$ such that $\#T \leq k$, where T^c is the complement set of T in $\{1, \dots, N\}$. In going further, we shall say that Φ has the *null space property* in X of order k with constant C if (3.8) holds for all $\eta \in \mathcal{N}$ and $\#T \leq k$. Thus, we have

Corollary 3.3 *Suppose that X is an ℓ_p^N space, $k > 0$ an integer and Φ an encoding matrix. If Φ has the null space property (3.8) in X of order $2k$ with constant $C_0/2$, then there exists a decoder Δ so that (Φ, Δ) satisfies (1.11) with constant C_0 . Conversely, the validity of (1.11) for some decoder Δ implies that Φ has the null space property (3.8) in X of order $2k$ with constant C_0 .*

In the next two sections, we shall use this corollary in order to study instance optimality in the case where the X norm is ℓ_1 and ℓ_2 respectively.

4 The case $X = \ell_1$

In this section, we shall study the null space property (3.8) in the case where $X = \ell_1$. We shall make use of the restricted isometry property (3.1) introduced by Candès and Tao. We begin with the following lemma whose proof is inspired by results in [4].

Lemma 4.1 *Let $a = \ell/k$, $b = \ell'/k$ with $\ell, \ell' \geq k$ integers. If Φ is any matrix which satisfies the RIP of order $(a+b)k$ with $\delta = \delta_{(a+b)k} < 1$. Then Φ satisfies the null space property in ℓ_1 of order ak with constant $C_0 = 1 + \frac{\sqrt{a}(1+\delta)}{\sqrt{b}(1-\delta)}$.*

Proof: It is enough to prove (3.8) in the case when T is the set of indices of the largest ak entries of η . Let $T_0 = T$, T_1 denote the set of indices of the next bk largest entries of η , T_2 the next bk largest, and so on. The last set T_s defined this way may have less than bk elements.

We define $\eta_0 := \eta_{T_0} + \eta_{T_1}$. Since $\eta \in \mathcal{N}$, we have $\Phi\eta_0 = -\Phi(\eta_{T_2} + \dots + \eta_{T_s})$, so that

$$\begin{aligned} \|\eta_T\|_{\ell_2} &\leq \|\eta_0\|_{\ell_2} \leq (1-\delta)^{-1}\|\Phi\eta_0\|_{\ell_2} = (1-\delta)^{-1}\|\Phi(\eta_{T_2} + \dots + \eta_{T_s})\|_{\ell_2} \\ &\leq (1-\delta)^{-1}\sum_{j=2}^s\|\Phi\eta_{T_j}\|_{\ell_2} \leq (1+\delta)(1-\delta)^{-1}\sum_{j=2}^s\|\eta_{T_j}\|_{\ell_2}, \end{aligned}$$

where we have used both bounds in (3.1). Now for any $i \in T_{j+1}$ and $i' \in T_j$, we have $|\eta_i| \leq |\eta_{i'}|$ so that $|\eta_i| \leq (bk)^{-1}\|\eta_{T_j}\|_{\ell_1}$. It follows that

$$\|\eta_{T_{j+1}}\|_{\ell_2} \leq (bk)^{-1/2}\|\eta_{T_j}\|_{\ell_1}, \quad j = 1, 2, \dots, s-1, \quad (4.1)$$

so that

$$\|\eta_T\|_{\ell_2} \leq (1+\delta)(1-\delta)^{-1}(bk)^{-1/2}\sum_{j=1}^{s-1}\|\eta_{T_j}\|_{\ell_1} \leq (1+\delta)(1-\delta)^{-1}(bk)^{-1/2}\|\eta_{T^c}\|_{\ell_1}. \quad (4.2)$$

By the Cauchy-Schwartz inequality $\|\eta_T\|_{\ell_1} \leq (ak)^{1/2}\|\eta_T\|_{\ell_2}$, and we therefore obtain

$$\|\eta\|_{\ell_1} = \|\eta_T\|_{\ell_1} + \|\eta_{T^c}\|_{\ell_1} \leq \left(1 + \frac{\sqrt{a}(1+\delta)}{\sqrt{b}(1-\delta)}\right)\|\eta_{T^c}\|_{\ell_1} \quad (4.3)$$

which verifies the null space property with the constant C_0 . \square

Combining Corollary 3.3 and Lemma 4.1 (with $a = 2$ and $b = 1$), we have therefore proved the following.

Theorem 4.2 *Let Φ be any matrix which satisfies the RIP of order $3k$. Define the decoder Δ for Φ as in (3.4) for $X = \ell_1$. Then (1.11) holds in $X = \ell_1$ with constant $C_0 = 2(1 + \sqrt{2}\frac{1+\delta}{1-\delta})$. Generally speaking, we cannot derive a constant of the type $1 + \epsilon$ from an analysis based Lemma 4.1, since it requires that the null space property holds with constant $C_0/2$ which is therefore larger than 1.*

As it was mentioned in the previous section, one can build matrices Φ which satisfy the RIP of order k under the condition $n \geq ck \log(N/n)$ where c is some fixed constant. We therefore conclude that instance optimality of order k in the ℓ_1 norm can be achieved at the price of $\mathcal{O}(k \log(N/n))$ measurements.

Remark 4.3 *More generally, if Φ satisfies the RIP of order $(2+b)k$, and Δ is defined by (3.4) for $X = \ell_1$, then (1.11) holds in $X = \ell_1$ with constant $C_0 = 2(1 + \sqrt{2/b} \frac{1+\delta}{1-\delta})$. Therefore, if we make b large, the constant C_0 in (1.11) is of the type $2 + \epsilon$ under a condition of the type $n \geq c \frac{k}{\epsilon^2} \log(N/n)$.*

Note that on the other hand, since instance optimality of order k in any norm X always implies that the reconstruction is exact when $x \in \Sigma_k$, it cannot be achieved with less than $2k$ measurements according to Lemma 3.1.

Before addressing the ℓ_2 case, let us briefly discuss the decoder Δ which achieves (1.11) for such a Φ . According to the proof of Lemma 3.2, one can build Δ as the solution of the minimization problem (3.4). It is not clear to us whether this minimization problem can be solved in polynomial time in N . The following result shows that it is possible to define Δ by ℓ_1 minimization if Φ satisfies the RIP with some additional control on the constants in (3.1).

Theorem 4.4 *Let Φ be any matrix which satisfies the RIP of order $3k$ with $\delta_{3k} \leq \delta < (\sqrt{2} - 1)^2/3$. Define the decoder Δ for Φ as in (1.10). Then, (Φ, Δ) satisfies (1.11) in $X = \ell_1$ with $C_0 = \frac{2\sqrt{2}+2-(2\sqrt{2}-2)\delta}{\sqrt{2}-1-(\sqrt{2}+1)\delta}$.*

Proof: We apply Lemma 4.1 with $a = 1$, $b = 2$ to see that Φ satisfies the null space property in ℓ_1 of order k with constant $C = 1 + \frac{1+\delta}{\sqrt{2}(1-\delta)} < 2$. This means that for any $\eta \in \mathcal{N}$ and T such that $\#T \leq k$, we have

$$\|\eta\|_{\ell_1} \leq C \|\eta_{T^c}\|_{\ell_1}, \quad (4.4)$$

and therefore

$$\|\eta_T\|_{\ell_1} \leq (C - 1) \|\eta_{T^c}\|_{\ell_1}. \quad (4.5)$$

Let $x^* = \Delta(\Phi x)$ be the solution of (1.10) so that $\eta = x^* - x \in \mathcal{N}$ and

$$\|x^*\|_{\ell_1} \leq \|x\|_{\ell_1}. \quad (4.6)$$

Denoting by T the set of indices of the largest k coefficients of x , we can write

$$\|x_T^*\|_{\ell_1} + \|x_{T^c}^*\|_{\ell_1} \leq \|x_T\|_{\ell_1} + \|x_{T^c}\|_{\ell_1}. \quad (4.7)$$

It follows that

$$\|x_T\|_{\ell_1} - \|\eta_T\|_{\ell_1} + \|\eta_{T^c}\|_{\ell_1} - \|x_{T^c}\|_{\ell_1} \leq \|x_T\|_{\ell_1} + \|x_{T^c}\|_{\ell_1}, \quad (4.8)$$

and therefore

$$\|\eta_{T^c}\|_{\ell_1} \leq \|\eta_T\|_{\ell_1} + 2\|x_{T^c}\|_{\ell_1} = \|\eta_T\|_{\ell_1} + 2\sigma_k(x)_{\ell_1}. \quad (4.9)$$

Using (4.5) and the fact that $C < 2$ we thus obtain

$$\|\eta_{T^c}\|_{\ell_1} \leq \frac{2}{2-C} \sigma_k(x)_{\ell_1}. \quad (4.10)$$

We finally use again (4.4) to conclude that

$$\|x - x^*\|_{\ell_1} \leq \frac{2C}{2-C} \sigma_k(x)_{\ell_1}, \quad (4.11)$$

which is the announced result. \square

5 The case $X = \ell_2$

In this section, we shall show that instance-optimality is not a very viable concept in $X = \ell_2$ in the sense that it will not even hold for $k = 1$ unless $n \geq cN$. We know from Corollary 3.3 that if Φ is a matrix of size $n \times N$ which satisfies

$$\|x - \Delta(\Phi x)\|_{\ell_2} \leq C_0 \sigma_k(x)_{\ell_2}, \quad x \in \mathbb{R}^N, \quad (5.1)$$

for some decoder Δ , then its null space \mathcal{N} will need to have the property

$$\|\eta\|_{\ell_2}^2 \leq C_0^2 \|\eta_{T^c}\|_{\ell_2}^2, \quad \#T \leq 2k. \quad (5.2)$$

Theorem 5.1 *For any matrix Φ of dimension $n \times N$, property (5.2) with $k = 1$ implies that $N \leq C_0^2 n$.*

Proof: We start from (5.2) with $k = 1$ from which we trivially derive

$$\|\eta\|_{\ell_2}^2 \leq C_0^2 \|\eta_{T^c}\|_{\ell_2}^2, \quad \#T \leq 1, \quad (5.3)$$

or equivalently for all $j \in \{1, \dots, N\}$,

$$\sum_{i=1}^N |\eta_i|^2 \leq C_0^2 \sum_{i \neq j} |\eta_i|^2. \quad (5.4)$$

From this, we derive that for all $j \in \{1, \dots, N\}$,

$$|\eta_j|^2 \leq (C_0^2 - 1) \sum_{i \neq j} |\eta_i|^2 = (C_0^2 - 1) (\|\eta\|_{\ell_2}^2 - |\eta_j|^2), \quad (5.5)$$

and therefore

$$|\eta_j|^2 \leq A \|\eta\|_{\ell_2}^2, \quad (5.6)$$

with $A = 1 - \frac{1}{C_0^2}$.

Let $(e_j)_{j=1, \dots, N}$ be the canonical basis of \mathbb{R}^N so that $\eta_j = \langle \eta, e_j \rangle$ and let v_1, \dots, v_{N-n} be an orthonormal basis for \mathcal{N} . Denoting by $P = P_{\mathcal{N}}$ the orthogonal projection onto \mathcal{N} , we apply (5.6) to $\eta := P(e_j) \in \mathcal{N}$ and find that for any $j \in \{1, \dots, N\}$

$$|\langle P(e_j), e_j \rangle|^2 \leq A. \quad (5.7)$$

This means

$$\sum_{i=1}^{N-n} |\langle e_j, v_i \rangle|^2 \leq A, \quad j = 1, \dots, N. \quad (5.8)$$

We sum (5.8) over $j \in \{1, \dots, N\}$ and find

$$N - n = \sum_{i=1}^{N-n} \|v_i\|_{\ell_2}^2 \leq AN. \quad (5.9)$$

It follows that $(1 - A)N \leq n$. That is, $N \leq nC_0^2$ as desired. \square

The above result means that when measuring the error in ℓ_2 , the comparison between compressed sensing and best k -term approximation on a general vector of \mathbb{R}^n is strongly in favor of best k -term approximation. However, this conclusion should be moderated in two ways. On the one hand, we shall see in §8 that one can obtain mixed-norm estimates of the form (1.15) from which one finds that compressed sensing compares favorably with best k -term approximation over sufficiently concentrated classes of vectors. On the other hand, we shall prove in the next section that (5.1) can be achieved with n of the same order as k up to a logarithmic factor, if one accepts that this result holds with high probability.

6 The case $X = \ell_2$ in probability

In order to formulate the results of this section, we let Ω be a probability space with probability measure P and let $\Phi = \Phi(\omega)$, $\omega \in \Omega$ be an $n \times N$ random matrix. We seek results of the following type: for any $x \in \mathbb{R}^N$, if we draw Φ at random with respect to P , then

$$\|x - \Delta(\Phi x)\|_{\ell_2} \leq C_0 \sigma_k(x)_{\ell_2} \quad (6.1)$$

holds for this particular x with high probability for some decoder Δ (dependent on the draw Φ). We shall even give explicit decoders which will yield this type of inequality. It should be understood that Φ is drawn independently for each x in contrast to building a Φ such that (6.1) holds simultaneously for all $x \in \mathbb{R}^N$ which was our original definition of instance optimality.

Two simple instances of random matrices which are often considered in compressed sensing are

- (i) Gaussian matrices: $\Phi_{i,j} = \mathcal{N}(0, \frac{1}{n})$ are i.i.d. Gaussian variables of variance $1/n$.
- (ii) Bernoulli matrices: $\Phi_{i,j} = \frac{\pm 1}{\sqrt{n}}$ are i.i.d. Bernoulli variables of variance $1/n$.

In order to establish such results, we shall need that the random matrix Φ has two properties which we now describe. The first of these relates to the restricted isometry property which we know plays a fundamental role in the performance of the matrix Φ in compressed sensing.

Definition 6.1 We say that the random matrix Φ satisfies RIP of order k with constant δ and probability $1 - \epsilon$ if there is a set $\Omega_0 \subset \Omega$ with $P(\Omega_0) \geq 1 - \epsilon$ such that for all $\omega \in \Omega_0$, the matrix $\Phi(\omega)$ satisfies (3.1) with constant $\delta_k \leq \delta$.

This property has been shown for random matrices of the above Gaussian or Bernoulli type. Namely, given any $c > 0$ and $\delta > 0$, there is a constant $c_0 > 0$ such that for all $n \geq c_0 k \log(N/n)$ this property will hold with $\epsilon \leq e^{-cn}$, see [2, 5, 9, 17].

The RIP controls the behavior of Φ on Σ_k , or equivalently on all the k dimensional spaces spanned by any subset of $\{e_1, \dots, e_N\}$ of cardinality k . On the other hand, for a general vector $x \in \mathbb{R}^N$, the image vector Φx might have a much larger norm than x . However, for standard constructions of random matrices the probability that Φx has large norm is small. We formulate this by the following definition.

Definition 6.2 We say that the random matrix Φ has the boundedness property with constant C and probability $1 - \epsilon$, if for each $x \in \mathbb{R}^N$, there is a set $\Omega_0(x) \subset \Omega$ with $P(\Omega_0(x)) \geq 1 - \epsilon$ such that for all $\omega \in \Omega_0(x)$,

$$\|\Phi(\omega)x\|_{\ell_2} \leq C\|x\|_{\ell_2}. \quad (6.2)$$

Note that the property which is required in this definition is clearly weaker than asking that the spectral norm $\|\Phi\| := \sup_{\|x\|_{\ell_2}=1} \|\Phi x\|_{\ell_2}$ is not greater than C with probability $1 - \epsilon$.

Again, this property has been shown for various random families of matrices and in particular for the Gaussian or Bernoulli families. Namely, given any $C > 1$, this property will hold with constant C and $\epsilon \leq 2e^{-\beta n}$ with $\beta = \beta(C) > 0$, see [1] or the discussion in [2]. Thus, the standard constructions of random matrices will satisfy both of these properties.

We now describe our process for decoding $y = \Phi x$, when $\Phi = \Phi(\omega)$ is our given realization of the random matrix. Let $T \subset \{1, \dots, N\}$ be any subset of column indices with $\#(T) = k$ and let X_T be the linear subspace of \mathbb{R}^N which consists of all vectors supported on T . For this T , we define

$$x_T^* := \underset{z \in X_T}{\text{Argmin}} \|\Phi z - y\|_{\ell_2}. \quad (6.3)$$

In other words, x_T^* is chosen as the least squares minimizer of the residual in approximation by elements of X_T . Notice that x_T^* is supported on T . If Φ satisfies RIP of order k then the matrix $\Phi_T^t \Phi_T$ is nonsingular and the nonzero entries of x_T^* are given by

$$(\Phi_T^t \Phi_T)^{-1} \Phi_T^t y. \quad (6.4)$$

To decode y , we search over all subsets T of cardinality k and choose

$$T^* := \underset{\#(T)=k}{\text{Argmin}} \|y - \Phi x_T^*\|_{\ell_2}. \quad (6.5)$$

Our decoding of y is now given by

$$x^* = \Delta(y) := x_{T^*}^*. \quad (6.6)$$

The main result of this section is the following.

Theorem 6.3 *Assume that Φ is a random matrix which satisfies RIP of order $2k$ with constant δ and probability $1 - \epsilon$ and also satisfies the boundedness property with constant C and probability $1 - \epsilon$. Then, for each $x \in \mathbb{R}^N$, there exists a set $\Omega(x) \subset \Omega$ with $P(\Omega(x)) \geq 1 - 2\epsilon$ such that for all $\omega \in \Omega(x)$ and $\Phi = \Phi(\omega)$, the estimate (6.1) holds with $C_0 = 1 + \frac{2C}{1-\delta}$. Here the decoder $\Delta = \Delta(\omega)$ is given by (6.6).*

Proof: Let $x \in \mathbb{R}^N$ be arbitrary and let $\Phi = \Phi(\omega)$ be the draw of the matrix Φ from the random ensemble. We denote by T the set of indices corresponding to the k largest coefficients of x . Thus

$$\|x - x_T\|_{\ell_2} = \sigma_k(x)_{\ell_2}. \quad (6.7)$$

We consider the set $\Omega' := \Omega_0 \cap \Omega(x - x_T)$ where Ω_0 is the set in the definition of RIP in probability and $\Omega(x - x_T)$ is the set in the definition of boundedness in probability for the vector $x - x_T$. Then $P(\Omega') \geq 1 - 2\epsilon$. For any $\omega \in \Omega'$, we have

$$\|x - x^*\|_{\ell_2} \leq \|x - x_T\|_{\ell_2} + \|x_T - x^*\|_{\ell_2} \leq \sigma_k(x)_{\ell_2} + \|x_T - x^*\|_{\ell_2}. \quad (6.8)$$

We bound the second term by

$$\begin{aligned} \|x_T - x^*\|_{\ell_2^N} &\leq (1 - \delta)^{-1} \|\Phi(x_T - x^*)\|_{\ell_2} \\ &\leq (1 - \delta)^{-1} (\|\Phi(x - x_T)\|_{\ell_2} + \|\Phi(x - x^*)\|_{\ell_2}) \\ &= (1 - \delta)^{-1} (\|y - \Phi x_T\|_{\ell_2} + \|y - \Phi x^*\|_{\ell_2}) \\ &\leq 2(1 - \delta)^{-1} \|y - \Phi x_T\|_{\ell_2} = 2(1 - \delta)^{-1} \|\Phi(x - x_T)\|_{\ell_2} \\ &\leq 2C(1 - \delta)^{-1} \|x - x_T\|_{\ell_2} = 2C(1 - \delta)^{-1} \sigma_k(x)_{\ell_2}. \end{aligned}$$

where the first inequality uses the RIP and the fact that $x_T - x^*$ is a vector with support of size less than $2k$, the third inequality uses the minimality of T^* and the fourth inequality uses the boundedness property in probability for $x - x_T$. \square

By virtue of the remarks on the properties of Gaussian and Bernoulli matrices, we derive the following quantitative result.

Corollary 6.4 *If Φ is a random matrix of either Gaussian or Bernoulli type, then for any $\epsilon > 0$ and $C_0 > 3$, there exists a constant c_0 such that if $n \geq c_0 k \log(N/n)$ the following holds: for every $x \in \mathbb{R}^N$, there exists a set $\Omega(x) \subset \Omega$ with $P(\Omega(x)) \geq 1 - 2\epsilon$ such that (6.1) holds for all $\omega \in \Omega(x)$ and $\Phi = \Phi(\omega)$.*

Remark 6.5 *Our analysis yields a constant of the form $C_0 = 3 + \eta$, where η can be made arbitrarily small at the expense of raising n , and it is not clear to us how to improve this constant down to $1 + \eta$ as in [8, 11, 12, 13].*

A variant of the above results deals with the situation where the vector x itself is drawn from a probability measure Q on \mathbb{R}^N . In this case, the following result shows that we can first pick the matrix Φ so that (6.1) will hold with high probability on the choice of x . In other words, only a few pathological signals are not reconstructed up to the accuracy of best k -term approximation.

Corollary 6.6 *If Φ a random matrix of either Gaussian or Bernoulli type, then for any $\epsilon > 0$ and $C_0 > 3$, there exists a constant c_0 such that if $n \geq c_0 k \log(N/n)$ the following holds: there exists a matrix Φ and a set $\Omega(\Phi) \subset \Omega$ with $Q(\Omega(\Phi)) \geq 1 - 2\epsilon$ such that (6.1) holds for all $x \in \Omega(\Phi)$.*

Proof: Consider random matrices of Gaussian or Bernoulli type, and denote by P their probability law. We consider the law $P \otimes Q$ which means that we draw independently Φ according to P and x according to Q . We denote by Ω_x and Ω_Φ the events that (6.1) does not hold given x , and Φ respectively. The event Ω_0 that (6.1) does not hold is therefore given by

$$\Omega_0 = \cup_x \Omega_x = \cup_\Phi \Omega_\Phi. \quad (6.9)$$

According to Corollary 6.4 we know that for all $x \in \mathbb{R}^N$,

$$P(\Omega_x) \leq \epsilon, \quad (6.10)$$

and therefore

$$P \otimes Q(\Omega_0) \leq \epsilon \quad (6.11)$$

By Chebyshev's inequality, we have for all $t > 0$,

$$P(\{\Phi : Q(\Omega_\Phi) \geq t\}) \leq \frac{\epsilon}{t}, \quad (6.12)$$

and in particular

$$P(\{\Phi : Q(\Omega_\Phi) \geq 2\epsilon\}) \leq \frac{1}{2}. \quad (6.13)$$

This shows that there exists a matrix Φ such that $Q(\Omega_\Phi) \leq 2\epsilon$, which means that for such a Φ the estimate (6.1) holds with probability larger than $1 - 2\epsilon$ over x . \square

We close this section with a few remarks comparing the results of this section with other results in the literature. The decoder defined by (6.3) is not computationally realistic since it requires a combinatorial search over all subset T of cardinality T . A natural question is therefore to obtain a decoder with similar approximation properties and more reasonable computational cost. Let us mention that fast decoding methods have been obtained for certain random constructions of matrices by Cormode and Muthukrishnan [8] and by Gilbert and coworkers [13, 18], that yield approximation properties which are similar to Theorem 6.3. Our results differ from theirs in the following two ways. First, we give general criteria for instance optimality to hold in probability. In this context we have not been concerned about the decoder. Our results can hold in particular for standard random classes of matrices such as the Gaussian and Bernoulli constructions. Secondly, when applying our results to these standard random classes, we obtain the range of n given by $n \geq ck \log(N/n)$ which is slightly wider than the range in these other works. That latter range is also treated in [18] but the corresponding results are confined to k -sparse signals. It is shown there that orthogonal matching pursuit identifies the support of such a sparse signal with high probability and that the orthogonal projection will then recover it precisely. More recently, we have shown in [7] that orthogonal matching pursuit when used as a decoder for general random matrices will provide the instance-optimal performance of (6.1). Thus OMP provides a practical decoder for general applications of random matrices in compressed sensing.

7 The case $X = \ell_p$ with $1 < p < 2$

In this section we shall discuss instance optimality in the case $X = \ell_p$ when $1 < p < 2$. We therefore discuss the validity of

$$\|x - \Delta(\Phi x)\|_{\ell_p} \leq C_0 \sigma_k(x)_{\ell_p}, \quad x \in \mathbb{R}^N, \quad (7.1)$$

depending on the value of n . Our first result is a generalization of Lemma 4.1.

Lemma 7.1 *Let Φ be any matrix which satisfies the RIP of order $2k + \tilde{k}$ with $\delta_{2k+\tilde{k}} \leq \delta < 1$ and*

$$\tilde{k} := k \left(\frac{N}{k} \right)^{2-2/p}. \quad (7.2)$$

Then, for any $1 \leq p < 2$, Φ satisfies the null space property in ℓ_p of order $2k$ with constant $C_0 = 2^{\frac{1}{p}-\frac{1}{2}} \frac{1+\delta}{1-\delta}$.

Proof: The proof is very similar to Lemma 4.1 so we sketch it. The idea is to take once again $T_0 = T$ the set of $2k$ largest coefficients of η and to take the other sets T_j of size \tilde{k} .

In the same way, we obtain

$$\|\eta_{T_0}\|_{\ell_2} \leq (1 + \delta)(1 - \delta)^{-1} \sum_{j=2}^s \|\eta_{T_j}\|_{\ell_2}. \quad (7.3)$$

Now if $j \geq 1$, for any $i \in T_{j+1}$ and $i' \in T_j$, we have $|\eta_i| \leq |\eta_{i'}|$ so that $|\eta_i|^p \leq \tilde{k}^{-1} \|\eta_{T_j}\|_{\ell_p}^p$. It follows that

$$\|\eta_{T_{j+1}}\|_{\ell_2} \leq (\tilde{k})^{1/2-1/p} \|\eta_{T_j}\|_{\ell_p}, \quad (7.4)$$

so that

$$\begin{aligned} \|\eta_T\|_{\ell_p} &\leq (2k)^{1/p-1/2} \|\eta_T\|_{\ell_2} \\ &\leq (1 + \delta)(1 - \delta)^{-1} (2k)^{1/p-1/2} \tilde{k}^{1/2-1/p} \sum_{j=1}^s \|\eta_{T_j}\|_{\ell_p} \\ &\leq (1 + \delta)(1 - \delta)^{-1} (2k)^{1/p-1/2} \tilde{k}^{1/2-1/p} s^{1-1/p} \|\eta_{T^c}\|_{\ell_p} \\ &\leq (1 + \delta)(1 - \delta)^{-1} (2k)^{1/p-1/2} \tilde{k}^{1/2-1/p} (N/\tilde{k})^{1-1/p} \|\eta_{T^c}\|_{\ell_p} \\ &= 2^{1/p-1/2} (1 + \delta)(1 - \delta)^{-1} \|\eta_{T^c}\|_{\ell_p}, \end{aligned} \quad (7.5)$$

where we have used twice Hölder's inequality and the relation between N , k and \tilde{k} . \square

The corresponding generalization of Theorem 4.2 is now the following.

Theorem 7.2 *Let Φ be any matrix which satisfies the RIP of order $2k + \tilde{k}$ with $\delta_{2k+\tilde{k}} \leq \delta < 1$ and \tilde{k} as in (7.2). Define the decoder Δ for Φ as in (3.4) for $X = \ell_p$. Then (7.1) holds with constant $C_0 = 2^{1/p+1/2} (1 + \delta)/(1 - \delta)$.*

Recall from our earlier remarks that an $n \times N$ matrix Φ can have RIP of order \tilde{k} provided that $\tilde{k} \leq c_0 n / \log(N/n)$. We therefore conclude from Theorem 7.2 and (7.2) that instance optimality of order k in the ℓ_p norm can be achieved at the price of $\mathcal{O}(k(N/k)^{2-2/p} \log(N/n))$ measurements so that the order of $\mathcal{O}(k(N/k)^{2-2/p} \log(N/k))$ measurements suffices, which is now significantly higher than k except in the case where $p = 1$. In the following, we prove that this price cannot be avoided.

Theorem 7.3 For any $s < 2 - 2/p$ and any matrix Φ of dimension $n \times N$, property (7.1) implies that

$$n \geq ck \left(\frac{N}{k} \right)^s, \quad (7.6)$$

with $c = \left(\frac{C_1}{C_0} \right)^{\frac{2/q-1}{1/q-1/p}}$ where C_0 is the constant in (7.1) and C_1 the lower constant in (2.17) and q is defined by the relation $s = 2 - 2/q$.

Proof: We shall use the results of §2 concerning the Gelfand width and the rate of best k -term approximation. If (1.11) holds, we find that for any compact class $K \subset \mathbb{R}^N$

$$E_n(K)_{\ell_p} \leq C_0 \sigma_k(K)_{\ell_p}. \quad (7.7)$$

We now consider the particular classes $K := U(\ell_q^N)$ with $1 \leq q < p$, so that in view of (2.6) and (2.17), the inequality (7.7) becomes

$$C_1 (N^{1-1/q} n^{-1/2})^{\frac{1/q-1/p}{1/q-1/2}} \leq C_0 k^{1/p-1/q}, \quad (7.8)$$

which gives (7.6) with $s = 2 - 2/q$ and $c = \left(\frac{C_1}{C_0} \right)^{\frac{2/q-1}{1/q-1/p}}$. \square

Remark 7.4 In the above proof the constant c blows up as q approaches p and therefore we cannot directly conclude that a condition of the type $n \geq ck(N/k)^{2-2/p}$ is necessary for (7.1) to hold although this seems plausible.

8 Mixed-norm instance optimality

In this section, we extend the study of instance optimality to more general estimates of the type

$$\|x - \Delta(\Phi x)\|_X \leq C_0 k^{-s} \sigma_k(x)_Y, \quad x \in \mathbb{R}^N, \quad (8.1)$$

which we refer to as mixed-norm instance optimality. We have in mind the situation where $X = \ell_p$ and $Y = \ell_q$ with $1 \leq q \leq p \leq 2$ and $s = 1/q - 1/p$. We are thus interested in estimates of the type

$$\|x - \Delta(\Phi x)\|_{\ell_p} \leq C_0 k^{1/p-1/q} \sigma_k(x)_{\ell_q}, \quad x \in \mathbb{R}^N. \quad (8.2)$$

The interest in such estimates stems from the following fact. Considering the classes $K = U(\ell_r^N)$ for $r < q$, we know from (2.8) that

$$k^{1/p-1/q} \sigma_k(K)_{\ell_q} \sim k^{1/p-1/q} k^{1/q-1/r} = k^{1/p-1/r} \sim \sigma_k(K)_{\ell_p}. \quad (8.3)$$

Therefore the estimate (8.2) yields the same rate of approximation as (7.1) over such classes, and on the other hand we shall see that it is valid for smaller values of n .

Our first result is a trivial generalization of Lemma 3.2 and Theorem 3.3 to the case of mixed norm instance-optimality, so we state it without proof. We say that Φ has the *mixed null space property* in (X, Y) of order k with constant C and exponent s if

$$\|\eta\|_X \leq C k^{-s} \|\eta_{T^c}\|_Y, \quad (8.4)$$

$\eta \in \mathcal{N}$ and $\#(T) \leq k$.

Theorem 8.1 *Given a norm $\|\cdot\|_X$, an integer $k > 0$ and an encoding matrix Φ . If Φ has the mixed null space property in (X, Y) of order $2k$ with constant $C_0/2$ and exponent s , then there exists a decoder Δ so that (Φ, Δ) satisfies (8.1) with constant C_0 . Conversely, the validity of (8.1) for some decoder Δ implies that Φ has the null space property in (X, Y) of order $2k$ with constant C_0 and exponent s .*

We next give a straightforward generalization of Lemma 7.1.

Lemma 8.2 *Let Φ be any matrix which satisfies the RIP of order $2k + \tilde{k}$ with $\delta_{2k+\tilde{k}} \leq \delta < 1$ and*

$$\tilde{k} := k \left(\frac{N}{k} \right)^{2-2/q}. \quad (8.5)$$

Then Φ satisfies the mixed null space property in (ℓ_p, ℓ_q) of order $2k$ with constant $C_0 = 2^{\frac{1}{p} + \frac{1}{2} \frac{1+\delta}{1-\delta}} + 2^{\frac{1}{p} - \frac{1}{q}}$ and exponent $s = 1/q - 1/p$.

Proof: As in the proof of Lemma 7.1, we take $T_0 = T$ the set of $2k$ largest coefficients of η and to take the other sets T_j of size \tilde{k} . By similar arguments, we arrive to the chain of inequalities

$$\begin{aligned} \|\eta_T\|_{\ell_p} &\leq (2k)^{1/p-1/2} \|\eta_T\|_{\ell_2} \\ &\leq (1+\delta)(1-\delta)^{-1} (2k)^{1/p-1/2} \tilde{k}^{1/2-1/q} \sum_{j=1}^s \|\eta_{T_j}\|_{\ell_q} \\ &\leq (1+\delta)(1-\delta)^{-1} (2k)^{1/q-1/2} \tilde{k}^{1/2-1/q} s^{1-1/q} \|\eta_{T^c}\|_{\ell_q} \\ &\leq (1+\delta)(1-\delta)^{-1} (2k)^{1/q-1/2} \tilde{k}^{1/2-1/q} (N/\tilde{k})^{1-1/q} \|\eta_{T^c}\|_{\ell_q} \\ &= 2^{1/p-1/2} (1+\delta)(1-\delta)^{-1} k^{-s} \|\eta_{T^c}\|_{\ell_q}, \end{aligned} \quad (8.6)$$

where we have used Hölder's inequality both with ℓ_q and ℓ_p as well as the relation between N , k and \tilde{k} .

It remains to bound the tail $\|\eta_{T^c}\|_{\ell_p}$. To this end, we infer from (2.4) that

$$\|\eta_{T^c}\|_{\ell_p} \leq \|\eta\|_{\ell_q} (2k)^{\frac{1}{p} - \frac{1}{q}} \leq (\|\eta_T\|_{\ell_q} + \|\eta_{T^c}\|_{\ell_q}) (2k)^{\frac{1}{p} - \frac{1}{q}}.$$

Invoking (7.5) for $p = q$ yields now

$$\|\eta_T\|_{\ell_q} \leq 2^{1/q-1/2} (1+\delta)(1-\delta)^{-1} \|\eta_{T^c}\|_{\ell_q}$$

so that

$$\|\eta_{T^c}\|_{\ell_p} \leq \left(2^{\frac{1}{p} - \frac{1}{2}} (1+\delta)(1-\delta)^{-1} + 2^{\frac{1}{p} - \frac{1}{q}} \right) \|\eta_{T^c}\|_{\ell_q} k^{\frac{1}{p} - \frac{1}{q}}. \quad (8.7)$$

Combining (8.7) and (8.6), finishes the proof. \square

We see that considering mixed-norm instance optimality in (ℓ_p, ℓ_q) in contrast to instance optimality in ℓ_q is beneficial since the value of \tilde{k} is smaller in (8.5) than in (7.2). The corresponding generalization of Theorem 7.2 is now the following.

Theorem 8.3 *Let Φ be any matrix which satisfies the RIP of order $2k + \tilde{k}$. Define the decoder Δ for Φ as in (3.4) for $X = \ell_p$. Then (8.2) holds with constant $C_0 = 2^{\frac{1}{p} + \frac{3}{2}} \frac{1+\delta}{1-\delta} + 2^{1 + \frac{1}{p} - \frac{1}{q}}$.*

By the same reasoning that followed Theorem 7.2 concerning the construction of matrices which satisfy RIP, we conclude that mixed instance optimality of order k in the ℓ_p and ℓ_q norm can be achieved at the price of $\mathcal{O}(k(N/k)^{2-2/q} \log(N/k))$ measurements. In particular, we see that when $q = 1$, this type of mixed norm estimate can be obtained with n larger than k only by a logarithmic factor. Such a result was already observed in [4] in the case $p = 2$ and $q = 1$. In view of (8.3) this implies in particular that compressed sensing behaves as good as best k -term approximation on classes such as $K = U(\ell_r^N)$ for $r < 1$.

One can prove that the above number of measurements is also necessary. This is expressed by a straightforward generalization of Theorem 7.3 that we state without proof.

Theorem 8.4 *For any matrix Φ of dimension $n \times N$, property (8.2) implies that*

$$n \geq ck \left(\frac{N}{k} \right)^{2-2/q}, \quad (8.8)$$

with $c = \left(\frac{C_1}{C_0} \right)^{\frac{2/q-1}{1/q-1/p}}$ where C_0 is the constant in (7.1) and C_1 the lower constant in (2.17).

Remark 8.5 *In general, there is no direct relationship between (7.1) and (8.2). We give an example to bring out this fact. Let us consider a fixed value of $1 < p \leq 2$ and values of N and $k < N/2$. We define x so that its first k coordinates are 1 and its remaining $N - k$ coordinates are in $(0, 1)$. Then $\sigma_k(x)_{\ell_r} = \|z\|_{\ell_r}$ where z is obtained from x by setting the first k coordinates of x equal to zero. We can choose z so that $1/2 \leq \|z\|_{\ell_r} \leq 2$, for $r = p, q$. In this case, the right side in (8.2) is smaller than the right side of (7.1) by the factor $k^{1/p-1/q}$ so an estimate in the mixed-norm instance-optimality sense is much better for this x . On the other hand, if we take all nonzero coordinates of z to be a with $a \in (0, 1)$, then the right side of (7.1) will be smaller than the right side of (8.2) by the factor $(N/k)^{1/p-1/q}$ which show that for this x the instance-optimality estimate is much better.*

References

- [1] D. Achilioptas, Database-friendly random projections, Proc. ACM Symposium on the Principles of Database Systems, pp. 274-281, 2001
- [2] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, A simple proof of the restricted isometry property for random matrices, Constructive Approximation, to appear.
- [3] E. J. Candès, J. Romberg and T. Tao, Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information, *IEEE Trans. Inf. Theory*, **52**(2006), 489–509.

- [4] E. Candès, J. Romberg, and T. Tao, Stable signal recovery from incomplete and inaccurate measurements, *Comm. Pure and Appl. Math.*, **59**(2006), 1207–1223.
- [5] E. Candès and T. Tao, Decoding by linear programming, *IEEE Trans. Inf. Theory* **51**(2005), 4203–4215.
- [6] E. Candès and T. Tao, Near optimal signal recovery from random projections:universal encoding strategies, *IEEE Trans. Inf. Theory* **52**(2006), 5406–5425.
- [7] A. Cohen, W. Dahmen, and R. DeVore, Near optimal approximation of arbitrary vectors from highly incomplete measurements, preprint.
- [8] G. Cormode and S. Muthukrishnan, Towards an algorithmic theory of compressed sensing, Technical Report 2005-25, DIMACS, 2005, also, Proc. SIROCCO, 2006.
- [9] D. Donoho, Compressed Sensing, *EEE Trans. Information Theory*, **52** (2006), 1289-1306.
- [10] A.Garnaev, E.D. Gluskin, The widths of a Euclidean ball, *Doklady AN SSSR*, 277 (1984), 1048-1052.
- [11] A. Gilbert, Y. Kotidis, S. Muthukrishnan and M. Strauss, How to summarize the universe: Dynamic maintenance of quantiles, *Proc. VLDB*, 2002, 454465.
- [12] A. Gilbert, S. Guha, Y. Kotidis, P. Indyk, S. Muthukrishnan and M. Strauss, Fast, small space algorithm for approximate histogram maintenance, *ACM STOC*, 2002, 389–398.
- [13] A. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan and M. Strauss, Near-optimal sparse fourier estimation via sampling, *ACM STOC*, 2002, 152161.
- [14] E.D. Gluskin, Norms of random matrices and widths of finite-dimensional sets, *Math. USSR Sbornik*, 48(1984), 173–182.
- [15] B. Kashin, The widths of certain finite dimensional sets and classes of smooth functions, *Izvestia* **41**(1977), 334–351.
- [16] G.G. Lorentz, M. von Golitschek and Yu. Makovoz, *Constructive Approximation:Advanced Problems*, Springer Grundlehren, vol. 304, Springer Berlin Heidelberg, 1996.
- [17] A. Pajor and N. Tomczak-Jaegermann, *Subspaces of small codimension of finite dimensional Banach spaces*, *Proc. Amer. Math. Soc.*, vol. 97, 1986, pp. 637–642.
- [18] J.A. Tropp, A.C. Gilbert, Signal recovery from partial information via orthogonal matching pursuit, preprint, April 2005.

Albert Cohen, Laboratoire Jacques-Louis Lions, Université Pierre et Marie Curie 175, rue du Chevaleret, 75013 Paris, France, cohen@ann.jussieu.fr

Wolfgang Dahmen, Institut für Geometrie und Praktische Mathematik, RWTH Aachen, Templergraben 55, D-52056 Aachen Germany, dahmen@igpm.rwth-aachen.de

Ronald DeVore, Industrial Mathematics Institute, University of South Carolina, Columbia, SC 29208, devore@math.sc.edu