On Theory of Compressive Sensing via ℓ_1-Minimization: Simple Derivations and Extensions

Yin Zhang
(CAAM Technical Report TR08-11)
Department of Computational and Applied Mathematics
Rice University, Houston, Texas, 77005.

July, 2008 (Updated September, 2008)

Abstract

Compressive (or compressed) sensing (CS) is an emerging methodology in computational signal processing that has recently attracted intensive research activities. At present, the basic CS theory includes recoverability and stability: the former quantifies the central fact that a sparse signal of length n can be exactly recovered from far fewer than n measurements via ℓ_1-minimization or other recovery techniques, while the latter specifies the stability of a recovery technique in the presence of measurement errors and inexact sparsity. So far, most analyses in CS rely heavily on the Restricted Isometry Property (RIP) for matrices.

In this paper, we present an alternative, non-RIP analysis for CS via ℓ_1-minimization. Our purpose is three-fold: (a) to introduce an elementary treatment of the CS theory free of RIP and easily accessible to a broad audience; (b) to extend the current recoverability and stability results so that prior knowledge can be utilized to enhance recovery via ℓ_1-minimization; and (c) to substantiate a property called uniform recoverability of ℓ_1-minimization; that is, for almost all random measurement matrices recoverability is asymptotically identical. With the aid of two classic results, the non-RIP approach enables us to quickly derive from scratch all basic results for the extended theory.

Key words. Compressive Sensing, ℓ_1-minimization, non-RIP analysis, recoverability and stability, prior information, uniform recoverability.

1 Introduction

1.1 What is Compressive Sensing?

The flows of data (e.g., signals and images) around us are enormous today and rapidly growing. However, the number of salient features hidden in massive data are usually much smaller than their sizes. Hence data are compressible. In data processing, the traditional practice is to measure (sense) data in full length and then compress the resulting measurements before storage or transmission. In such a scheme, recovery of data is generally straightforward. This traditional data-acquisition

process can be described as "full sensing plus compressing". *Compressive sensing* (CS), also known as compressed sensing or compressive sampling, represents a paradigm shift in which the number of measurements is reduced during acquisition so that no additional compression is necessary. The price to pay is that more sophisticated recovery procedures become necessary.

In this paper, we will use the term "signal" to represent generic data (so an image is a signal). Let $\bar{x} \in \mathbb{R}^n$ represent a discrete signal and $b \in \mathbb{R}^m$ a vector of linear measurements formed by taking inner products of $\bar{x}$ with a set of linearly independent vectors $a_i \in \mathbb{R}^n$, $i = 1, 2, \cdots, m$. In matrix format, the measurement vector is $b = A\bar{x}$, where $A \in \mathbb{R}^{m \times n}$ has rows a_i^T, $i = 1, 2, \cdots, m$. This process of obtaining b from an unknown signal $\bar{x}$ is often called *encoding*, while the process of recovering $\bar{x}$ from the measurement vector b is called *decoding.*

When the number of measurements m is equal to n, decoding simply entails solving a linear system of equations, i.e., $\bar{x} = A^{-1}b$. However, in many applications, it is much more desirable to take fewer measurements provided one can still recover the signal. When $m < n$, the linear system $Ax = b$ is typically under-determined, permitting infinitely many solutions. In this case, is it still possible to recover $\bar{x}$ from b through a computationally tractable procedure?

If we know that the measurement b is from a highly sparse signal (i.e., it has very few nonzero components), then a reasonable decoding model is to look for the sparsest signal among all those that produce the measurement b; that is,

$$\min\{\|x\|_0 : Ax = b\}, \tag{1}$$

where the quantity $\|x\|_0$ denotes the number of non-zeros in x. Model (1) is a combinatorial optimization problem with a prohibitive complexity if solved by enumeration, and thus does not appear tractable. An alternative model is to replace the "ℓ_0-norm" by the ℓ_1-norm and solve a computationally tractable linear program:

$$\min\{\|x\|_1 : Ax = b\}. \tag{2}$$

This approach, often called basis pursuit, was popularized by Chen, Donoho and Sauders [8] in signal processing, though similar ideas existed earlier in other areas such as geo-sciences (see Santosa and Symes [27], for example).

In most applications, sparsity is hidden in a signal $\bar{x}$ so that it becomes sparse only under a "sparsifying" basis Φ; that is, $\Phi\bar{x}$ is sparse instead of $\bar{x}$ itself. In this case, one can do a change of variable $z = \Phi x$ and replace the equation $Ax = b$ by $(A\Phi^{-1})z = b$. For an orthonormal basis Φ, the null space of $A\Phi^{-1}$ is a rotation of that of A, and such a rotation does not alter the success rate of CS recovery (as we will see later that the probability measure for success or failure is rotation-invariant). For simplicity and without loss of generality, we will assume $\Phi = I$ throughout this paper.

Fortunately, under favorable conditions the combinatorial problem (1) and the linear program (2) can share common solutions. Specifically, if the signal $\bar{x}$ is sufficiently sparse and the measurement matrix A possesses certain nice attributes (to be specified later), then $\bar{x}$ will solve both (1) and (2) for $b = A\bar{x}$. This property is called *recoverability*, which, along with the fact that (2) can be solved efficiently in theory, establishes the theoretical soundness of the decoding model (2).

Generally speaking, compressive sensing refers to the following two-step approach: choosing a measurement matrix $A \in \mathbb{R}^{m \times n}$ with $m < n$ and taking measurement $b = A\bar{x}$ on a sparse signal $\bar{x}$, and then reconstructing $\bar{x} \in \mathbb{R}^n$ from $b \in \mathbb{R}^m$ by some means. Since $m < n$, the measurement b is already compressed during sensing, hence the name "compressive sensing" or CS. Using the basis pursuit model (2) to recover $\bar{x}$ from b represents a fundamental instance of CS, but certainly not the only one. Other recovery techniques include greedy-type algorithms (see [28], for example).

In this paper, we will exclusively focus on ℓ_1-minimization decoding models, including (2) as a special case, because ℓ_1-minimization has the following two advantages: (a) the flexibility to incorporate prior information into decoding models, and (b) uniform recoverability. These two advantages will be introduced and studied in this paper.

1.2 Current Theory for CS via ℓ_1-minimization

Basic theory of CS presently consists of two components: recoverability and stability. Recoverability addresses the central questions: what types of measurement matrices and recovery procedures ensure exact recovery of all k-sparse signals (those having exactly k-nonzeros) and what is the best order m for the sparsity k? On the other hand, stability addresses the robustness issues in recovery when measurements are noisy and/or sparsity is inexact.

There are a number of earlier works that have laid the groundwork for the existing CS theory, especially pioneering works by Dohono and his co-workers (see the survey paper [3] for a list of references on these early works). From these early works, it is known that certain matrices can guarantee recovery for sparsity k up to the order of $\sqrt{m}$ (for example, see Donoho and Elad [11]). In recent seminal works by Candés and Tao [4, 5], it is shown that for a standard normal random matrix $A \in \mathbb{R}^{m \times n}$, recoverability is ensured with high probability for sparsity k up to the order of $m/\log(n/m)$, which is the best recoverability order available. Later the same order has been extended by Baraniuk *et al.* [1] to a few other random matrices such as Bernoulli matrices with ± 1 entries.

In practice, it is almost always the case that either measurements are noisy or signal sparsity is inexact, or both. Here inexact sparsity refers to the situation where a signal contains a small number of significant components in magnitude, while the magnitudes of the rest are small but not necessarily zero. Such approximately sparse signals are compressible too. The subject of CS stability studies the issues concerning how accurately a CS approach can recover signals under these circumstances. Stability results have been established for the ℓ_1-minimization model (2) and its extension

$$\min\{\|x\|_1 : \|Ax - b\|_2 \leq \delta\}. \tag{3}$$

Consider model (2) for $b = A\hat{x}$ where $\hat{x}$ is approximately k-sparse so that it has only k significant components. Let $\hat{x}(k)$ be a so-called k-term approximation of $\hat{x}$ obtained by setting the $n - k$ insignificant components of $\hat{x}$ to zero, and let x^* be the optimal solution of (2). Existing stability results for model (2) include the following two types of error bounds,

$$\|x^* - \hat{x}\|_2 \leq Ck^{-1/2}\|\hat{x} - \hat{x}(k)\|_1, \tag{4}$$

3

$$\|x^* - \hat{x}\|_1 \leq C\|\hat{x} - \hat{x}(k)\|_1, \tag{5}$$

where the sparsity level k can be up to the order of $m/\log(n/m)$ depending on what type of measurement matrices are in use, and C denotes a generic constant independent of dimensions whose value may vary from one place to another. These results are established by Candés and Tao [4] and Candés, Romberg and Tao [6] (see also Donoho [10], and Cohen, Dahmen and DeVore [9]). For the extension model (3), the following stability result is obtained by Candés, Romberg and Tao [6]:

$$\|x^* - \hat{x}\|_2 \leq C(\delta + k^{-1/2}\|\hat{x} - \hat{x}(k)\|_1). \tag{6}$$

In the case where $\hat{x}$ is exactly k-sparse so that $\hat{x} = \hat{x}(k)$, the above stability results reduce to the exact recoverability: $x^* = \hat{x}$ (also $\delta = 0$ is required in (6)). Therefore, when combined with relevant random matrix properties, the stability results imply recoverability in the case of solving model (2). More recently, stability results have also been established for some greedy algorithms by Needell and Vershynin [25] and Needell and Tropp [24]. Yet, there still exist CS recovery methods that have been shown to possess recoverability but with stability unknown.

Existing results in CS recoverability and stability are mostly based on analyzing properties of the measurement matrix A. The most widely used analytic tool is the so-called Restricted Isometry Property (RIP) of A, first introduced in [5] for the analysis of CS recoverability (but an earlier usage can be found in [20]). Given $A \in \mathbb{R}^{m \times n}$, the k-th RIP parameter of A, $\delta_k(A)$, is defined as the smallest quantity δ that satisfies for some $R > 0$

$$(1 - \delta)R \leq \left(\frac{\|Ax\|_2}{\|x\|_2}\right)^2 \leq (1 + \delta)R, \quad \forall x, \|x\|_0 = k. \tag{7}$$

The smaller $\delta_k(A)$ is, the better RIP is for that k value. Roughly speaking, RIP measures the "overall conditioning" of the set of $m \times k$ submatrices of A.

All the above stability results (including those in [25, 24]) have been obtained under various assumptions on the RIP parameters of A. For example, the error bounds (4) and (6) were first obtained under the assumption $\delta_{3k}(A) + 3\delta_{4k}(A) < 2$. Consequently, stability constants, represented by C above, obtained by existing RIP-based analyses all depend on the RIP parameters of A. We will see, however, that as far as the results within the scope of this paper are concerned, the dependency on RIP can be removed.

Donoho established stable recovery results in [10] under three conditions on measurement matrix A (conditions CS1-CS3). Although these conditions do not directly use RIP properties, they are still matrix-based. For example, condition CS1 requires that the minimum singular values of all $m \times k$ sub-matrices, with $k < \rho m/\log(n)$ for some $\rho > 0$, of $A \in \mathbb{R}^{m \times n}$ be uniformly bounded below from zero. Consequently, the stability results in [10] are all dependent on matrix A.

Another analytic tool, mostly used by Donoho and his co-authors (for example, see [12]) is to study the combinatorial and geometric properties of the polytope formed by the columns of A. While the RIP approach uses sufficient conditions for recoverability, the "polytope approach" uses a necessary and sufficient condition. Although the latter approach can lead to tighter recoverability constants, the former has so far produced stability results such as (4)-(6).

4

1.3   Contributions of this Paper

Since existing RIP-based analyses are usually rather intricate and lengthy, we seek to devise an elementary approach to produce a more accessible and concise treatment. Here an interesting question arises: Is RIP really an indispensable property for CS analyses? In any recovery model using the equation $Ax = b$, the pair (A, b) carries all the information about the unknown signal. Obviously, $Ax = b$ is equivalent to $GAx = Gb$ for any nonsingular matrix $G \in \mathbb{R}^{m \times m}$. Numerical considerations aside, (GA, Gb) ought to carry exactly the same amount of information as (A, b) does. However, the RIP properties of A and GA can be vastly different. One can easily choose G to make the RIP of GA arbitrarily bad no matter how good the RIP of A is. This observation indicates that RIP is far from indispensable. (See 4.3 for further discussions on the RIP issue).

One contribution of this paper is to use a non-RIP analysis to develop the theory for CS via ℓ_1-minimization. Incidentally, this non-RIP analysis also gives stronger results than those by previous RIP analyses in two aspects described below.

In practice, *a priori* knowledge often exists about the signal to be recovered. Beside sparsity, existing CS theory does not incorporate such prior-information in decoding models, with the exception of when the signs of a signal are known (see [13, 30], for example). Can we extend the existing theory to explicitly include prior-information into our decoding models? In particular, it is desirable to analyze the following general model:

$$\min\{\|x\|_1 : \|Ax - b\| \le \delta, x \in \Omega\}, \tag{8}$$

where $\| \cdot \|$ is a generic norm, $\delta \ge 0$ and the set $\Omega \subset \mathbb{R}^n$ represents prior information about the signal under construction.

So far, different types of measurement matrices have required different analyses, and only a few random matrix types, such as Gaussian and Bernoulli distributions with rather restrictive conditions such as zero mean, have been studied in details. Hence a framework that covers a wide range of measurement matrices becomes highly desirable. Moreover, an intriguing challenge is to show that a large collection of random matrices shares an asymptotic and identical recovery behavior, which appears to be true from empirical observations by different groups (see [14, 15], for example). We will examine this phenomenon that we call *uniform recoverability*, meaning that recoverability is essentially invariant with respect to different types of random matrices.

In summary, this paper consists of three contributions to the theory of CS via ℓ_1-minimization: (i) an elementary, non-RIP treatment; (ii) a flexible theory that allows any form of prior information to be explicitly included; and (iii) a theoretical explanation of the uniform recoverability phenomenon.

This paper is self-contained, in that it includes derivations for all results, with the exception of invoking two classic results without proofs. To make the paper more accessible to a broad audience while limiting its length, we keep discussions at a minimal level on issues not of immediate relevance, such as historical connections and technical details on some deep mathematical notions, and we do not pursue generality.

This paper is not a comprehensive survey on the subject of CS (see [3] for a recent survey on RIP-based CS theory), and does not cover every aspect of CS. In particular, this paper does

not discuss in any detail CS applications and algorithms. We refer the reader to the CS resource website at Rice University [31] for more complete and up-to-date information on CS research and practice.

1.4 Notation and Organization

The ℓ_1-norm and the ℓ_2-norm in $\mathbb{R}^n$ are denoted by $\|\cdot\|_1$ and $\|\cdot\|_2$, respectively. For any vector $v \in \mathbb{R}^n$ and $\alpha \subset \{1, \cdots, n\}$, v_α is the sub-vector of v whose elements are indexed by those in α. Scalar operations, such as absolute value, are applied to vectors component-wise. The support of a vector x is denoted as

$$\text{supp}(x) \triangleq \{i : x_i \neq 0\},$$

and $\|x\|_0 \triangleq |\text{supp}(x)|$ is the number of non-zeros in x where $|\cdot|$ denotes cardinality of a set. For $\mathcal{F} \subset \mathbb{R}^n$ and $\bar{x} \in \mathbb{R}^n$, define $\mathcal{F} - \bar{x} \triangleq \{x - \bar{x} : x \in \mathcal{F}\}$. For random variables, we use the symbol *iid* for the phrase "independently identically distributed".

This paper is organized as follows. In section 2, we introduce some basic conditions for CS recovery via ℓ_1-minimization. An extended recoverability result is presented in Section 3 for standard normal random matrices. Section 4 contains two stability results. We explain the uniform recoverability phenomenon in Section 5. The last section contains concluding remarks.

2 Sparsest Point and ℓ_1-Minimization

In this section we introduce general conditions that relate finding the sparsest point to ℓ_1-minimization. Given a set $\mathcal{F} \subset \mathbb{R}^n$, we consider the following equivalence relationship:

$$\arg\min_{x \in \mathcal{F}} \|x\|_0 = \{\bar{x}\} = \arg\min_{x \in \mathcal{F}} \|x\|_1. \tag{9}$$

The problem on the left is a combinatorial problem of finding the sparsest point in $\mathcal{F}$, while the one on the right is a continuous, ℓ_1-minimization problem. The singleton set in the middle indicates that there is a *common and unique* minimizer $\bar{x}$ for both problems.

2.1 Preliminaries

For $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $\Omega \subset \mathbb{R}^n$, we will study equivalence (9) on sets of the form:

$$\mathcal{F} = \{x : Ax = b\} \cap \Omega. \tag{10}$$

For any $\bar{x} \in \mathcal{F}$, the following identity will be useful:

$$\mathcal{F} - \bar{x} \equiv \mathcal{N}(A) \cap (\Omega - \bar{x}), \tag{11}$$

where $\mathcal{N}(A)$ denotes the null space of the matrix A. This identity can be verified as follows:

$$\begin{aligned}
\mathcal{F} &= \{x : Ax = b, x \in \Omega\} \\
&= \{\bar{x} + v : Av = 0, v \in \Omega - \bar{x}\} \\
&= \bar{x} + \{v : Av = 0, v \in \Omega - \bar{x}\} \\
&= \bar{x} + \mathcal{N}(A) \cap (\Omega - \bar{x}).
\end{aligned}$$

We start from the following simple but important observation.

Lemma 1. *For $x, y \in \mathbb{R}^n$ and $\alpha = \text{supp}(x)$, $\|x\|_1 < \|y\|_1$ if*

$$\|y - x\|_1 > 2\|(y - x)_\alpha\|_1. \tag{12}$$

Moreover,

$$\sqrt{\|x\|_0} < \frac{1}{2}\frac{\|y - x\|_1}{\|y - x\|_2} \tag{13}$$

implies both (12) and

$$\sqrt{\|y\|_0} \geq \frac{1}{2}\frac{\|y - x\|_1}{\|y - x\|_2}. \tag{14}$$

Proof. Since $\alpha = \text{supp}(x)$, we have $\|x\|_1 = \|x_\alpha\|_1$ and $x_\beta = 0$ where $\beta = \{1, \cdots, n\} \setminus \alpha$. Let $y = x + v$. We calculate

$$\begin{aligned}
\|y\|_1 &= \|x + v\|_1 = \|x_\alpha + v_\alpha\|_1 + \|0 + v_\beta\|_1 \\
&= \|x\|_1 + (\|v_\beta\|_1 - \|v_\alpha\|_1) + (\|x_\alpha + v_\alpha\|_1 - \|x_\alpha\|_1 + \|v_\alpha\|_1),
\end{aligned} \tag{15}$$

where in the right-hand side we have added and subtracted the terms $\|x\|_1$ and $\|v_\alpha\|_1$. In the above identity, the last term in parentheses is nonnegative due to the triangle inequality:

$$\|x_\alpha + v_\alpha\|_1 \geq \|x_\alpha\|_1 - \|v_\alpha\|_1. \tag{16}$$

Hence, $\|x\|_1 < \|y\|_1$ if $\|v_\beta\|_1 > \|v_\alpha\|_1$ which is equivalent to $\|v\|_1 > 2\|v_\alpha\|_1$. This proves the first part of the lemma. In view of the relationship between the 1-norm and 2-norm,

$$\|v_\alpha\|_1 \leq \sqrt{|\alpha|}\|v_\alpha\|_2 \leq \sqrt{|\alpha|}\|v\|_2 = \sqrt{\|x\|_0}\|v\|_2.$$

Hence, $\|v\|_1 > 2\|v_\alpha\|_1$ if $\|v\|_1 > 2\sqrt{\|x\|_0}\|v\|_2$, which proves that (13) implies (12). Finally, if (13) holds but (14) does not, then due to the symmetry in the right-hand side of (13), or (14), with respect to x and y, a contradiction would arise in that $\|x\|_1 < \|y\|_1$ and $\|y\|_1 < \|x\|_1$. This completes the proof. $\qquad\square$

2.2 Sufficient Conditions

We now introduce two sufficient conditions, (17)-(18), for recoverability. Both are rather straight-forward observations, and the latter, for the case $\Omega = \mathbb{R}^n$, is well-known.

Proposition 1. *For any $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, and $\Omega \subset \mathbb{R}^n$, equivalence (9) holds uniquely at $\bar{x} \in \mathcal{F}$, where $\mathcal{F}$ is defined as in (10), if the sparsity of $\bar{x}$ satisfies*

$$\sqrt{\|\bar{x}\|_0} < \min \left\{ \frac{1}{2} \frac{\|v\|_1}{\|v\|_2} \; : \; v \in \mathcal{N}(A) \cap (\Omega - \bar{x}) \setminus \{0\} \right\}. \tag{17}$$

Moreover, the condition

$$\sqrt{\|\bar{x}\|_0} < \min \left\{ \frac{1}{2} \frac{\|v\|_1}{\|v\|_2} \; : \; v \in \mathcal{N}(A) \setminus \{0\} \right\}, \tag{18}$$

corresponding to $\Omega = \mathbb{R}^n$, implies (17).

Proof. Since $\mathcal{F} \equiv \{\bar{x} + v : v \in \mathcal{F} - \bar{x}\}$ for any $\bar{x}$, the right half of equivalence (9) holds uniquely at $\bar{x} \in \mathcal{F}$ if and only if

$$\|\bar{x} + v\|_1 > \|\bar{x}\|_1, \; \forall\, v \in (\mathcal{F} - \bar{x}) \setminus \{0\}. \tag{19}$$

In view of the identity (11) and condition (13) in Lemma 1, (17) is clearly a sufficient condition for (19) to hold. Obviously, (18) implies (17) because the minimum in (17) is taken over a subset of that in (18).

Now we show that the left half of equivalence (9) also holds uniquely at $\bar{x}$ provided it satisfies (17). Note that (17) is equivalent to

$$\sqrt{\|\bar{x}\|_0} < \frac{1}{2} \frac{\|y - \bar{x}\|_1}{\|y - \bar{x}\|_2}, \; \forall\, y \in \mathcal{F}, y \neq \bar{x}.$$

Hence, it follows from the last part of Lemma 1 that

$$\sqrt{\|y\|_0} \geq \frac{1}{2} \frac{\|y - \bar{x}\|_1}{\|y - \bar{x}\|_2} > \sqrt{\|\bar{x}\|_0}, \; \forall\, y \in \mathcal{F}, y \neq \bar{x}.$$

Consequently, $\bar{x}$ must be the sparsest point in $\mathcal{F}$. $\qquad\square$

Remark 1. *It is worth noting that for some prior-information set Ω, the right-hand side of (17) could be significantly larger than that of (18), suggesting that adding prior information can never hurt but potentially help raise the lower bound on recoverable sparsity. Since $0 \in \mathcal{N}(A) \cap (\Omega - \bar{x})$ and $\|v\|_1/\|v\|_2$ is scale-invariant, a necessary condition for the right-hand side of (17) to be larger than that of (18) is that the origin is on the boundary of $\Omega - \bar{x}$ (which holds true for the nonnegative orthant, for example).*

2.3   A Necessary and Sufficient Condition

For the case $\Omega = \mathbb{R}^n$, we consider the situation of using a fixed measurement matrix A for the recovery of signals of a fixed sparsity level but all possible values and sparsity patterns.

Proposition 2. *Given $A \in \mathbb{R}^{m \times n}$ and any integer $k \geq 1$, the equivalence*

$$\{\bar{x}\} = \arg \min\{\|x\|_1 : Ax = A\bar{x}\} \tag{20}$$

holds for all $\bar{x} \in \mathbb{R}^n$ such that $\|\bar{x}\|_0 \leq k$ if and only if

$$\|v\|_1 > 2\|v_\alpha\|_1, \ \forall\, v \in \mathcal{N}(A) \setminus \{0\}, \tag{21}$$

holds for all index sets $\alpha \subset \{1, \cdots, n\}$ such that $|\alpha| = k$.

Proof. In view of the fact $\{x : Ax = A\bar{x}\} \equiv \{\bar{x} + v : v \in \mathcal{N}(A)\}$, (20) holds at $\bar{x}$ if and only if $\|\bar{x} + v\|_1 > \|\bar{x}\|_1, \ \forall\, v \in \mathcal{N}(A) \setminus \{0\}$. As is proved in the first part of Lemma 1, (21) is a sufficient condition for (20) with $\alpha = \mathrm{supp}(\bar{x})$. It is also necessary because given any $v \in \mathcal{N}(A)$ the triangle inequality (16) always holds as an equality at some $\bar{x}$ with $\mathrm{supp}(\bar{x}) = \alpha$. This is trivially true at any $v_\alpha = 0$; otherwise, $\bar{x}_\alpha = -v_\alpha$ makes (16) to hold as an equality. $\qquad\square$

We already know from Proposition 1 that if k satisfies

$$1 \leq \sqrt{k} < \min\left\{\frac{1}{2}\frac{\|v\|_1}{\|v\|_2} : \ v \in \mathcal{N}(A) \setminus \{0\}\right\},$$

then $\bar{x}$ in (20) is also the sparsest point in the set $\{x : Ax = A\bar{x}\}$.

3   Recoverability

Let us restate the equivalence (9) in a more explicit form:

$$\arg \min_{x \in \Omega}\{\|x\|_0 : Ax = A\bar{x}\} = \{\bar{x}\} = \arg \min_{x \in \Omega}\{\|x\|_1 : Ax = A\bar{x}\}, \tag{22}$$

where Ω can be any nonempty and closed subset of $\mathbb{R}^n$. When Ω is a convex set, the problem on the right-hand side is a convex program that can be efficiently solved at least in theory.

Recoverability addresses conditions under which the equivalence (22) holds. The conditions involved include the properties of the measurement matrix A and the degree of sparsity in signal $\bar{x}$. Clearly, the prior information set Ω can also affect recoverability. However, since we allow Ω to be any subset of $\mathbb{R}^n$, the results we obtain in this paper are the "worst-case" results in terms of varying Ω.

3.1 Kashin-Garnaev-Gluskin Inequality

We will make use of a classic result established in the late 1970's and early 1980's by Kashin [20], and Garnaev and Gluskin [16] in a particular form, which provides a lower bound on the ratio of the ℓ_1-norm to the ℓ_2-norm when it is restricted to subspaces of a given dimension. We know that in the entire space $\mathbb{R}^n$, the ratio can vary from 1 to $\sqrt{n}$, namely,

$$1 \leq \frac{\|v\|_1}{\|v\|_2} \leq \sqrt{n}, \quad \forall v \in \mathbb{R}^n \setminus \{0\}.$$

Roughly speaking, this ratio is small for sparse vectors that have many zero (or near-zero) elements. However, it turns out that in most subspaces this ratio can have much larger lower bounds than 1. In other words, most subspaces do not contain excessively sparse vectors.

For $p < n$, let $\mathbf{G}(n, p)$ denote the set of all p-dimensional subspaces of $\mathbb{R}^n$ (which is called a Grassmannian). It is known that there exists a unique rotation-invariant probability measure, say $\mathbf{Prob}(\cdot)$, on $\mathbf{G}(n, p)$ (see [23], for example). From our perspective, we will bypass the technical details on how such a probability measure is defined. Instead, it suffices to just mention that drawing a member from $\mathbf{G}(n, p)$ uniformly at random amounts to generating an n by p random matrix of *iid* entries from the standard normal distribution $\mathcal{N}(0, 1)$ whose range space will be a random member of $\mathbf{G}(n, p)$ (see Sec. (3.5) of [2] for a concise introduction).

Theorem 1 (Kashin, Garnaev and Gluskin). *For any two natural numbers $m < n$, there exists a set of $(n - m)$-dimensional subspaces of $\mathbb{R}^n$, say $\mathbb{S} \subset \mathbf{G}(n, n - m)$, such that*

$$\mathbf{Prob}\,(\mathbb{S}) \geq 1 - e^{-c_0(n-m)}, \tag{23}$$

and for any $\mathcal{V} \in \mathbb{S}$,

$$\frac{\|v\|_1}{\|v\|_2} \geq \frac{c_1 \sqrt{m}}{\sqrt{1 + \log(n/m)}}, \quad \forall v \in \mathcal{V} \setminus \{0\}, \tag{24}$$

where the constants $c_i > 0$, $i = 0, 1$, are independent of the dimensions.

This theorem ensures that a subspace $\mathcal{V}$ drawn from $\mathbf{G}(n, n-m)$ at random will satisfy inequality (24) with high probability when $n - m$ is large. From here on, we will call (24) the KGG inequality.

When $m < n$, the orthogonal complement of each and every member of $\mathbf{G}(n, m)$ is a member of $\mathbf{G}(n, n - m)$, establishing a one-to-one correspondence between $\mathbf{G}(n, m)$ and $\mathbf{G}(n, n - m)$. As a result, if $A \in \mathbb{R}^{m \times n}$ is a random matrix with *iid* standard normal entries, then the range space of A^T is a uniformly random member of $\mathbf{G}(n, m)$, while at the same time the null space of A is a uniformly random member of $\mathbf{G}(n, n-m)$ in which the KGG inequality holds with high probability when $n - m$ is large.

Remark 2. *Theorem 1 contains only a part of the seminal result on the so-called Kolmogorov or Gelfand width in approximation theory, first obtained by Kashin [20] in a weaker form and later improved to the current order by Garnaev and Gluskin [16]. In its original form, the result does not explicitly state the probability estimate, but gives both lower and upper bounds of the same order*

for the involved Kolmogorov width (see [10] for more information on the Kolmogorov and Gelfand widths and their duality in a case of interest). Theorem 1 is formulated after a description given by Gluskin and Milman in [19] (see the paragraph around inequality (5), i.e., the KGG inequality (24), on page 133). As will be seen, this particular form of the KGG result enables a greatly simplified CS analysis.

The connections between compressive sensing and the works of Kashin [20] and Garnaev and Gluskin [16] are well known. Candés and Tao [4] used the KGG result to establish that the order of stable recovery is optimal in a sense, though they derived their stable recovery results via an RIP-based analysis. In [10] Donoho pointed out that the KGG result implies an abstract stable recovery result (Theorem 1 in [10] for the case of $p = 1$) in an "indirect manner". In both cases, the authors used the original form of the KGG result, but in terms of the Gelfand width.

The approach taken in this paper is based on examining the ratio $\|v\|_1/\|v\|_2$ (or a variant of it in the case of uniform recoverability) in the null space of A, while relying on the KGG result in the form of Theorem 1 to supply the order of recoverable sparsity and the success probability in a direct manner. This approach was used in [29] to obtain a rather limited recoverability result. In this paper, we also use it to derive stability and uniform recoverability results.

3.2 An Extended Recoverability Result

The following recoverability result follows directly from the sufficient condition (17) and the KGG inequality (24) for $\mathcal{V} = \mathcal{N}(A)$ (also see the discussions before and after Theorem 1). It extends the result in [5] from $\Omega = \mathbb{R}^n$ to any $\Omega \subset \mathbb{R}^n$. We call a random matrix *standard normal* if it has *iid* entries drawn from the standard normal distribution $\mathcal{N}(0,1)$.

Theorem 2. *Let $m < n$, $\Omega \subset \mathbb{R}^n$, and $A \in \mathbb{R}^{m \times n}$ be either standard normal or any rank-m matrix such that $BA^T = 0$ where $B \in \mathbb{R}^{(n-m) \times m}$ is standard normal. Then with probability greater than $1 - e^{-c_0(n-m)}$, the equivalence (22) holds at $\bar{x} \in \Omega$ if the sparsity of $\bar{x}$ satisfies*

$$\|\bar{x}\|_0 < \frac{c_1^2}{4} \frac{m}{1 + \log(n/m)}, \tag{25}$$

where $c_0, c_1 > 0$ are some absolute constants independent of the dimensions m and n.

An oft-missed subtlety is that recoverability is entirely dependent on the properties of a subspace but independent of its representations. In Theorem 2, we have added a second case for A, where A^T can be any basis matrix for the null space of B, to bring attention to this subtle point.

Remark 3. *For the case $\Omega = \mathbb{R}^n$, the sparsity bound given in (25) is first established by Candés and Tao [5] for standard normal random matrices. In Baraniuk et al. [1], the same order is extended to a few other random matrices such as Bernoulli matrices whose entries are ± 1. Weaker results have been established for partial Fourier and other partial orthonormal matrices with random rows by Candés, Romberg and Tao [7], and Rudelson and Vershynin [26]. Moreover, an in-depth study on the asymptotic form for the constant in (25) can be found in the work of Donoho and Tanner [12].*

We emphasize that in terms of Ω the sparsity order in the right-hand side of (25) is a worst-case lower bound for (17) since it actually bounds (18) from below. In principle, larger lower bounds may exist for (17) corresponding to certain prior-information sets $\Omega \neq \mathbb{R}^n$, though specific investigations will be necessary to obtain such better lower bounds. For Ω equal to the nonnegative orthant of $\mathbb{R}^n$, a better bound has been obtained by Donoho and Tanner [13] (also see [30] for an alternative proof).

3.3   Why is the Extension Useful

The extended recoverability result gives a theoretical guarantee that adding prior information can never hurt but possibly enhance recoverability. The extended decoding model (3) includes many useful variations for the prior information set Ω.

For example, a sparse signal $\bar{x}$ under construction is often known to be close to a known "prior signal" x_p; say that $\bar{x}$ is an magnetic resonance image (MRI) taken from a patient today, while x_p is the one taken a week earlier from the same person. Given the closeness of the two, we may consider solving the model:

$$\min\{\|x\|_1 : Ax = b, \|x - x_p\|_1 \leq \delta\},$$

which has a prior-information set $\Omega = \{x : \|x - x_p\|_1 \leq \delta\}$ for some $\delta > 0$. When $\delta > \|\bar{x} - x_p\|_1$, adding the prior information will not raise the lower bound on recoverable sparsity as is given in (18), but nevertheless it will still help control errors in recovery, which arguably is more important in practice. For an appropriate $\mu > 0$, the above model is equivalent to

$$\min\{\|x\|_1 + \mu\|x - x_p\|_1 : Ax = b\},$$

which has a "mixed-norm" objective.

As another example, consider a two-dimensional image $\bar{x}$ known to contain blocky regions of almost constant values. In this case, the total variation (TV) of $\bar{x}$ should be small, which is defined as $\mathrm{TV}(\bar{x}) \triangleq \sum_i \|D_i \bar{x}\|_2$ where $D_i \bar{x} \in \mathbb{R}^2$ is a finite-difference gradient approximation of $\bar{x}$ evaluated at the pixel i. This prior information is characterized by the set $\Omega = \{x : \mathrm{TV}(x) \leq \delta\}$ for some $\delta > 0$, and leads to the model

$$\min\{\|x\|_1 : Ax = b, \mathrm{TV}(x) \leq \delta\}.$$

For an appropriate $\mu > 0$, the above model is equivalent to

$$\min\{\|x\|_1 + \mu\mathrm{TV}(x) : Ax = b\}. \tag{26}$$

We now present a computational example to illustrate that prior-information can indeed significantly raise the recoverable sparsity level. Consider the first two images on the left side of Figure 1, called images 1 and 2, where image 1 is supposed to be from full-body magnetic resonance imaging. Both images are 1281×400 in size, with pixel values between 0 and 255. Image 2 is the reverse of image 1 whose pixel values are obtained by subtracting those of image 1 from 255, so white becomes black and vice versa.
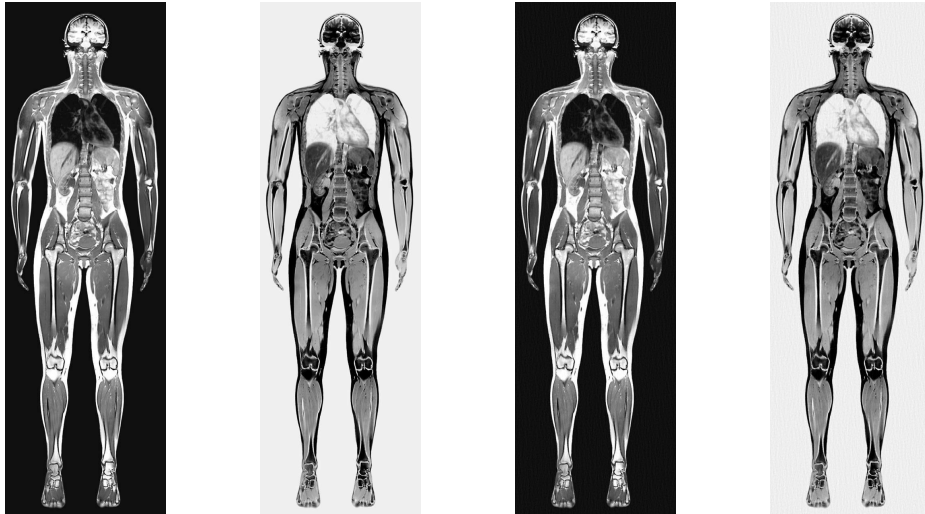
Figure 1: Simulations of compressive sensing with prior information. From left to right, images 3 and 4 were reconstructed from 20% of the Fourier coefficients of images 1 and 2, respectively, via model (26).

Image 1 is approximately but not highly sparse. The background pixel values of image 1 are 16 instead of 0 (even though the background appears black). About 44% of pixels in image 1 have pixel values greater than 32. Image 2 is much less sparse with over 90% of pixels having pixel values greater than 32. However, both images have blocky structures and hence small TV values. This prior information makes model (26) suitable for recovering these two images from under-sampled measurements. For the measurement matrix A in (26), we used partial Fourier matrices, each formed by one fifth of the rows of a two-dimensional discrete Fourier matrix. Specifically, a measurement vector b consisted of 20% of the Fourier coefficients of either image 1 or 2. These Fourier coefficients were randomly selected but biased towards those associated with lower frequency basis vectors.

We approximately solved model (26) and recovered images 3 and 4 from 20% of the Fourier coefficients of images 1 and 2, respectively. These two recovered images appear almost identical to their corresponding originals, though slight quality degradations can be found upon close examination. Without using the prior information (i.e., the TV term in model (26)), we have found that the quality of the recovered image 4 would be significantly inferior.

Since the pixel values are real while the Fourier coefficients are complex-valued, in theory 50% of the Fourier coefficients are enough for exact recovery. The use of only 20% of the coefficients represents a 60% reduction in the amount of data required for recovery. In fact, simulations in this example indicates how CS may be applied to MRI where scanned data are essentially Fourier coefficients of images under construction. A 60% reduction in MRI scanning time would represent significant improvements in MRI practice. We refer the reader to the work of Lustig, Donoho, Santos and Pauly [22], and the references therein, for more information on the applications of CS to MRI.

13

4 Stability

In practice, a measurement vector b is most likely inaccurate due to various sources of imprecisions and/or contaminations. A more realistic measurement model should be $b = A\bar{x} + r$ where $\bar{x}$ is a desirable and sparse signal, and $r \in \mathbb{R}^m$ can be interpreted either as noise or as spurious measurements generated by a noise-like signal. Now given an under-sampled and imprecise measurement vector $b = A\bar{x} + r$, can we still approximately recover the sparse signal $\bar{x}$? What kind of errors should we expect? These are questions our stability analysis should address.

4.1 Preliminaries

Assume that $A \in \mathbb{R}^{m \times n}$ is of rank m and $b = A\bar{x} + r$, where $\bar{x} \in \Omega \subset \mathbb{R}^n$. Since the sparse signal of interest, $\bar{x}$, does not satisfy the equation $Ax = b$, we relax our goal to the inequality $\|Ax - b\| \leq \delta$ in some norm $\|\cdot\|$, with $\delta \geq \|r\|$ so that $\bar{x}$ satisfies the inequality. An alternative interpretation for the imprecise measurement b is that it is generated by a signal $\hat{x}$ that is approximately sparse so that $b = A\hat{x}$ for $\hat{x} = \bar{x} + p$ where p is small and satisfies $Ap = r$.

Consider the QR-factorization: $A^T = UR$, where $U \in \mathbb{R}^{n \times m}$ satisfies $U^T U = I$ and $R \in \mathbb{R}^{m \times m}$ is upper triangular. Obviously, $Ax = b$ is equivalent to $U^T x = R^{-T} b$, and

$$\|Ax - b\|_M = \|U^T x - R^{-T} b\|_2,$$

where

$$\|q\|_M = (q^T M q)^{1/2} \quad \text{and} \quad M = (R^T R)^{-1}. \tag{27}$$

We will make use of the following two projection matrices:

$$P_r = A^T (AA^T)^{-1} A \equiv UU^T \quad \text{and} \quad P_n = I - P_r, \tag{28}$$

where P_r is the projection onto the range space of A^T and P_n onto the null space of A. In addition, we will use the constant

$$C_\nu = \frac{1 + \nu\sqrt{2 - \nu^2}}{1 - \nu^2} > 1, \quad \nu \in (0, 1). \tag{29}$$

It is easy to see that as ν approaches 1, $C_\nu \approx 1/(1 - \nu)$. For example, $C_\nu \approx 2.22$ for $\nu = 0.5$, $C_\nu \approx 4.34$ for $\nu = 0.75$ and $C_\nu \approx 10.43$ for $\nu = 0.9$.

4.2 Two Stability Results

Given the imprecise measurement $b = A\bar{x} + r$, consider the decoding model

$$x_\delta^* = \arg\min_{x \in \Omega}\{\|x\|_1 : \|Ax - b\|_M \leq \delta\} = \arg\min_{x \in \mathcal{F}(\delta)} \|x\|_1, \tag{30}$$

where $\Omega \subset \mathbb{R}^n$ is a closed, prior-information set, $\delta \geq \|r\|_2$, the weighted norm is defined in (27), and

$$\mathcal{F}(\delta) = \{x : \|Ax - b\|_M \leq \delta, x \in \Omega\}. \tag{31}$$

In general, $x_\delta^* \neq \bar{x}$ and is not strictly sparse. We will show that x_δ^* is close to $\bar{x}$ under suitable conditions. To our knowledge, stability of this model has not been previously investigated. Our result below says that if $\mathcal{F}(\delta)$ contains a sufficiently sparse point $\bar{x}$, then the distance between x_δ^* and $\bar{x}$ is of order δ. Consequently, If $\bar{x} \in \mathcal{F}(0)$, then $x_0^* = \bar{x}$.

In our stability analysis below, we will make use of the following sparsity condition:

$$ k = \frac{\nu^2}{4} \left(\frac{\|u\|_1}{\|u\|_2} \right)^2, \quad \text{for some } \nu \in (0, 1). \tag{32} $$

Theorem 3. *Let $\delta \geq \|A\bar{x} - b\|_M$ for some $\bar{x} \in \Omega$. Assume that $k = \|\bar{x}\|_0$ satisfies (32) for $u = P_n(x_\delta^* - \bar{x})$ whenever $P_n(x_\delta^* - \bar{x}) \neq 0$. Then for $p = 1$ or 2*

$$ \|x_\delta^* - \bar{x}\|_p \leq \gamma_p(C_\nu + 1)(\|A\bar{x} - b\|_M + \delta), \tag{33} $$

where $\gamma_1 = \sqrt{n}$, $\gamma_2 = 1$ and C_ν is defined in (29).

Remark 4. *We quickly add that if A is a standard normal random matrix, then the KGG inequality (24) implies that with high probability the right-hand side of (32) for $u = P_n(x_\delta^* - \bar{x})$ is at least of the order $m/\log(n/m)$. This same comment also applies to the next theorem.*

The proof of this theorem, as well as that of Theorem 4 below, will be given in Subsection 4.4. Next we consider the special case $\delta = 0$, namely,

$$ x_0^* = \arg \min_{x \in \Omega} \{\|x\|_1 : Ax = b\}. $$

A k-term approximation of $x \in \mathbb{R}^n$, denoted by $x(k)$, is obtained from x by setting its $n - k$ smallest elements in magnitude to zero. Obviously, $\|x\|_1 \equiv \|x(k)\|_1 + \|x - x(k)\|_1$. Due to measurement errors, there may be no sparse signal that satisfies the equation $Ax = b$. In this case, we show that if the observation b is generated by a signal $\hat{x}$ that has a good, k-term approximation $\hat{x}(k)$, then the distance between x_0^* and $\hat{x}$ is bounded by the error in the k-term approximation. Consequently, if $\hat{x}$ is itself k-sparse so that $\hat{x} = \hat{x}(k)$, then $x_0^* = \hat{x}$.

Theorem 4. *Let $\hat{x} \in \Omega$ satisfy $Ax = b$ and $\hat{x}(k)$ be a k-term approximation of $\hat{x}$. Assume that k satisfies*

$$ \|\hat{x} - \hat{x}(k)\|_1 \leq \|\hat{x}\|_1 - \|x_0^*\|_1 \tag{34} $$

and (32) for $u = P_n(x_0^* - \hat{x}(k))$ whenever $P_n(x_0^* - \hat{x}(k)) \neq 0$. Then for $p = 1$ or 2

$$ \|x_0^* - \hat{x}(k)\|_p \leq (C_\nu + 1)\|P_r(\hat{x} - \hat{x}(k))\|_p, \tag{35} $$

where C_ν is defined in (29).

It follows from (35) and the triangle inequality that for $p = 1$ or 2

$$ \|x_0^* - \hat{x}\|_p \leq (C_\nu + 1)\|P_r(\hat{x} - \hat{x}(k))\|_p + \|\hat{x} - \hat{x}(k)\|_p, \tag{36} $$

which has the same type of right-hand side as those in (4) and (5).

15

Condition (34) can always be met for k sufficiently large, while condition (32) demands that k be sufficiently small. Together, the two require that the measurement b be observed from a signal $\hat{x}$ that has a sufficiently good k-term approximation for a sufficiently small k. This requirement seems very reasonable.

In the special case of $\Omega = \mathbb{R}^n$, the error bounds in Theorems 3 and 4 bear similarities with existing stability results by Candés, Romberg and Tao [6] (see (4)–(6) in Section 1 and also results in [9]), but substantial differences exist in the the norms used in the two sides, the constants involved, and the conditions required. Overall, Theorems 3 and 4 do not contain, nor are contained in, the existing results, though they all state the same fact that CS recovery based on ℓ_1-minimization is stable to some extent.

4.3 RIP Issue Revisited

A main difference between the existing stability results and ours lies in the constants involved. In our error bound (35) the constant is given by an explicit formula depending only on the number $\nu \in (0,1)$ representing the relative sparsity level of the signal under construction. Relatively speaking, the sparser the signal is, the smaller the constant is, and the more stable the recovery is supposed to be.

On the other hand, the constants in the existing stability results depend on RIP parameters of the matrix A. The better the RIP parameters, the smaller those constants. However, does a smaller RIP-dependent constant imply a more stable recovery?

Given a measurement matrix A and a signal $\bar{x}$, which is either exactly or approximately sparse, suppose that we try to recover $\bar{x}$ by solving

$$\min\{\|x\|_1 : GAx = GA\bar{x}\} \tag{37}$$

with a varying nonsingular matrix $G \in \mathbb{R}^{m \times m}$. Under the assumption of exact arithmetics, it is obvious that both recoverability and stability should remain exactly the same as long as G stays nonsingular. Indeed, this is what we have attained in our recoverability and stability results, which depend only on the subspaces associated with A while independent of matrix representations. This is not the case, however, with the RIP-based results since, by definition (7), the RIP parameters are matrix-dependent. This matrix dependency (including that for the results in [10]) suggests, falsely, that stability of the decoding model (37) should vary with G.

On the other hand, Theorem 4 requires condition (34), which is not required by RIP-based results. The relative strengths and weaknesses of the available stability results remain an issue for further investigation.

4.4 Proofs of Stability Results

Our stability results follow directly from the following simple lemma.

Lemma 2. *Let $x, y \in \mathbb{R}^n$ such that $\|y\|_1 \le \|x\|_1$, and let $y - x = u + w$. where $u^T w = 0$. Whenever $u \ne 0$, assume that $k = \|x\|_0$ satisfies (32). Then for $p = 1$ or 2*

$$\|u\|_p \le C_\nu \|w\|_p, \tag{38}$$

16

$$\|y - x\|_p \leq (C_\nu + 1)\|w\|_p. \tag{39}$$

Proof. If $u = 0$, both (38) and (39) are trivially true. If $u \neq 0$, condition (32) and the assumption $\|y\|_1 \leq \|x\|_1$ imply that $w \neq 0$; otherwise, by Lemma 1, (32) would imply $\|y\|_1 > \|x\|_1$. For $u \neq 0$ and $w \neq 0$,

$$\frac{\|u + w\|_1}{\|u + w\|_2} \geq \left(\frac{1 - \|w\|_1/\|u\|_1}{\sqrt{1 + (\|w\|_2/\|u\|_2)^2}} \right) \frac{\|u\|_1}{\|u\|_2},$$

which follows from the triangle inequality $\|u + w\|_1 \geq \|u\|_1 - \|w\|_1$. Furthermore,

$$\frac{\|u + w\|_1}{\|u + w\|_2} \geq \left(\frac{1 - \eta(u, w)}{\sqrt{1 + \eta(u, w)^2}} \right) \frac{\|u\|_1}{\|u\|_2} \triangleq \phi(\eta(u, w)) \frac{\|u\|_1}{\|u\|_2}, \tag{40}$$

where $\phi(t) = (1 - t)/\sqrt{1 + t^2}$ and

$$\eta(u, w) \triangleq \max \left\{ \frac{\|w\|_1}{\|u\|_1}, \frac{\|w\|_2}{\|u\|_2} \right\}. \tag{41}$$

If $1 - \eta(u, w) \leq 0$, then (38) trivially holds. Therefore, we assume that $\eta(u, w) < 1$.

If $\phi(\eta(u, w)) > \nu$, then it follows from (32) and (40) that

$$\sqrt{\|x\|_0} \leq \frac{\nu}{2} \frac{\|u\|_1}{\|u\|_2} \leq \frac{\nu/\phi(\eta(u, w))}{2} \frac{\|u + w\|_1}{\|u + w\|_2} < \frac{1}{2} \frac{\|u + w\|_1}{\|u + w\|_2},$$

which would imply $\|x\|_1 < \|y\|_1$ by Lemma 1, contradicting the assumption of the lemma. Therefore, $\phi(\eta(u, w)) \leq \nu$ must hold. It is easy to verify that

$$\phi(t) = \frac{1 - t}{\sqrt{1 + t^2}} \leq \nu \quad \text{and} \quad t < 1 \iff \frac{1}{C_\nu} \leq t < 1,$$

where $1/C_\nu$ is the root of the quadratic $q(t) \triangleq (1 - t)^2 - \nu^2(1 + t^2)$ that is smaller than 1 (noting that $\phi(t) \leq \nu$ is equivalent to $q(t) \leq 0$ for $t < 1$). We conclude that there must hold $\eta(u, w) \geq 1/C_\nu$, which implies (38) in view of the definition of $\eta(u, w)$ in (41). Finally, (39) follows directly from the relationship $y - x = u + w$, the triangle inequality, and (38). $\square$

Clearly, whether the estimates of the lemma hold for $p = 1$ or $p = 2$ depends on which ratio is larger in (41); or equivalently, which ratio is larger between $\|u\|_1/\|u\|_2$ and $\|w\|_1/\|w\|_2$. If u is from a random, $(n - m)$-dimensional subspace of $\mathbb{R}^n$, then w is from its orthogonal complement — a random, m-dimensional subspace. When $n - m \gg m$, the KGG result indicates that it is more likely that $\|u\|_1/\|u\|_2 < \|w\|_1/\|w\|_2$; or equivalently, $\|w\|_2/\|u\|_2 < \|w\|_1/\|u\|_1$. In this case, $p = 1$ is more likely than $p = 2$.

Corollary 1. *Let $U \in \mathbb{R}^{n \times m}$ with $m < n$ have orthonormal columns so that $U^T U = I$. Let $x, y \in \mathbb{R}^n$ satisfy $\|y\|_1 \leq \|x\|_1$ and $\|U^T y - d\|_2 \leq \delta$ for some $d \in \mathbb{R}^m$. Define $u \triangleq (I - UU^T)(y - x)$ and $w \triangleq UU^T(y - x)$. In addition, let $k = \|x\|_0$ satisfy (32) whenever $u \neq 0$. Then*

$$\|y - x\|_p \leq \gamma_p(C_\nu + 1)(\|U^T x - d\|_2 + \delta), \quad p = 1 \text{ or } 2, \tag{42}$$

where $\gamma_1 = \sqrt{n}$ and $\gamma_2 = 1$.

Proof. Noting that $y = x + u + w$ and $U^T u = 0$, we calculate

$$\delta \geq \|U^T(x + u + w) - d\|_2 = \|U^T w - (d - U^T x)\|_2 \geq \|U^T w\|_2 - \|U^T x - d\|_2,$$

which implies

$$\|w\|_2 = \|U^T w\|_2 \leq \delta + \|U^T x - d\|_2. \tag{43}$$

Combining (43) with (39), we arrive at (42) for either $p = 1$ or 2, where in the case of $p = 1$ we use the inequality $\|w\|_1 \leq \sqrt{n}\|w\|_2$. $\qquad\square$

Proof of Theorem 3

Proof. Since $\bar{x} \in \mathcal{F}(\delta)$ and x_δ^* minimizes the 1-norm in $\mathcal{F}(\delta)$, we have $\|x_\delta^*\|_1 \leq \|\bar{x}\|_1$. The proof then follows from applying Corollary 1 to $y = x_\delta^*$ and $x = \bar{x}$, and the fact that the weighted norm defined in (27) satisfies $\|Ax - b\|_M = \|U^T x - R^{-T} b\|_2$. $\qquad\square$

Proof of Theorem 4

Proof. We note that condition (34) is equivalent to $\|x_0^*\|_1 \leq \|\hat{x}(k)\|_1$. Upon applying Lemma 2 to $y = x_0^*$ and $x = \hat{x}(k)$ with $u = P_n(y - x)$ and $w = P_r(y - x)$, and also noting $P_r x_0^* = P_r \hat{x}$, we have

$$\|x_0^* - \hat{x}(k)\|_p \leq (C_\nu + 1)\|P_r(\hat{x}(k) - \hat{x})\|_p,$$

which completes the proof. $\qquad\square$

5 Uniform Recoverability

The recoverability result, Theorem 2, is derived only for standard normal random matrices. It has been empirically observed (see [14, 15], for example) that recovery behavior of many different random matrices seems to be identical. We call this phenomenon *uniform recoverability*. In this section, we provide a theoretical explanation to this property.

5.1 Preliminaries

We will consider the simple case where $\Omega = \mathbb{R}^n$ and $\delta = 0$ so that we can make use of the necessary and sufficient condition in Proposition 2. We first translate this necessary and sufficient condition into a form more conducive to our analysis.

For $0 < k < n$, we define the following function that maps a vector (of any size n) to a scalar:

$$\lambda_k(v) \triangleq \frac{\|v\|_1 - 2\|v(k)\|_1}{\|v\|_2}, \quad v \neq 0, \tag{44}$$

where $v(k)$ is a k-term approximation of v whose nonzero elements are the k largest elements of v in magnitude. It is important to note that $\lambda_k(v)$ is invariant with respect to multiplications by scalars (or scale-invariant) and permutations of v. Moreover, $\lambda_k(v)$ is continuous, and achieves its minimum and maximum (since its domain can be restricted to the unit sphere).

Proposition 3. *Given* $A \in \mathbb{R}^{m \times n}$ *with* $m < n$, *the equivalence (20), i.e.,*

$$\{\bar{x}\} = \arg \min\{\|x\|_1 : Ax = A\bar{x}\},$$

holds for all $\bar{x}$ *with* $\|\bar{x}\|_0 \leq k$ *if and only if*

$$0 < \Lambda_k(A) \triangleq \min\{\lambda_k(v) : v \in \mathcal{N}(A) \setminus \{0\}\}. \tag{45}$$

Proof. For any fixed $v \neq 0$, the condition $\|v\|_1 > 2\|v_\alpha\|_1$ in (21) for all index sets α with $|\alpha| \leq k$ is clearly equivalent to $\lambda_k(v) > 0$. After taking the minimum over all $v \neq 0$ in $\mathcal{N}(A)$, we see that $\Lambda_k(A) > 0$ is equivalent to the necessary and sufficient condition in Proposition 2. $\square$

For notational convenience, given any $A \in \mathbb{R}^{m \times n}$ with $m < n$, let us define the set

$$\mathbf{sub}(A) = \left\{ B \in \mathbb{R}^{m \times (m+1)} : B \text{ is a submatrix of } A \right\}.$$

In other words, each member of $\mathbf{sub}(A)$ is formed by $m + 1$ columns of A in their original order. Clearly, the cardinality of $\mathbf{sub}(A)$ is n choose $m + 1$. We say that the set $\mathbf{sub}(A)$ has full rank if every member of $\mathbf{sub}(A)$ has rank m. It is well known that for most distributions, $\mathbf{sub}(A)$ will have full rank with high probability.

5.2 Results for Uniform Recoverability

When $A \in \mathbb{R}^{m \times n}$ is randomly chosen from a probability space, $\Lambda_k(A)$ is a random variable whose sign, according to Proposition 3, determines the success or failure of recovery for all $\bar{x}$ with $\|\bar{x}\|_0 \leq k$. In this setting, the following theorem indicates that $\Lambda_k(A)$ is a sample minimum of another random variable $\lambda_k(d(B))$ where $B \in \mathbb{R}^{m \times (m+1)}$ is from the same probability space as A, and $d(B) \in \mathbb{R}^{m+1}$ is defined by

$$[d(B)]_i \triangleq |\det(B_i)|, \quad i = 1, \cdots, m + 1, \tag{46}$$

where $B_i \in \mathbb{R}^{m \times m}$ is the submatrix of B obtained by deleting the i-th column of B.

Theorem 5. *Let* $A \in \mathbb{R}^{m \times n}$ *($m < n$) with* $\mathbf{sub}(A)$ *of full rank (i.e., every member of* $\mathbf{sub}(A)$ *has rank m). Then for* $k \leq m$

$$\Lambda_k(A) = \min\left\{\lambda_k(d(B)) : B \in \mathbf{sub}(A)\right\}, \tag{47}$$

where $d(B) \in \mathbb{R}^{m+1}$ *is defined in (46).*

The proof of this theorem is left to Subsection 5.4.

Remark 5. *Theorem 5, together with Proposition 3, establishes that recoverability is determined by the properties of $d(B)$, not directly those of A. If distributions of $d(B)$ for different types of random matrices converge to the same limit as $m \to \infty$, then asymptotically there should be an identical recovery behavior for different types of random matrices.*

For any random matrix B, by definition the components of $d(B)$ are random determinants (in absolute value). It has been established by Girko [17] that a wide class of random determinants (in absolute value) does share a limit distribution (see the book by Girko [18] for earlier results).

Theorem 6 (Girko). *For any m, let the random elements $t_{ij}^{(m)}$, $1 \le i, j \le m$, of the matrix $T_m = [t_{ij}^{(m)}]$ be independent, $\mathbf{E}(t_{ij}^{(m)}) = \mu$, $\mathbf{Var}(t_{ij}^{(m)}) = 1$ and $\sup_{i,j,m} \mathbf{E}|t_{ij}^{(m)}|^{4+\delta} < \infty$ for some $\delta > 0$. Then*

$$\lim_{m \to \infty} \mathbf{Prob}\left\{ \frac{\log\left(\det(T_m)^2/[(m-1)!(1+m\mu^2)]\right)}{\sqrt{2\log m}} < t \right\} = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{t} e^{-\frac{x^2}{2}} dx. \qquad (48)$$

Theorem 6 says that for a wide class of random determinants squared, their logarithms, with proper scalings, converge in distribution to the standard normal law; or the limit distribution of the random determinants squared is log-normal.

Remark 6. *Since the elements of $d(B)$ are random determinants in absolute value, they all have the same limit distribution as long as B satisfies the conditions of Theorem 6.*

The elements of $d(B)$ are not independent in limit, because any two of them are determinants of $m \times m$ matrices that share $m-1$ common columns. However, the dependency among the elements of $d(B)$ is purely algebraic rather than stochastic, and hence does not vary with the type of random matrices. To stress this point, we mention that for a wide range of random matrices B, the ratios $\det(B_i)/\det(B_j)$, $i \ne j$, converge to Cauchy distribution with the cumulative distribution function $1/2 + \arctan(t)/\pi$. This result can be found, in a slightly different form, in Theorem 15.1.1 of the book by Girko [18].

Remark 7. *We observe from (48) that the mean values μ only affect the scaling factor of the determinant, but not the asymptotic recoverability behavior since $\lambda_k(\cdot)$ is scale-invariant, implying that measurement matrices need not have zero mean, as is usually required in earlier theoretical results of this sort. In addition, the unit variance assumption in Theorem 6 is not restrictive because it can always be achieved by scaling.*

5.3 Numerical Illustration

To illustrate the uniformity of CS recovery, we sample the random variable $\lambda_k(d(B))$ for $B \in \mathbb{R}^{m \times (m+1)}$ whose entries are *iid* and randomly drawn from one of the two probability distributions: the standard normal distribution $\mathcal{N}(0,1)$ or the uniform distribution on the interval $[0,1]$. While the former has zero mean, the latter has mean $1/2$. In Figure 2, we plot the empirical density functions (namely, scaled histograms) of $\lambda_k(d(B))$ for the two random distributions with different values of m, k and sample size. Recall that successful recovery of all k-sparse signals is guaranteed for matrix A if and only if the sample minimum of $\lambda_k(d(B))$ over $\mathbf{sub}(A)$ is positive.

As can be seen from Figure 2, even at $m = 50$, the two empirical density functions for $\lambda_{10}(d(B))$, corresponding to the standard normal (solid line) and the uniform distributions (small circles) respectively, are already fairly close. At $m = 200$, the two empirical density functions for $\lambda_{40}(d(B))$ become largely indistinguishable in most places.
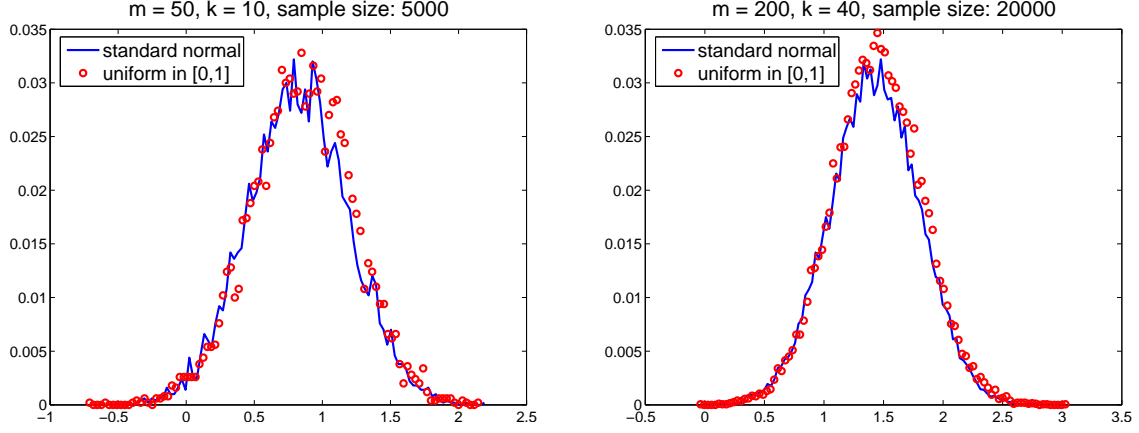
Figure 2: Empirical density functions of $\lambda_k(d(B))$ for 2 random distributions

5.4 Proof of Theorem 5

The following result, established by Eydelzon [15] in a slightly different form, will play a key role in the proof. For completeness, we include a proof for it.

Lemma 3. *Let $\mathcal{V}$ be an $(n-m)$-dimensional subspace of $\mathbb{R}^n$ and $k \leq m < n$. If $\hat{v} \in \mathcal{V}$ minimizes $\lambda_k(v)$ in $\mathcal{V}$, then $\hat{v}$ has at least $n - m - 1$ zeros, or equivalently at most $m + 1$ nonzeros.*

Proof. Let $\mathcal{V}$ be spanned by the orthonormal columns of $Q \in \mathbb{R}^{n \times (n-m)}$ (so $Q^T Q = I$) and $\hat{v} = Q\hat{s}$ minimizes $\lambda_k(v)$ in $\mathcal{V}$ for some $\hat{s} \in \mathbb{R}^{n-m}$. Assume, without loss of generality, that $\|\hat{v}\|_2 = \|Q\hat{s}\|_2 = \|\hat{s}\|_2 = 1$. We now prove the result by contradiction.

Suppose that $\hat{v}$ has at most $n - m - 2$ zeros, say, $\hat{v}_i = q_i^T \hat{s} = 0$ for $i = 1, 2, \cdots, n - m - 2$ where q_i^T is the i-th row of Q. Then there must exist a unit (in the 2-norm) vector $h \in \mathbb{R}^{n-m}$ that is perpendicular to $\hat{s}$ and q_i for $i = 1, 2, \cdots, n - m - 2$. By construction,

$$(\hat{v} + \tau Q h)_i = 0, \;\; i = 1, 2, \cdots, n - m - 2,$$

for any scalar value τ. By setting τ sufficiently small in absolute value, say $|\tau| \leq \epsilon$, we can ensure for $i > n - m - 2$ that $\text{sign}((\hat{v} + \tau Q h)_i) = \text{sign}(\hat{v}_i)$ so that

$$|(\hat{v} + \tau Q h)_i| = \text{sign}(\hat{v}_i)(\hat{v} + \tau Q h)_i = |\hat{v}_i| + \tau \, \text{sign}(\hat{v}_i)(Q h)_i, \; \forall i > n - m - 2.$$

Now we evaluate $\lambda_k(\cdot)$ at $v = \hat{v} + \tau Q h$ for $0 < |\tau| \leq \epsilon$ (with a yet undecided sign for τ),

$$
\begin{aligned}
\lambda_k(v) &= (\|v\|_1 - 2\|v(k)\|_1)/\sqrt{1 + \tau^2} \\
&= (\|\hat{v}\|_1 - 2\|\hat{v}(k)\|_1 + \tau \omega)/\sqrt{1 + \tau^2} \\
&= (\lambda_k(\hat{v}) + \tau \omega)/\sqrt{1 + \tau^2}
\end{aligned}
\tag{49}
$$

where $\sqrt{1 + \tau^2} = \|v\|_2 = \|\hat{s} + \tau h\|_2$ and for some index set J with $|J| = k$,

$$\omega = \sum_{i=n-m-1}^{n} \text{sign}(\hat{v}_i)(Q h)_i - 2 \sum_{i \in J} \text{sign}(\hat{v}_i)(Q h)_i.$$

21

If $\omega \neq 0$, we set $\mathrm{sign}(\tau) = -\mathrm{sign}(\omega)$ so that $\tau\omega < 0$. Now a contradiction, $\lambda_k(v) < \lambda_k(\hat{v})$, arises from (49). So $\hat{v}$ must have at least $n - m - 1$ zeros or at most $m + 1$ nonzeros. $\qquad\square$

Proof of Theorem 5

Proof. In view of Lemma 3 with $\mathcal{V} = \mathcal{N}(A)$, to find $\Lambda_k(A)$ it suffices to evaluate the minimum of $\lambda_k(\cdot)$ over all those vectors in $\mathcal{N}(A) \setminus \{0\}$ that have at most $m + 1$ nonzeros.

Without loss of generality, let $v \in \mathcal{N}(A) \setminus \{0\}$ so that $v_i = 0$ for all $i > m + 1$, and let $B = [b_1 \; b_2 \; \cdots \; b_{m+1}] \in \mathbf{sub}(A)$ consist of the first $m + 1$ columns of A. Then $Av = Bu = 0$, where u consists of the first $m + 1$ elements of v. Hence, $u \in \mathbb{R}^{m+1}$ spans the null space of B which is one-dimensional (recall that $\mathbf{sub}(A)$ has full rank).

Let B_i be the submatrix of B with its i-th column removed. Without loss of generality, we assume that $\det(B_1) \neq 0$. Clearly, the null space of B is spanned by the vector

$$u = \begin{pmatrix} -1 \\ B_1^{-1}b_1 \end{pmatrix} \in \mathbb{R}^{m+1},$$

where, by Crammer's rule, $(B_1^{-1}b_1)_i = \det(B_{i+1})/\det(B_1)$, $i = 1, 2, \cdots, m$. Since the function $\lambda_k(\cdot)$ is scale-invariant, to evaluate $\lambda_k(\cdot)$ at $v \in \mathcal{N}(A) \setminus \{0\}$ with $v_i = 0$ for $i > m + 1$, it suffices to evaluate it at $d(B) \triangleq |\det(B_1)u|$, which coincides with the definition in (46).

Obviously, the exactly same argument can be equally applied to all other nonzero vectors in $\mathcal{N}(A)$ which have at most $m + 1$ nonzeros in different locations, corresponding to different members of $\mathbf{sub}(A)$. This completes the proof. $\qquad\square$

6 Conclusions

CS is an emerging methodology with a solid theoretical foundation that is still evolving. Most previous analyses in the CS theory relied on the RIP of the measurement matrix A. These analyses can be called matrix-based. The non-RIP analysis presented in this paper, however, is subspace-based, and utilizes the classic KGG inequality to supply the order of recoverable sparsity. It should be clear from this non-RIP analysis that CS recoverability and stability are solely determined by the properties of the subspaces associated with A regardless of matrix representations.

The non-RIP approach used in this paper has enabled us to derive the extended recoverability and stability results immediately from a couple of remarkably simple observations (Lemmas 1 and 2) on the 2-norm versus 1-norm ratio in the null space of A. The obtained extensions include: (a) allowing the use of prior information in recovery, (b) establishing RIP-free formulas for stability constants, and (c) explaining the uniform recoverability phenomenon. In our view, these new results reinforce the theoretical advantages of ℓ_1-minimization-based CS decoding models.

As has been alluded to at the beginning, there are topics in the CS theory that are not covered in this work, one of which is that the recoverable sparsity order given in Theorem 2 can be shown to be optimal in a sense (see [3] for an argument). Nevertheless, it is hoped that this work will help enhance and enrich the theory of CS, make the theory more accessible, and stimulate more

activities in utilizing prior information and different measurement matrices in CS research and practice.

Acknowledgments

References

[1] R. Baraniuk, M. Davenport, R. DeVore and M. Wakin. A simple proof of the restricted isometry property for random matrices. To appear in Constructive Approximation. 2007.

[2] A. Barvinok. Math 710: Measure Concentration. Lecture notes, Department of Mathematics, University of Michigan, Ann Arbor, Michigan 48109-1109.

[3] E. Candès. Compressive sampling. International Congress of Mathematicians, Madrid, Spain, August 22-30, 2006 (Eur. Math. Soc., Zürich, 2006), Vol. 3, pp. 1433-1452.

[4] E. Candès and T. Tao. Near optimal signal recovery from random projections: universal encoding strategies. IEEE Transactions on Information Theory, 52 (2006), pp. 5406–5425.

[5] E. Candès and T. Tao. Decoding by linear programming. IEEE Transactions on Information Theory, Vol. 51, pp. 4203–4215, 2005.

[6] E. Candès, J. Romberg, and T. Tao, Stable signal recovery from incomplete and inaccurate information. Communications on Pure and Applied Mathematics, 2005 (2005), pp. 1207–1233.

[7] E. Candès, J. Romberg, and T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Inform. Theory 52 (2006), 489–509.

[8] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. SIAM J. Scientific Computing 20: 33-61, 1998.

[9] A. Cohen, W. Dahmen, and R. A. DeVore. Compressed sensing and k-term approximation. *Submitted.*, (2007).

[10] D. Donoho. Compressed sensing. IEEE Transactions on Information Theory, 52 (2006), pp. 1289–1306.

[11] D. Donoho and M. Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. Proc. Natl. Acad. Sci. U.S.A., 100(2003): 2197–2202.

[12] D. Donoho and J. Tanner. Counting faces of randomly-projected polytopes when the projection radically lowers dimension. Submitted to Journal of the AMS, (2005).

[13] D. Donoho and J. Tanner. Sparse Nonnegative Solutions of Underdetermined Linear Equations by Linear Programming. Proceedings of the National Academy of Sciences (USA), 2005 v102(27): 9446-9451.

[14] D. Donoho and Y. Tsaig. Extensions of compressed sensing. Signal Processing, 86(3), pp. 533-548, March 2006.

[15] Anatoly Eydelzon. A Study on Conditions for Sparse Solution Recovery in Compressive Sensing. PhD Thesis, Rice University, CAAM Technical Report TR07-12, (2007).

[16] A. Garnaev and E. D. Gluskin. The widths of a Euclidean ball. *Dokl. Akad. Nauk SSSR*, 277 (1984), pp. 1048–1052.

[17] V. L. Girko. A Refinement of the Central Limit Theorem for Random Determinants. Theory of Probability and its Applications. Vol. 42, No. 1, pp. 21-129. 1998, (translated from a Russian Journal).

[18] V. L. Girko. *Theory of Random Determinants*. Kluwer Academic Publishers, Dordrecht, Boston, London. 1990.

[19] E. Gluskin and V. Milman. Note on the Geometric-Arithmetic Mean Inequality. *Geometric Aspects of Functional Analysis Israel Seminar 2001-2002*. Lecture Notes in Mathematics, Springer, Berlin, Heidelberg. Vol.1807/2003, pp.131 - 135.

[20] B. S. Kashin. Diameters of certain finite-dimensional sets in classes of smooth functions. Izv. Akad. Nauk SSSR, Ser. Mat., 41 (1977), pp. 334–351.

[21] B. S. Kashin and V. N. Temlyakov. A Remark on Compressed Sensing. Mathematical Notes, 2007, Vol. 82, No. 6, pp. 748–755. Pleiades Publishing, Ltd., 2007.

[22] M. Lustig, D. Donoho, J. Santos and J. Pauly. Compressed Sensing MRI. IEEE Signal Processing Magazine, March (2008): 72–82.

[23] V. D. Milman and G. Schechtman. *Asymptotic Theory of Finite Dimensional Normed Spaces, With an Appendix by M. Gromov*. Lecture Notes in Mathematics 1200. Springer, (2001).

[24] D. Needell and J. A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. arXiv:0803.2392v2, 2008.

[25] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit. Submitted for publication, October 2007.

[26] M. Rudelson and R. Vershynin. Geometric approach to error correcting codes and reconstruction of signals. International Mathematical Research Notices, 64 (2005), pp. 4019–4041.

[27] F. Santosa and W. Symes. Linear inversion of band-limited reflection histograms. SIAM Journal on Scientific and Statistical Computing. 7 (1986), pp. 1307–1330.

[28] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. IEEE Trans. Info. Theory, 53(12): 4655–4666, 2007.

[29] Y. Zhang. A Simple Proof for Recoverability of ℓ_1-Minimization: Go Over or Under? Rice University CAAM Technical Report TR05-09, (2005).

[30] Y. Zhang. A Simple Proof for Recoverability of ℓ_1-Minimization (II): the Nonnegativity Case. Rice University CAAM Technical Report TR05-10, (2005).

[31] Compressive Sensing Resources. `http://www.dsp.ece.rice.edu/cs`.