

COMPRESSED SENSING: HOW SHARP IS THE RESTRICTED ISOMETRY PROPERTY?

JEFFREY D. BLANCHARD*, CORALIA CARTIS†, AND JARED TANNER‡

Abstract. Compressed Sensing (CS) seeks to recover an unknown vector with N entries by making far fewer than N measurements; it posits that the number of compressed sensing measurements should be comparable to the information content of the vector, not simply N . CS combines the important task of compression directly with the measurement task. Since its introduction in 2004 there have been hundreds of manuscripts on CS, a large fraction of which develop algorithms to recover a signal from its compressed measurements.

Because of the paradoxical nature of CS – exact reconstruction from seemingly undersampled measurements – it is crucial for acceptance of an algorithm that rigorous analyses verify the degree of undersampling the algorithm permits. The Restricted Isometry Property (RIP) has become the dominant tool used for the analysis in such cases.

We present here an asymmetric form of RIP which gives tighter bounds than the usual symmetric one. We give the best known bounds on the RIP constants for matrices from the Gaussian ensemble. Our derivations illustrate the way in which the combinatorial nature of CS is controlled. Our quantitative bounds on the RIP allow precise statements as to how aggressively a signal can be undersampled, the essential question for practitioners. We also document the extent to which RIP gives precise information about the true performance limits of CS, by comparing with approaches from high-dimensional geometry.

Key words. Compressed sensing, sparse approximation, restricted isometry property, phase transitions, convex relaxation, Gaussian matrices, singular values of random matrices.

AMS subject classifications. Primary: 41A46, 94A12, 94A20. Secondary: 15A52, 60F10, 90C25.

1. Introduction. Consider the task of measuring an unknown vector $x \in \mathbb{R}^N$ by taking inner products with vectors of one's choosing. The obvious choice would be to ask for the inner product of x with respect to each of the N canonical unit vectors e_j (the j^{th} entry of e_j being one and all others zero). But what if it is known a priori that x is k -sparse – i.e. has only $k < N$ nonzero entries? Can't one then do better? If the nonzero entries of x are indexed by the set K ($x(j) \neq 0$ if $j \in K$ and $x(j) = 0$ for $j \in K^c$), then only k inner products are needed: those with the canonical unit vectors e_j for $j \in K$. However, what if K is unknown? Is it still possible to make fewer than N measurements of x ?

Questions of this form must have been around for millennia. Consider this puzzle: “A counterfeit coin is hidden in a batch of N otherwise similar coins; it is distinguished from the others by its slightly heavier weight. How many balance weighings are needed to find the counterfeit?” Abstractly, this concerns the special case where K is an unknown singleton and the nonzero value is nonnegative; the balance is abstractly the same as an inner product which gives weight +1 to the coefficients placed in the “right” pan and -1 to the coefficients placed in the “left” pan. Many people quickly find that roughly $\log(N)$ measurements suffice to find the position and value of the nonzero, each time putting half the remaining coins in one pan, half in the other, and discarding from further consideration the coins that turn out on the light side. Lighthearted as puzzles can sometimes seem, they can lead to serious applications.

During World War Two, efficient screening of large groups of soldiers for certain infections was based on the principle of group testing, in which blood from many soldiers is combined in a single tube and tested for presence of an infectious agent. If an infection is found, one studies that group and by dyadic subdivision eventually isolates the infecteds [26, 34].

*Department of Mathematics and Statistics, Grinnell College, Grinnell, Iowa, USA (jeff@math.utah.edu). JDB acknowledges support from NSF DMS (VIGRE) grant number 0602219 while a VIGRE postdoctoral fellow at Department of Mathematics, University of Utah.

†School of Mathematics, University of Edinburgh, Edinburgh, UK (coralia.cartis@ed.ac.uk).

‡School of Mathematics, University of Edinburgh, Edinburgh, UK (jared.tanner@ed.ac.uk). JT acknowledges support from the Alfred P. Sloan Foundation, Leverhulme Trust, and thanks John E. and Marva M. Warnock for their generous support in the form of an endowed chair.

More advanced mathematics can do much better than such common-sense ideas. Those with a physical bent may quickly see that, if N is prime, again assuming a singleton K and a nonnegative x , it will be enough, in fact, to make only 2 inner products, with respectively a sine and cosine of frequency $2\pi/N$; the phase of the corresponding complex Fourier coefficient immediately reveals the position of the nonzero. Note here that, for large N , we are doing dramatically better than common-sense (2 measurements rather than $\log(N)$).

Advanced mathematics is better than the common-sense approach in another way: common-sense uses adaptive measurements, where the next measurement vector is selected after viewing all previous measurements. In the advanced approach, adaptivity is unnecessary: one simply makes 2 measurements defined a priori and later combines the two to reconstruct.

Compressed Sensing (CS) embodies the advanced approach: it designs a special matrix A of size $n \times N$, measures x via $y = Ax$, giving n measurements of the N vector x in parallel, and reconstructs x from (y, A) using computationally efficient and stable algorithms. The key point is that n can be taken much smaller than N , and much closer to k . For example, if x is known to be k -sparse and nonnegative $n = 2k + 1$ suffices [21] and if x is only known to be k -sparse, roughly $n = 2 \log(N/n) \cdot k$ will suffice, if k/N is small [22].

Since the release of the seminal CS papers in 2004, [10, 8, 17], a great deal of excitement has been generated in signal processing and applied mathematics research, with hundreds of papers on the theory, applications, and extensions of compressed sensing (more than 400 of these are collected at Rice's online Compressive Sensing Resources archive dsp.rice.edu/cs). Many applications have been proposed, including magnetic resonance imaging [40, 41], radar [45], and single-pixel cameras [28] to name a few. In the MRI applications, it has been reported that diagnostic quality images can be obtained in 1/7 the recording time using CS approaches, [39]. For a recent review of CS see the special issue containing [28, 40] and for a review of sparse approximation see [5].

In CS the matrix A and reconstruction algorithm are referred to as an encoder/decoder pair and much of the research has focused on their construction; that is, how should the measurement matrix A be selected and what are the most computationally efficient and robust algorithms for recovering x given y and A ? The two most prevalent encoders in the literature construct A by drawing its entries independently and identically from a Gaussian normal distribution, or by randomly sampling its rows without replacement from amongst the rows of a Fourier matrix. These encoders are popular as they are amenable to analysis, and they can be viewed as models of matrices with mean-zero entries and fast matrix-vector products, respectively. The most widely-studied decoder has been ℓ^1 -minimization,

$$(1.1) \quad \min_{z \in \mathbb{R}^N} \|z\|_1 \text{ subject to } Az = y,$$

which is the convex relaxation of the computationally intractable decoder, [42], seeking the sparsest solution in agreement with the measurements

$$(1.2) \quad \min_{z \in \mathbb{R}^N} \|z\|_0 \text{ subject to } Az = y.$$

Following the usual convention in the CS community, $\|z\|_0$ counts the number of nonzero entries in z . Many other encoder/decoder pairs are also being actively studied, with new alternatives being proposed regularly; see Section 3.

Here we do not review these exciting activities, but focus our attention on how to interpret the existing theoretical guarantees; in particular, we believe an important task for theory is to correctly predict the triples (k, n, N) for which a given encoder/decoder will successfully recover the measured signal, or a suitable approximation thereof. To exemplify this, we restrict our attention to a now-standard encoder/decoder pair: A Gaussian and ℓ^1 -minimization. This pair offers the cleanest mathematical structure, giving us the chance to make the strongest and clearest statements which can be made at this time, for example by drawing on the existing wealth of knowledge in random

matrix theory and high-dimensional convex geometry. In this paper we focus almost exclusively on the most widely used tool for analyzing the performance of encoder/decoder pairs, the *Restricted Isometry Property* (RIP) introduced by Candès and Tao [11].

DEFINITION 1.1 (Restricted Isometry Property). *A matrix A of size $n \times N$ is said to satisfy the RIP with RIP constant $R(k, n, N; A)$ if, for every $x \in \chi^N(k) := \{x \in \mathbb{R}^N : \|x\|_0 \leq k\}$,*

$$(1.3) \quad R(k, n, N; A) := \min_{c \geq 0} c \text{ subject to } (1 - c)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + c)\|x\|_2^2.$$

As suggested by the name, the RIP constants measure how much the matrix A acts like an isometry when “restricted” to k columns; it describes the most significant distortions of the ℓ^2 norm of any k -sparse vector. Typically, $R(k, n, N; A)$ is measured for matrices with unit ℓ^2 -norm columns, and in this special case $R(1, n, N) = 0$. Specifically, the RIP constant $R(k, n, N; A)$ is the maximum distance from 1 of all the eigenvalues of the $\binom{N}{k}$ submatrices, $A_K^T A_K$, derived from A , where K is an index set of cardinality k which restricts A to those columns indexed by K .

It is important to note that the RIP is predominantly used to establish theoretical performance guarantees when either the measurement vector y is corrupted with noise or the vector x is not strictly k -sparse. Proving that an algorithm is stable to noisy measurements is essential for applications since measurements are rarely free from noise. In this paper, we focus on the ideal noiseless case with the hopes of investigating the best possible theoretical results. For the noisy case, see [1] for ℓ^q -minimization for $q \in (0, 1]$ and [3] for greedy algorithms.

For many CS encoder/decoder pairs it has been shown that if the RIP constants for the encoder remain bounded as n and N increase with $n/N \rightarrow \delta \in (0, 1)$, then the decoder can be guaranteed to recover the sparsest x for k up to a critical threshold, which can be expressed as a fraction of n , $\rho(\delta) \cdot n$. Typically each encoder/decoder pair has a different $\rho(\delta)$. Little is generally known about the magnitude of $\rho(\delta)$ for encoder/decoder pairs, making it difficult for a practitioner to know how aggressively they may undersample, or which decoder has stronger performance guarantees. (For a recent review of compressed sensing algorithms, including which have $\rho(\delta) > 0$, see [43, Section 7].) In this paper, we endeavor to be as precise as possible about the value of the RIP constants for the Gaussian ensemble, and show how this gives quantitative values for $\rho(\delta)$ for the ℓ^1 -minimization decoder. Similar results for other decoders are available in [3].

To quantify the sparsity/undersampling trade off, we adopt a *proportional-growth* asymptotic, in which we consider sequences of triples (k, n, N) where all elements grow large in a coordinated way, $n \sim \delta N$ and $k \sim \rho n$ for some constants $\delta, \rho > 0$. This defines a two-dimensional phase space (δ, ρ) in $[0, 1]^2$ for asymptotic analysis.

DEFINITION 1.2 (Proportional-Growth Asymptotic). *A sequence of problem sizes (k, n, N) is said to grow proportionally if, for $(\delta, \rho) \in [0, 1]^2$, $\frac{n}{N} \rightarrow \delta$ and $\frac{k}{n} \rightarrow \rho$ as $n \rightarrow \infty$.*

Ultimately, we want to determine, as precisely as possible, which subset of this phase space corresponds to successful recovery and which subset corresponds to unsuccessful recovery. This is the phase-transition framework advocated by Donoho et al [16, 18, 20, 21, 24]; see Section 3 for a precise definition. By translating the sufficient RIP conditions into the proportional-growth asymptotic, we find lower bounds on the phase-transition for (δ, ρ) in $[0, 1]^2$. An answer to this question plays the role of an *undersampling theorem*: to what degree can we undersample a signal and still be able to reconstruct it?

The central aims of this paper are:

- to shed some light on the behavior of the RIP constants of a matrix ensemble with as much precision as possible;
- to advocate a unifying framework for the comparison of theoretical CS results by showing the reader how to interpret and compare some of the existing recovery guarantees for the prevalent ℓ^1 decoder;
- to introduce a reader new to this topic to the type of large deviation analysis calculations often encountered in CS and applicable to many areas faced with combinatorial challenges.

In pursuit of these goals, we sharpen the use of the RIP and squeeze the most out of it, quantifying what can currently be said in the proportional-growth asymptotic and thereby making precise the undersampling theorems the RIP implies. We proceed in Section 2 along two main avenues. First, we concentrate on Gaussian matrices, using bounds on their singular values we develop the sharpest known bounds on their RIP constants; in fact, these are the the best known bounds of any class of matrices in the proportional-growth asymptotic with $n < N$. Second, we use an asymmetric definition of the RIP where the lower and upper eigenvalues are treated separately, and in doing so further improve the conditions in which the RIP implies CS decoders recover the measured signal. In Section 3 we combine these two improvements to exhibit a region of the (δ, ρ) phase space where RIP analysis shows that undersampling will be successful for the ℓ^1 -minimization decoder (1.1).

The RIP is not the only tool used to analyze the performance of CS decoders. The different methods of analysis lead to results that are rather difficult to compare. In Section 3.2, we describe in the proportional-growth asymptotic, with A Gaussian and the ℓ^1 -minimization decoder, two alternative methods bounding the phase transition: the polytope analysis [16, 18, 22] of Donoho and Tanner and the geometric functional analysis techniques of Rudelson and Vershynin [46]. By translating these two methods of analysis and the RIP analysis into the proportional-growth asymptotic, we can readily compare the results obtained by these three techniques by comparing the regions of the (δ, ρ) phase space where each method of analysis has guaranteed successful recovery. In particular, we find that for the Gaussian encoder, the RIP, despite its popularity, is currently dramatically weaker than the other two approaches in the strength of conclusions that it can offer. However, this limitation is counterbalanced by RIP being successfully applied to a broad class of encoder/decoder pairs, and seamlessly also proving stability to noisy measurements and compressible signals.

We conclude with a discussion of some other important and related topics not addressed in the current paper. We briefly discuss comparisons of results when noise is present in the measurements or the signal x is not perfectly k -sparse, average case analysis versus the theoretical worst case analysis presented here, and the potential to improve the phase transition curves through improved analysis or improved bounds.

2. Bounds on RIP for Gaussian Random Matrices. Let $K \subset \{1, \dots, N\}$ be an index set of cardinality k which specifies the columns of A chosen for a submatrix, A_K , of size $n \times k$. Explicitly computing $R(k, n, N; A)$ would require enumerating all $\binom{N}{k}$ subsets K of the columns of A , forming each matrix $G_K = A_K^T A_K$, and calculating their largest and smallest eigenvalues. We have never seen this done except for small sizes of N and k , so not much is known about the RIP constants of deterministic matrices. Fortunately, analysis can penetrate where computation becomes intractable. Associated with a random matrix ensemble is an, as of yet unknown, probability density function for $R(k, n, N)$. Let us focus on the Gaussian ensemble where much is already known about its eigenvalues. We say that an $n \times N$ random matrix A is drawn from the *Gaussian ensemble* of random matrices if the entries are sampled independently and identically from the standard normal distribution, $\mathcal{N}(0, n^{-1})$. (The n^{-1} scaling in the Gaussian ensemble cause the ℓ^2 norm of its columns to have expectation 1.) We say that a $k \times k$ matrix $W_{n,k}$ is a *Wishart* matrix if it is the Gram matrix $X^T X$ of an $n \times k$ matrix X from the Gaussian ensemble. The largest and smallest eigenvalues of a Wishart matrix are random variables, denoted here $\Lambda_{n,k}^{max} = \lambda^{max}(W_{n,k})$ and $\Lambda_{n,k}^{min} = \lambda^{min}(W_{n,k})$. These random variables tend to defined limits, in expectation, as n and k increase in a proportional manner. With $\frac{k}{n} \rightarrow \rho$ as $n \rightarrow \infty$, we have $\mathcal{E}(\Lambda_{n,k}^{max}) \rightarrow (1 + \sqrt{\rho})^2$ and $\mathcal{E}(\Lambda_{n,k}^{min}) \rightarrow (1 - \sqrt{\rho})^2$; [33, 48], see Figure 2.1. Explicit formulas bounding $\Lambda_{n,k}^{max}$ and $\Lambda_{n,k}^{min}$ are available [30]. An empirical approximation of the probability density functions of $\Lambda_{n,k}^{max}$ and $\Lambda_{n,k}^{min}$ is shown in Figure 2.2.

The asymmetric way that the expected eigenvalues $\Lambda_{n,k}^{max}$ and $\Lambda_{n,k}^{min}$ deviate from 1 suggests that the symmetric treatment used by the traditional RIP is missing an important part of the picture. We generalize the RIP to an asymmetric form and derive the sharpest recovery conditions implied by the RIP.

DEFINITION 2.1 (Asymmetric Restricted Isometry Property). *For a matrix A of size $n \times N$,*

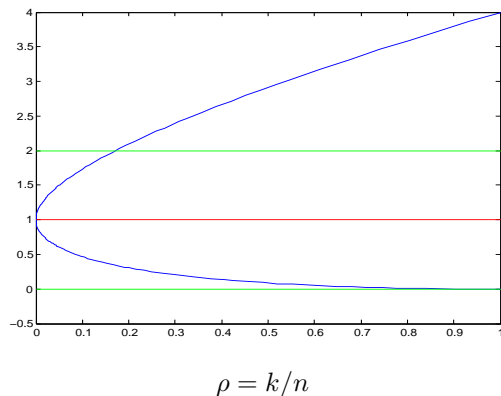


FIG. 2.1. Expected values of the largest and smallest eigenvalues of a Wishart matrix $W_{n,k}$ with $\rho = \frac{k}{n}$. Note the asymmetry with respect to 1.

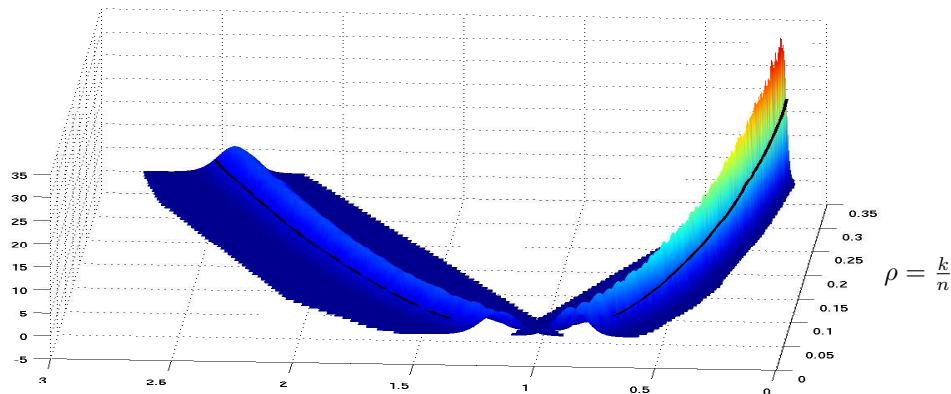


FIG. 2.2. **Empirical Distributions of the Largest and Smallest Eigenvalues of a Wishart Matrix.** A collection of frequency histograms of $\Lambda_{n,k}^{max}$ and $\Lambda_{n,k}^{min}$: x-axis – size of the eigenvalue; y-axis – number of occurrences; z-axis – ratio $\rho = \frac{k}{n}$ of the Wishart parameters. Overlays: curves depicting the expected values $(1 \pm \sqrt{\rho})^2$ of $\Lambda_{n,k}^{max}$ and $\Lambda_{n,k}^{min}$. Here $n = 200$. At this value of n it is evident that $\Lambda_{n,k}^{max}$ and $\Lambda_{n,k}^{min}$ lie near, but not on curves. For larger n , the concentration would be tighter.

the asymmetric RIP constants $L(k, n, N; A)$ and $U(k, n, N; A)$ are defined as:

$$(2.1) \quad L(k, n, N; A) := \min_{c \geq 0} c \text{ subject to } (1 - c)\|x\|_2^2 \leq \|Ax\|_2^2, \text{ for all } x \in \chi^N(k);$$

$$(2.2) \quad U(k, n, N; A) := \min_{c \geq 0} c \text{ subject to } (1 + c)\|x\|_2^2 \geq \|Ax\|_2^2, \text{ for all } x \in \chi^N(k).$$

(A similar change in the definition of the RIP constants was used independently by Foucart and Lai in [32], motivated by different concerns.)

REMARK 1. Although both the smallest and largest singular values of $A_K^T A_K$ affect the stability of the reconstruction algorithms, the smaller eigenvalue is dominant for compressed sensing in that it allows distinguishing between sparse vectors from their measurement by A . In fact, it is often incorrectly stated that $R(2k, n, N) < 1$ is a necessary condition to ensure that there are no two k -sparse vectors, say x and x' , with the same measurements $Ax = Ax'$; the actual necessary condition is $L(2k, n, N) < 1$.

We see from (2.1) and (2.2) that $(1 - L(k, n, N)) = \min_K \lambda^{min}(G_K)$ and $(1 + U(k, n, N)) =$

$\max_K \lambda^{\max}(G_K)$ with $G_K = A_K^T A_K$. A standard large deviation analysis of bounds on the probability density functions of $\Lambda_{n,k}^{\max}$ and $\Lambda_{n,k}^{\min}$ allows us to establish upper bounds of $L(k, n, N)$ and $U(k, n, N)$ which are exponentially unlikely to be exceeded.

DEFINITION 2.2 (Asymptotic RIP Bounds). *Let A be a matrix of size $n \times N$ drawn from the Gaussian ensemble and consider the proportional-growth asymptotic ($\frac{n}{N} \rightarrow \delta$ and $\frac{k}{n} \rightarrow \rho$ as $n \rightarrow \infty$). Let $H(p) := p \log(1/p) + (1-p) \log(1/(1-p))$ denote the usual Shannon Entropy with base e logarithms, and let*

$$(2.3) \quad \psi_{\min}(\lambda, \rho) := H(\rho) + \frac{1}{2} [(1-\rho) \log \lambda + 1 - \rho + \rho \log \rho - \lambda],$$

$$(2.4) \quad \psi_{\max}(\lambda, \rho) := \frac{1}{2} [(1+\rho) \log \lambda + 1 + \rho - \rho \log \rho - \lambda].$$

Define $\lambda^{\min}(\delta, \rho)$ and $\lambda^{\max}(\delta, \rho)$ as the solution to (2.5) and (2.6), respectively:

$$(2.5) \quad \delta \psi_{\min}(\lambda^{\min}(\delta, \rho), \rho) + H(\rho \delta) = 0 \quad \text{for} \quad \lambda^{\min}(\delta, \rho) \leq 1 - \rho$$

$$(2.6) \quad \delta \psi_{\max}(\lambda^{\max}(\delta, \rho), \rho) + H(\rho \delta) = 0 \quad \text{for} \quad \lambda^{\max}(\delta, \rho) \geq 1 + \rho.$$

Define $\mathcal{L}(\delta, \rho)$ and $\mathcal{U}(\delta, \rho)$ as

$$(2.7) \quad \mathcal{L}(\delta, \rho) := 1 - \lambda^{\min}(\delta, \rho) \quad \text{and} \quad \mathcal{U}(\delta, \rho) := \min_{\nu \in [\rho, 1]} \lambda^{\max}(\delta, \nu) - 1.$$

To facilitate ease of calculating $\mathcal{L}(\delta, \rho)$ and $\mathcal{U}(\delta, \rho)$, web forms for their calculation are available at ecos.maths.ed.ac.uk.

In the proportional growth asymptotic, the probability that $\mathcal{L}(\delta, \rho)$ and $\mathcal{U}(\delta, \rho)$ bound the random variables $L(k, n, N)$ and $U(k, n, N)$, respectively, tends to 1 as $n \rightarrow \infty$. In statistical terminology, the coverage probability of the upper confidence bounds $\mathcal{L}(\delta, \rho)$ and $\mathcal{U}(\delta, \rho)$ tends to one as $n \rightarrow \infty$. In fact, all probabilities presented in this manuscript converge to their limit “exponentially in n ”; that is, the probability for finite n approaches its limit as n grows with discrepancy bounded by a constant multiple of $e^{-n\beta}$ for some fixed $\beta > 0$.

THEOREM 2.3 (Validity of RIP Bounds). *Fix $\epsilon > 0$. Under the proportional-growth asymptotic, Definition 1.2, sample each $n \times N$ matrix A from the Gaussian ensemble. Then*

$$\text{Prob}(L(k, n, N; A) < \mathcal{L}(\delta, \rho) + \epsilon) \rightarrow 1 \quad \text{and} \quad \text{Prob}(U(k, n, N; A) < \mathcal{U}(\delta, \rho) + \epsilon) \rightarrow 1$$

exponentially in n .

REMARK 2. *Extensive empirical estimates of $L(k, n, N)$ and $U(k, n, N)$ show that the bounds $\mathcal{L}(\delta, \rho)$ and $\mathcal{U}(\delta, \rho)$ are rather sharp; in fact, they are no more than twice the actual upper bounds on $L(k, n, N)$ and $U(k, n, N)$, see Figure 2.4 and Table 2.1, and are much closer for the region applicable for CS decoders, $\rho \ll 1$. The empirically observed lower bounds on $L(k, n, N)$ and $U(k, n, N)$ are calculated through the following process. The number of rows, n , is fixed at one of the values in Table 2.1. For each n , 47 values of N are selected so that n/N ranges from 1/20 to 20/21. For each (n, N) a matrix A of size $n \times N$ is drawn from $\mathcal{N}(0, n^{-1})$ and either the algorithm from [27] or [38] is applied to determine support sets of size $k = 1, 2, \dots, n-1$ which are candidates for the support sets that maximize $L(k, n, N; A)$ or $U(k, n, N; A)$. The largest or smallest eigenvalue of each resulting $n \times k$ submatrix is calculated and recorded. The above process is repeated for some number of matrices, see the caption of Table 2.1, and the maximum value recorded. The empirical calculation of RIP constants are lower bounds on the true RIP constants as the support sets calculated by [27] and [38] may not be the support sets which maximize the RIP constants.*

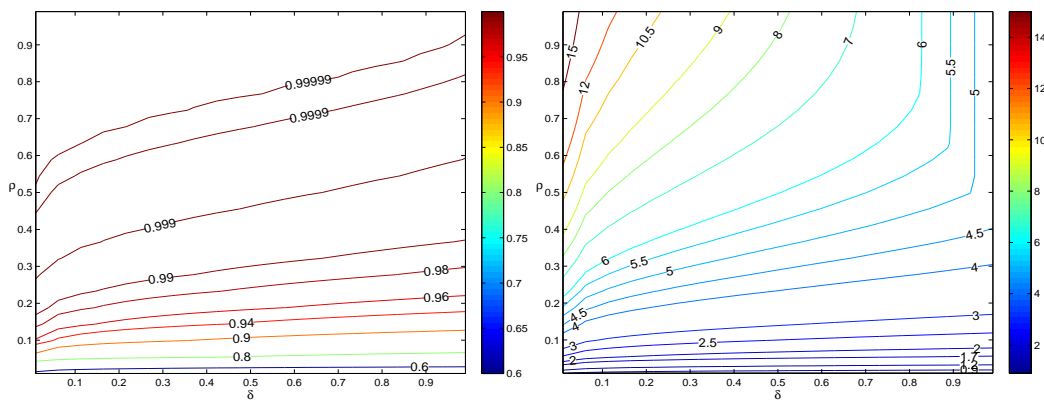


FIG. 2.3. **The RIP bounds of Eq. (2.7).** Level sets of $\mathcal{L}(\delta, \rho)$ (left panel) and $\mathcal{U}(\delta, \rho)$ (right panel) over the phase space $(\delta, \rho) \in [0, 1]^2$. For large matrices from the Gaussian ensemble, it is overwhelmingly unlikely that the RIP constants $L(k, n, N; A)$ and $U(k, n, N; A)$ will be greater than these values.

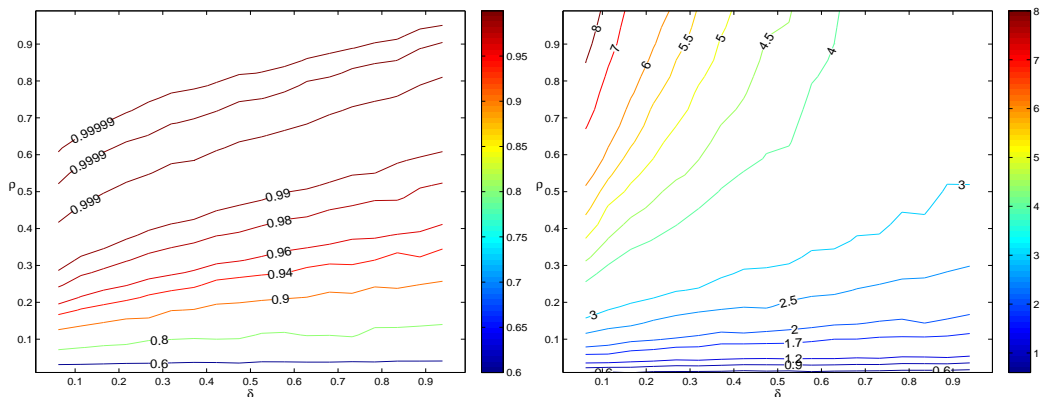


FIG. 2.4. **Empirically observed lower estimates of RIP bounds of RIP constants.** Although there is no computationally tractable method for calculating the RIP constants of a matrix, there are efficient algorithms which perform local searches for extremal eigenvalues of submatrices; allowing for observable lower bounds on the RIP constants. Algorithms for lower bounding $L(k, n, N)$, [27], and $U(k, n, N)$, [38], were applied to dozens of A drawn Gaussian $\mathcal{N}(0, n^{-1})$ with $n = 400$ and N increasing from 420 to 8000. Level sets of the observed $L(k, n, N; A)$ (left panel) and $U(k, n, N; A)$ (right panel).

n	$\max \frac{\mathcal{L}(\delta, \rho)}{L(k, n, N)}$	$\max \frac{\mathcal{U}(\delta, \rho)}{U(k, n, N)}$
200	1.22	1.83
400	1.32	1.81

TABLE 2.1

The maximum ratio of the RIP bounds in Theorem 2.3 to empirically observed values. For each of the ratios n/N tested, multiple matrices were drawn and empirical low bounds on their RIP constants calculated. For $n = 200$ between 9 and 175 matrices were drawn for each n/N , and for $n = 400$ between 7 and 489 matrices were drawn for each n/N . Our bounds are numerically found to be within a multiple of 1.83 of empirically observed lower bounds.

2.1. Proof of Theorem 2.3. In order to prove Theorem 2.3, this section employs a type of large deviation technique often encountered in CS and applicable in fact, to many areas faced with combinatorial challenges.

We first establish some useful lemmas concerning the extreme eigenvalues of Wishart matrices. The matrix A generates $\binom{N}{k}$ different Wishart matrices $G_k = A_K^T A_K$. Exponential bounds on the tail probabilities of the largest and smallest eigenvalues of such Wishart matrices can be combined

with exponential bounds on $\binom{N}{k}$ to control the chance of large deviations using the union bound. This large deviation analysis technique is characteristic of proofs in compressed sensing. By using the exact probability density functions on the tail behavior of the extreme eigenvalues of Wishart matrices the overestimation of the union bound is dramatically reduced. We focus on the slightly more technical results for the bound on the most extreme of the largest eigenvalues, $\mathcal{U}(\delta, \rho)$, and prove these statements in full detail. Corresponding results for $\mathcal{L}(\delta, \rho)$ are stated with their similar proofs omitted.

The probability density function, $f_{max}(k, n; \lambda)$, for the largest eigenvalue of the $k \times k$ Wishart matrix $A_K^T A_K$ was determined by Edelman in [29]. For our analysis, a simplified upper bound suffices.

LEMMA 2.4 (Lemma 4.2, pp. 550 [29]). *Let A_K be a matrix of size $n \times k$ whose entries are drawn i.i.d from $\mathcal{N}(0, n^{-1})$. Let $f_{max}(k, n; \lambda)$ denote the probability density function for the largest eigenvalue of the Wishart matrix $A_K^T A_K$ of size $k \times k$. Then $f_{max}(k, n; \lambda)$ satisfies:*

$$(2.8) \quad f_{max}(k, n; \lambda) \leq \left[(2\pi)^{1/2} (n\lambda)^{-3/2} \left(\frac{n\lambda}{2} \right)^{(n+k)/2} \frac{1}{\Gamma(\frac{k}{2})\Gamma(\frac{n}{2})} \right] \cdot e^{-n\lambda/2} =: g_{max}(k, n; \lambda).$$

For our purposes, it is sufficient to have a precise characterization of $g_{max}(k, n; \lambda)$'s exponential (with respect to n) behavior.

LEMMA 2.5. *Let $k/n = \rho \in (0, 1)$ and define*

$$\psi_{max}(\lambda, \rho) := \frac{1}{2} [(1 + \rho) \log \lambda + 1 + \rho - \rho \log \rho - \lambda].$$

Then

$$(2.9) \quad f_{max}(k, n; \lambda) \leq p_{max}(n, \lambda) \exp(n \cdot \psi_{max}(\lambda, \rho))$$

where $p_{max}(n, \lambda)$ is a polynomial in n, λ .

Proof. Let $g_{max}(k, n; \lambda)$ be as defined in (2.8) and let $\rho_n = k/n$. To extract the exponential behavior of $g_{max}(k, n; \lambda)$ we write $\frac{1}{n} \log(g_{max}(k, n; \lambda)) = \Phi_1(k, n; \lambda) + \Phi_2(k, n; \lambda) + \Phi_3(k, n; \lambda)$ where

$$\begin{aligned} \Phi_1(k, n; \lambda) &= \frac{1}{2n} \log(2\pi) - \frac{3}{2n} \log(n\lambda) \\ \Phi_2(k, n; \lambda) &= \frac{1}{2} \left[(1 + \rho_n) \log \left(\frac{\lambda n}{2} \right) - \lambda \right] \\ \Phi_3(k, n; \lambda) &= -\frac{1}{n} \log \left(\Gamma \left(\frac{k}{2} \right) \Gamma \left(\frac{n}{2} \right) \right). \end{aligned}$$

Clearly, $\lim_{n \rightarrow \infty} \Phi_1(k, n; \lambda) = 0$ and can be subsumed as part of $p_{max}(n, \lambda)$. To simplify Φ_3 , we apply the second of Binet's log gamma formulas [52, Sec. 12.32], namely $\log(\Gamma(z)) = (z - 1/2) \log z - z + \log \sqrt{2\pi} + I$ where I is a convergent, improper integral. With $c(n, \rho)$ representing the constant and integral from Binet's formula we then have

$$\Phi_2(k, n; \lambda) + \Phi_3(k, n; \lambda) = \frac{1}{2} \left[(1 + \rho_n) \log \lambda - \left(\rho_n - \frac{1}{n} \right) \log \rho_n + \frac{2}{n} \log \frac{n}{2} + \rho_n + 1 - \lambda + \frac{1}{n} c(n, \rho_n) \right].$$

As $\lim_{n \rightarrow \infty} n^{-1} c(n, \rho_n) = 0$ it can be absorbed into $p_{max}(n, \lambda)$ and we have

$$\psi_{max}(\lambda, \rho) := \lim_{n \rightarrow \infty} \frac{1}{n} \log [g_{max}(k, n; \lambda)] = \frac{1}{2} [(1 + \rho) \log \lambda - \rho \log \rho + \rho + 1 - \lambda]$$

and the conclusion follows. \square

To bound $U(k, n, N)$, we must simultaneously account for all $\binom{N}{k}$ Wishart matrices $A_K^T A_K$ derived from A . Using a union bound this amounts to studying the exponential behavior of $\binom{N}{k} g_{max}(k, n; \lambda)$. In the proportional-growth asymptotic this can be determined by characterizing $\lim_{N \rightarrow \infty} N^{-1} \log \left[\binom{N}{k} g_{max}(k, n; \lambda) \right]$, which from Lemma 2.5 is given by

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \log \left[\binom{N}{k} g_{max}(k, n; \lambda) \right] &= \lim_{N \rightarrow \infty} \frac{1}{N} \log \left[\binom{N}{k} \right] + \lim_{N \rightarrow \infty} \frac{1}{N} \log [g_{max}(n, k; \lambda)] \\ &= H \left(\frac{k}{N} \right) + \delta \lim_{n \rightarrow \infty} \frac{1}{n} \log [g_{max}(n, k; \lambda)] \\ (2.10) \qquad \qquad \qquad &= H(\rho\delta) + \delta \psi_{max}(\lambda, \rho) =: \delta \psi_U(\delta, \rho; \lambda). \end{aligned}$$

Recall that $H(p) := p \log(1/p) + (1-p) \log(1/(1-p))$ is the usual Shannon Entropy with base e logarithms.

Equipped with Lemma 2.5 and (2.10), Proposition 2.6 establishes $\lambda^{max}(\delta, \rho) - 1$ as an upper bound on $U(k, n, N)$ in the proportional-growth asymptotic.

PROPOSITION 2.6. *Let $\delta, \rho \in (0, 1)$, and A be a matrix of size $n \times N$ whose entries are drawn i.i.d. from $\mathcal{N}(0, n^{-1})$. Define $\tilde{U}(\delta, \rho) := \lambda^{max}(\delta, \rho) - 1$ where $\lambda^{max}(\delta, \rho)$ is the solution to (2.6). Then for any $\epsilon > 0$, in the proportional-growth asymptotic*

$$Prob \left(U(k, n, N) > \tilde{U}(\delta, \rho) + \epsilon \right) \rightarrow 0$$

exponentially in n .

Proof. Throughout this proof δ and ρ are fixed, and we focus our attention on λ , often abbreviating $\psi_U(\delta, \rho; \lambda)$ in (2.10) as $\psi_U(\lambda)$. We first verify that (2.6) has a unique solution. Since

$$\frac{d}{d\lambda} \psi_U(\lambda) = \frac{1}{2} \left(\frac{1+\rho}{\lambda} - 1 \right),$$

$\psi_U(\lambda)$ is strictly decreasing on $[1+\rho, \infty)$ and is strictly concave. Combined with

$$\psi_U(1+\rho) = \delta^{-1} H(\rho\delta) + \frac{1}{2} \left[(1+\rho) \log(1+\rho) + \rho \log \frac{1}{\rho} \right] > 0$$

and $\lim_{\lambda \rightarrow \infty} \psi_U(\lambda) = -\infty$, there is a unique solution to (2.6), namely $\lambda^{max}(\delta, \rho)$.

Select $\epsilon > 0$ and let (k, n, N) be such that $\frac{n}{N} = \delta_n$, $\frac{k}{n} = \rho_n$. First, we write the probability statement in terms of $\lambda^{max}(\delta_n, \rho_n)$:

$$\begin{aligned} Prob \left[U(k, n, N) > \tilde{U}(\delta_n, \rho_n) + \epsilon \right] &= Prob [U(k, n, N) > \lambda^{max}(\delta_n, \rho_n) - 1 + \epsilon] \\ &= Prob [1 + U(k, n, N) > \lambda^{max}(\delta_n, \rho_n) + \epsilon] \\ &= \binom{N}{k} \int_{\lambda^{max}(\delta_n, \rho_n) + \epsilon}^{\infty} f_{max}(k, n; \lambda) d\lambda \\ (2.11) \qquad \qquad \qquad &\leq \binom{N}{k} \int_{\lambda^{max}(\delta_n, \rho_n) + \epsilon}^{\infty} g_{max}(k, n; \lambda) d\lambda. \end{aligned}$$

To bound the integral in (2.11) in terms of $g_{max}(\delta, \rho; \lambda^{max}(\delta_n, \rho_n))$ we write $g_{max}(k, n; \lambda)$ in terms of n, ρ_n , and λ as $g_{max}(k, n; \lambda) = \varphi(n, \rho_n) \lambda^{-\frac{3}{2}} \lambda^{\frac{n}{2}(1+\rho_n)} e^{-\frac{n}{2}\lambda}$ where

$$\varphi(n, \rho_n) = (2\pi)^{\frac{1}{2}} n^{-\frac{3}{2}} \left(\frac{n}{2} \right)^{\frac{n}{2}(1+\rho_n)} \frac{1}{\Gamma\left(\frac{n}{2}\rho_n\right) \Gamma\left(\frac{n}{2}\right)}.$$

Since $\lambda^{\max}(\delta_n, \rho_n) > 1 + \rho_n$, the quantity $\lambda^{\frac{n}{2}(1+\rho_n)} e^{-\frac{n}{2}\lambda}$ is strictly decreasing in λ on $[\lambda^{\max}(\delta, \rho_n), \infty)$. Therefore we have

$$\begin{aligned}
\int_{\lambda^{\max}(\delta_n, \rho_n) + \epsilon}^{\infty} g_{\max}(k, n; \lambda) d\lambda &\leq \varphi(n, \rho_n) (\lambda^{\max}(\delta_n, \rho_n) + \epsilon)^{\frac{n}{2}(1+\rho_n)} e^{-\frac{n}{2}(\lambda^{\max}(\delta_n, \rho_n) + \epsilon)} \int_{\lambda^{\max}(\delta_n, \rho_n) + \epsilon}^{\infty} \lambda^{-\frac{3}{2}} d\lambda \\
&= (\lambda^{\max}(\delta_n, \rho_n) + \epsilon)^{\frac{3}{2}} g_{\max}(k, n; \lambda^{\max}(\delta_n, \rho_n) + \epsilon) \int_{\lambda^{\max}(\delta_n, \rho_n) + \epsilon}^{\infty} \lambda^{-\frac{3}{2}} d\lambda \\
(2.12) \qquad \qquad \qquad &= 2 (\lambda^{\max}(\delta_n, \rho_n) + \epsilon) g_{\max}(k, n; \lambda^{\max}(\delta_n, \rho_n) + \epsilon).
\end{aligned}$$

Therefore, combining (2.11) and (2.12) we obtain

$$\begin{aligned}
\text{Prob} \left[U(k, n, N) > \tilde{\mathcal{U}}(\delta_n, \rho_n) + \epsilon \right] &\leq 2 (\lambda^{\max}(\delta_n, \rho_n) + \epsilon) \binom{N}{k} g_{\max}(k, n; \lambda^{\max}(\delta_n, \rho_n) + \epsilon) \\
&\leq p_{\max}(n, \lambda^{\max}(\delta_n, \rho_n)) \exp \left[n \cdot \psi_U(\lambda^{\max}(\delta_n, \rho_n) + \epsilon) \right] \\
(2.13) \qquad \qquad \qquad &\leq p_{\max}(n, \lambda^{\max}(\delta_n, \rho_n)) \exp \left[n\epsilon \cdot \frac{d}{d\lambda} \psi_U(\lambda) \Big|_{\lambda = (\lambda^{\max}(\delta_n, \rho_n))} \right],
\end{aligned}$$

with the last inequality following from the strict concavity of $\psi_U(\lambda)$. Since $\frac{d}{d\lambda} \psi_U(\lambda^{\max}(\delta, \rho)) < 0$ is strictly bounded away from zero and $\lim_{n \rightarrow \infty} \lambda^{\max}(\delta_n, \rho_n) = \lambda^{\max}(\delta, \rho)$, we arrive at, for any $\epsilon > 0$

$$\lim_{n \rightarrow \infty} \text{Prob} \left[U(k, n, N) > \tilde{\mathcal{U}}(\delta, \rho) + \epsilon \right] \rightarrow 0.$$

□

The term $H(\rho\delta)$ in (2.10), from the union bound over all $\binom{N}{k}$ matrices $A_K^T A_K$, results in an overly pessimistic bound in the vicinity of $\rho\delta = 1/2$. As we are seeking the least upper bound on $U(k, n, N)$ we note that any upper bound for $U(j, n, N)$ for $j > k$ is also an upper bound for $U(k, n, N)$, and replace the bound $\tilde{\mathcal{U}}(\delta, \rho)$ with the minimum of $\tilde{\mathcal{U}}(\delta, \nu)$ for $\nu \in [\rho, 1]$.

PROPOSITION 2.7. *Let $\delta, \rho \in (0, 1)$, and define $\mathcal{U}(\delta, \rho) := \min_{\nu \in [\rho, 1]} \tilde{\mathcal{U}}(\delta, \nu)$ with $\tilde{\mathcal{U}}(\delta, \nu)$ defined as in Proposition 2.6. For any $\epsilon > 0$, in the proportional-growth asymptotic*

$$\text{Prob}(U(k, n, N) > \mathcal{U}(\delta, \rho) + \epsilon) \rightarrow 0$$

exponentially in n .

Proof. By the definition of $\chi^N(k)$ in Definition 1.1, $U(j, n, N) \geq U(k, n, N)$ for $j = k + 1, k + 2, \dots, n$; combined with Proposition 2.6 for $\frac{j}{n} \rightarrow \nu$ as $n \rightarrow \infty$

$$\text{Prob} \left(U(j, n, N) > \tilde{\mathcal{U}}(\delta, \nu) + \epsilon \right) \rightarrow 0$$

exponentially in n , and taking a minimum over the compact set $\nu \in [\rho, 1]$ we arrive at the desired result. □

A similar approach leads to corresponding results for $\mathcal{L}(\delta, \rho)$. Edelman also determined the probability density function, $f_{\min}(k, n; \lambda)$, for the smallest eigenvalue of the $k \times k$ Wishart matrix $A_K^T A_K$ [29]. Here again, a simplified upper bound suffices:

LEMMA 2.8 (Prop. 5.2, pp. 553 [29]). *Let A_K be a matrix of size $n \times k$ whose entries are drawn i.i.d. from $\mathcal{N}(0, n^{-1})$. Let $f_{\min}(k, n; \lambda)$ denote the probability density function for the smallest eigenvalue of the Wishart matrix $A_K^T A_K$ of size $k \times k$. Then $f_{\min}(k, n; \lambda)$ satisfies:*

(2.14)

$$f_{\min}(k, n; \lambda) \leq \left(\frac{\pi}{2n\lambda} \right)^{1/2} \cdot e^{-n\lambda/2} \left(\frac{n\lambda}{2} \right)^{(n-k)/2} \cdot \left[\frac{\Gamma(\frac{n+1}{2})}{\Gamma(\frac{k}{2})\Gamma(\frac{n-k+1}{2})\Gamma(\frac{n-k+2}{2})} \right] =: g_{\min}(k, n; \lambda).$$

With Lemma 2.8, we establish a bound on the asymptotic behavior of the distribution of the smallest eigenvalue of Wishart matrix of size $k \times k$.

LEMMA 2.9. *Let $k/n = \rho \in (0, 1)$ and define*

$$\psi_{min}(\lambda, \rho) := H(\rho) + \frac{1}{2} [(1 - \rho) \log \lambda + 1 - \rho + \rho \log \rho - \lambda].$$

Then

$$(2.15) \quad f_{min}(k, n; \lambda) \leq p_{min}(n, \lambda) \exp(n \cdot \psi_{min}(\lambda, \rho))$$

where $p_{min}(n, \lambda)$ is a polynomial in n, λ .

With Lemma 2.9, the large deviation analysis yields

$$(2.16) \quad \lim_{N \rightarrow \infty} \frac{1}{N} \log \left[\binom{N}{k} g_{min}(k, n; \lambda) \right] = H(\rho\delta) + \delta\psi_{min}(\lambda, \rho).$$

Similar to the proof of Proposition 2.6, Lemma 2.9 and (2.16) are used to establish $\mathcal{L}(\delta, \rho)$ as an upper bound on $L(k, n, N)$ in the proportional-growth asymptotic.

PROPOSITION 2.10. *Let $\delta, \rho \in (0, 1]$, and A be a matrix of size $n \times N$ whose entries are drawn i.i.d. from $\mathcal{N}(0, n^{-1})$. Define $\mathcal{L}(\delta, \rho) := 1 - \lambda^{min}(\delta, \rho)$ where $\lambda^{min}(\delta, \rho)$ is the solution to (2.5). Then for any $\epsilon > 0$, in the proportional-growth asymptotic*

$$Prob(L(k, n, N) > \mathcal{L}(\delta, \rho) + \epsilon) \rightarrow 0$$

exponentially in n .

The bound $\mathcal{L}(\delta, \rho)$ is strictly increasing in ρ for any $\delta \in (0, 1)$, and as a consequence no tighter bound can be achieved by minimizing over matrices of larger size as was done in Proposition 2.7.

3. RIP Undersampling Theorems. The high level of interest in compressed sensing is due to the introduction of computationally efficient and stable algorithms which provably solve the seemingly intractable (1.2) even for k proportional to n . New compressed sensing decoders are being introduced regularly; broadly speaking, they fall into one of two categories: greedy algorithms and regularizations. Greedy algorithms are iterative, with each step selecting a locally optimal subset of entries in x which are adjusted to improve the desired error metric. Examples of greedy algorithms include Orthogonal Matching Pursuit (OMP) [50], Regularized OMP (ROMP) [44], Stagewise OMP (StOMP) [25], Compressive Sampling MP (CoSaMP) [43], Subspace Pursuit (SP) [15], and Iterated Hard Thresholding (IHT) [4]. Regularization formulations for sparse approximation began with the relaxation of (1.2) to the now ubiquitous (convex) ℓ^1 -minimization [14], (1.1), and has since been extended to non-convex ℓ^q -minimization for $q \in (0, 1)$, [35, 32, 13, 12, 47]. Although general-purpose convex optimization solvers may be employed to solve ℓ^1 -minimization (1.1), highly-efficient software has been recently designed specifically for ℓ^1 -minimization in the context of compressed sensing, see [14, 31, 51, 54]. Non-convex formulations have sometimes been able to offer substantial improvements, but at the cost of limited guarantees that the global minima can be found efficiently, so it remains unclear how practical they really are.

As stated at the end of the introduction, one of the central aims of this article is to advocate a unifying framework for the comparison of results in compressed sensing. Currently there is no general agreement in the compressed sensing community on such a framework, making it difficult to compare results obtained by different methods of analysis or to identify when new results are improvements over existing ones. Donoho has put forth the phase transition framework borrowed from the statistical mechanics literature and used successfully in a similar context by the combinatorial optimization community, see [36, 37]. This framework has been successfully employed in compressed sensing by Donoho et al, [20, 21, 24].

Fortunately, every compressed sensing algorithm that has an optimal recovery order of n proportional to k can be cast in the phase transition framework of Donoho et al., parametrized by two inherent problem size parameters¹:

- the *undersampling rate* of measuring x through n inner products with the rows of A , as compared to directly sampling each element of $x \in \mathbb{R}^N$:

$$\delta_n = n/N \in (0, 1)$$

- the *oversampling rate* of making n measurements as opposed to the optimal *oracle* rate of making k measurements when the oracle knows the support of x :

$$\rho_n = k/n \in (0, 1).$$

For each value of $\delta_n \in (0, 1)$ there is a largest value of ρ_n which guarantees successful recovery of x .

We now formalize the phase transition framework described above.

DEFINITION 3.1 (Strong Equivalence). *The event $StrongEquiv(A, \text{ALG})$ denotes the following property of an $n \times N$ matrix A : for every k -sparse vector x , the algorithm “ALG” exactly recovers x from the corresponding measurements $y = Ax$.*

For most compressed sensing algorithms and for a broad class of matrices, under the proportional-growth asymptotic there is a strictly positive function $\rho_S(\delta; \text{ALG}) > 0$ defining a region of the (δ, ρ) phase space which ensures successful recovery of every k -sparse vector $x \in \chi^N(k)$. This function, $\rho_S(\delta; \text{ALG})$, is called the *Strong phase transition function* [10, 16, 18].

DEFINITION 3.2 (Region of Strong Equivalence). *Consider the proportional-growth asymptotic with parameters $(\delta, \rho) \in (0, 1) \times (0, 1/2)$. Draw the corresponding $n \times N$ matrices A from the Gaussian ensemble and fix $\epsilon > 0$. Suppose that we are given a function $\rho_S(\delta; \text{ALG})$ with the property that, whenever $0 < \rho < (1 - \epsilon)\rho_S(\delta; \text{ALG})$, $\text{Prob}(StrongEquiv(A, \text{ALG})) \rightarrow 1$ as $n \rightarrow \infty$. We say that $\rho_S(\delta; \text{ALG})$ bounds a region of strong equivalence.*

REMARK 3. *The subscript S emphasizes that the phase transition function $\rho_S(\delta; \text{ALG})$ will define a region of the (δ, ρ) phase space which guarantees that the event $StrongEquiv(A, \text{ALG})$ is satisfied with probability on the draw of A converging to one exponentially in n . This notation has been established in the literature by Donoho and Tanner [16, 21] to distinguish strong equivalence (i.e. that every k -sparse vector x is successfully recovered) from weak equivalence (i.e. all but a small fraction of k -sparse vectors are successfully recovered). For example, [16, 21] study the event where ℓ^1 -minimization (1.1) exactly recovers x from the corresponding measurements $y = Ax$, except for a fraction $(1 - \epsilon)$ of the support sets.*

For the remainder of this section, we translate existing guarantees of $StrongEquiv(A, \ell^1)$ into bounds on the region of strong equivalence in the proportional-growth asymptotic; we denote $\rho_S(\delta; \ell^1) \equiv \rho_S(\delta)$. A similar presentation of other CS decoders is available in [3]. In order to make quantitative statements, the matrix or random matrix ensemble must first be specified, [2]; we again consider A drawn from the Gaussian ensemble.² In Section 3.1 we demonstrate how to incorporate the RIP bounds from Section 2 into results obtained from an RIP analysis. In Section 3.2 we compare bounds on the region of $StrongEquiv(A, \ell^1)$ proven by three distinct methods of analysis: eigenvalue analysis and the RIP [32], geometric functional analysis [46], and convex polytopes [16].

3.1. Region of $StrongEquiv(A, \ell^1)$ implied by the RIP. In this section, we incorporate the bounds on RIP constants established in Section 2 into a known condition implying $StrongEquiv(A, \ell^1)$ obtained from an RIP analysis. Following the pioneering work of Candès, Romberg, and Tao [8, 11], many different conditions on the RIP constants have been developed which ensure recovery of every

¹For some algorithms, such as ℓ^1 -regularization, these two parameters fully characterize the behavior of the algorithm for a particular matrix ensemble, whereas for other algorithms, such as OMP, the distribution of the nonzero coefficients also influences the behavior of the method.

²Similar results have been proven for other random matrix ensembles, but they are even less precise than those for the Gaussian distribution.

k -sparse vector via ℓ^1 -minimization, [6, 7, 9, 10, 46] to name a few. The current state of the art RIP conditions for ℓ^1 -minimization were developed by Foucart and Lai [32].

THEOREM 3.3 (Foucart and Lai [32]). *For any matrix A of size $n \times N$ with RIP constants $L(2k, n, N)$ and $U(2k, n, N)$, for $2k \leq n < N$. Define*

$$(3.1) \quad \mu^{FL}(k, n, N) := \frac{1 + \sqrt{2}}{4} \left(\frac{1 + U(2k, n, N)}{1 - L(2k, n, N)} - 1 \right).$$

If $\mu^{FL}(k, n, N) < 1$, then there is $StrongEquiv(A, \ell^1)$.

To translate this result into the phase transition framework for matrices from the Gaussian ensemble, we employ the RIP bounds (2.7) to the asymmetric RIP constants $L(2k, n, N)$ and $U(2k, n, N)$. It turns out that naively inserting these bounds into (3.1) yields a bound on $\mu^{FL}(k, n, N)$, see Lemma 3.6, and provides a simple way to obtain a bound on the region of strong equivalence.

DEFINITION 3.4 (RIP Region of $StrongEquiv(A, \ell^1)$). *Define*

$$(3.2) \quad \mu^{FL}(\delta, \rho) := \frac{1 + \sqrt{2}}{4} \left(\frac{1 + \mathcal{U}(\delta, 2\rho)}{1 - \mathcal{L}(\delta, 2\rho)} - 1 \right)$$

and $\rho_S^{FL}(\delta)$ as the solution to $\mu^{FL}(\delta, \rho) = 1$.

The function $\rho_S^{FL}(\delta)$ is displayed as the red curve in Figure 3.1.

THEOREM 3.5. *Fix $\epsilon > 0$. Consider the proportional-growth asymptotic, Definition 1.2, with parameters $(\delta, \rho) \in (0, 1) \times (0, 1/2)$. Draw the corresponding $n \times N$ matrices A from the Gaussian ensemble. If $\rho < (1 - \epsilon)\rho_S^{FL}(\delta)$, then $Prob(StrongEquiv(A, \ell^1)) \rightarrow 1$ as $n \rightarrow \infty$.*

Therefore the function $\rho_S^{FL}(\delta)$ bounds a region of strong equivalence for ℓ^1 -minimization.

Theorem 3.5 follows from Theorem 3.3 and the validity of the probabilistic bounds on the RIP constants, Theorem 2.3. In particular, Lemma 3.6 bounds $\mu^{FL}(k, n, N)$ in terms of the asymptotic RIP bounds $\mathcal{L}(\delta, 2\rho)$ and $\mathcal{U}(\delta, 2\rho)$, by the quantity $\mu^{FL}(\delta, (1 + \epsilon)\rho)$ defined in (3.3). If $\rho_\epsilon(\delta)$ is the solution to $\mu^{FL}(\delta, (1 + \epsilon)\rho) = 1$, then for $\rho < \rho_\epsilon(\delta)$ we achieve the desired bound, $\mu^{FL}(k, n, N) < 1$, to ensure $StrongEquiv(A, \ell^1)$. The statement of Theorem 3.5 follows from relating $\rho_\epsilon(\delta)$ to $\rho_S^{FL}(\delta)$, the solution to $\mu^{FL}(\delta, \rho) = 1$.

LEMMA 3.6. *Fix $\epsilon > 0$. Consider the proportional-growth asymptotic with parameters $(\delta, \rho) \in (0, 1) \times (0, 1/2)$. Draw the corresponding $n \times N$ matrices A from the Gaussian ensemble. Then*

$$(3.3) \quad Prob(\mu^{FL}(k, n, N) < \mu^{FL}(\delta, (1 + \epsilon)\rho)) \rightarrow 1$$

exponentially in n .

Proof. Theorem 2.3 and the form of $\mu^{FL}(\delta, \rho)$ imply a similar bound to the above with a modified dependence on ϵ . For any $c\epsilon > 0$, with $n/N \rightarrow \delta \in (0, 1)$ and $k/n \rightarrow \rho \in (0, 1/2]$, the probability, on the draw of A from the Gaussian ensemble, that

$$(3.4) \quad \mu^{FL}(k, n, N) < \frac{1 + \sqrt{2}}{4} \left(\frac{1 + \mathcal{U}(\delta, 2\rho) + c\epsilon}{1 - \mathcal{L}(\delta, 2\rho) - c\epsilon} - 1 \right).$$

is satisfied converges to one exponentially with n . Since $\mathcal{U}(\delta, \rho)$ is non-decreasing in ρ and $\mathcal{L}(\delta, \rho)$ is strictly increasing in ρ for any δ and $\rho \in (0, 1)$, it follows that the right-hand side of (3.4) can be bounded by the right-hand side of (3.3) for any fixed ϵ satisfying $0 < \epsilon < \frac{1}{2\rho} - 1$, by setting

$$c := \frac{\rho}{2} \left. \frac{\partial \mathcal{L}(\delta, z)}{\partial z} \right|_{z=2(1+\epsilon)\rho} > 0.$$

(The upper bound on ϵ is imposed so that the second argument of $\mathcal{U}(\delta, \cdot)$ and $\mathcal{L}(\delta, \cdot)$, $2(1 + \epsilon)\rho$, is in the admissible range of $(0, 1)$.) That the bound (3.3) is satisfied for all $\epsilon > 0$ sufficiently small, and

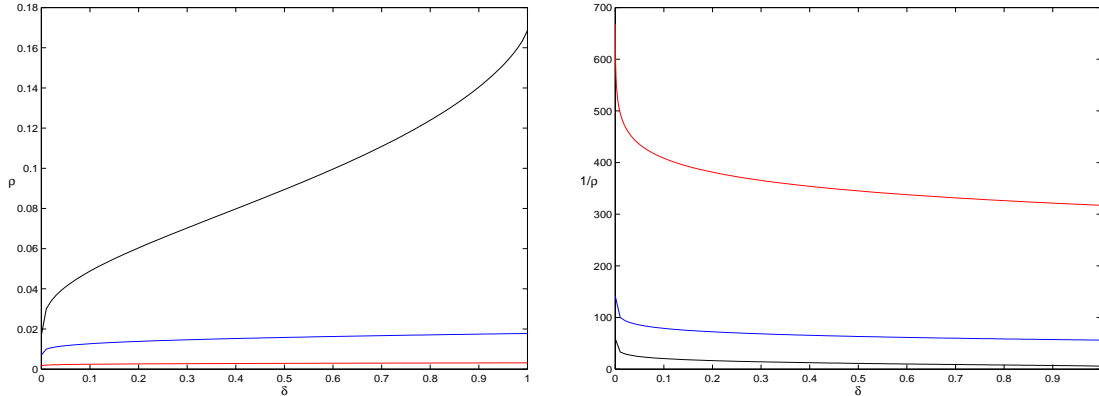


FIG. 3.1. Left panel: Three lower bounds on the Strong ℓ^1/ℓ^0 -equivalence phase transition, $\rho_S(\delta)$, for Gaussian random matrices from Theorem 3.7 ($\rho_S^D(\delta)$, black), Theorem 3.9 ($\rho_S^{RV}(\delta)$, blue), and Theorem 3.5 ($\rho_S^{FL}(\delta)$, red). Right panel: The inverse of the ℓ^1/ℓ^0 -equivalence phase transition lower bounds in the left panel.

that the right hand side of (3.3) is strictly increasing in ϵ establishes that (3.3) is in fact satisfied probability on the draw of A that converges to one exponentially in n for any $\epsilon \in \left(0, \frac{1}{2\rho} - 1\right)$. \square

Proof. [Theorem 3.5] Let $\rho_\epsilon(\delta)$ be the solution of $\mu^{FL}(\delta, (1+\epsilon)\rho) = 1$. Then, for any $\rho < \rho_\epsilon(\delta)$, Lemma 3.6 implies that $\mu^{FL}(k, n, N) < 1$, which by Theorem 3.3, ensures $StrongEquiv(A, \ell^1)$. To remove the dependence on the level curve $\rho_\epsilon(\delta)$, note that $\rho_\epsilon(\delta)$ is related to $\rho_S^{FL}(\delta)$, the solution of $\mu^{FL}(\delta, \rho) = 1$, by $(1+\epsilon)\rho_\epsilon(\delta) \equiv \rho_S^{FL}(\delta)$. Since $(1-\epsilon) < (1+\epsilon)^{-1}$ for all $\epsilon > 0$, we have $(1-\epsilon)\rho_S^{FL}(\delta) < \rho_\epsilon(\delta)$. Thus, provided $\rho < (1-\epsilon)\rho_S^{FL}(\delta)$, the statement of Theorem 3.5 is satisfied. \square

3.2. Comparison of bounds on $StrongEquiv(A, \ell^1)$. In this section we use the phase transition framework to readily compare bounds on the region of $StrongEquiv(A, \ell^1)$ obtained from vastly different methods of analysis. In Section 3.1, we have already determined the region of strong equivalence for ℓ^1 -minimization obtained by using the RIP. Here we look at two other examples, namely Donoho’s polytope results [16, 18] and the sufficient condition of Rudelson and Vershynin [46] obtained from geometric functional analysis. We do not go into great details about how the results were obtained, but simply point out that the methods of analysis are rather different. As a result, the original statements of the theorems take drastically different forms and are therefore difficult to compare even qualitatively. Translating the results into the phase transition framework, however, offers a direct, quantitative, and simple method of comparison.

Using polytope theory and the notion of central-neighborliness, Donoho [16] defined a function $\rho_S^D(\delta)$ which defines a region of the (δ, ρ) phase space ensuring $StrongEquiv(A, \ell^1)$ with probability on the draw of A converging to one exponentially in n . The phase transition function $\rho_S^D(\delta)$ is displayed as the black curve in Figure 3.1.

THEOREM 3.7 (Donoho [16]). *Fix $\epsilon > 0$. Consider the proportional-growth asymptotic, Definition 1.2, with parameters $(\delta, \rho) \in (0, 1) \times (0, 1/2)$. Sample each $n \times N$ matrix A from the Gaussian ensemble. Suppose $\rho < (1-\epsilon)\rho_S^D(\delta)$. Then $Prob(StrongEquiv(A, \ell^1)) \rightarrow 1$ as $n \rightarrow \infty$.*

Therefore $\rho_S^D(\delta)$ bounds a region of strong equivalence for ℓ^1 -minimization.

Rudelson and Vershynin [46] used an alternative geometric approach from geometric functional analysis (GFA) to determine regions of $StrongEquiv(A, \ell^1)$ for Gaussian and random partial Fourier matrices. For Gaussian matrices their elegantly simple proof involves employing Gordon’s “escape through the mesh theorem” on the nullspace of A . Their lower bound on the region of $StrongEquiv(A, \ell^1)$ is larger for the Gaussian ensemble than for the Fourier ensemble. We restate their condition for the Gaussian ensemble in the proportional growth asymptotic.

DEFINITION 3.8 (GFA Region of $StrongEquiv(A, \ell^1)$). Define

$$\gamma(\rho\delta) := \exp\left(\frac{\log(1 + 2\log(e/\rho\delta))}{4\log(e/\rho\delta)}\right),$$

$$(3.5) \quad \mu^{RV}(\delta, \rho) := \rho(12 + 8\log(1/\rho\delta) \cdot \gamma^2(\rho\delta)),$$

and $\rho_S^{RV}(\delta)$ as the solution to $\mu^{RV}(\delta, \rho) = 1$.

The function $\rho_S^{RV}(\delta)$ is displayed as the blue curve in Figure 3.1.

THEOREM 3.9 (Rudelson and Vershynin [46]). Fix $\epsilon > 0$. Consider the proportional-growth asymptotic, Definition 1.2, with parameters $(\delta, \rho) \in (0, 1) \times (0, 1/2)$. Sample each $n \times N$ matrix A from the Gaussian ensemble. Suppose $\rho < (1 - \epsilon)\rho_S^{RV}(\delta)$. Then $\text{Prob}(StrongEquiv(A, \ell^1)) \rightarrow 1$ as $n \rightarrow \infty$.

Therefore $\rho_S^{RV}(\delta)$ bounds a region of strong equivalence for ℓ^1 -minimization.

Versions of Theorems 3.7 and 3.9 exist for finite values of (k, n, N) , [23, 46], but in each case the recoverability conditions rapidly approach the stated asymptotic limiting functions $\rho_S(\delta)$ as (k, n, N) grow; we do not further complicate the discussion with their rates of convergence.

Since Theorems 3.5, 3.7, and 3.9 provide a region of $StrongEquiv(A, \ell^1)$, we now have three subsets of the exact region of $StrongEquiv(A, \ell^1)$. Although Theorems 3.5, 3.7, and 3.9 each have the same goal of quantifying the exact boundary of $StrongEquiv(A, \ell^1)$ for Gaussian random matrices, they are arrived at using substantially different methods of analysis. The efficacy of the bounds from the largest region of $StrongEquiv(A, \ell^1)$ to the smallest region are $\rho_S^D(\delta)$ of Donoho, $\rho_S^{RV}(\delta)$ of Rudelson and Vershynin, and $\rho_S^{FL}(\delta)$ of Foucart and Lai, see the left panel of Figure 3.1. From the inverse of $\rho_S(\delta)$, see the right panel of Figure 3.1, we can read the constant of proportionality where the associated method of analysis guarantees $StrongEquiv(A, \ell^1)$: from Theorems 3.7, 3.9, and 3.3 they are bounded below by: $n \geq 5.9k$, $n \geq 56k$, and $n \geq 317k$ respectively.

3.3. Further Considerations. The phase transition framework can also be used to quantify what has been proven about an encoder/decoder pair's speed of convergence, its degree of robustness to noise, and to make comparisons of these properties between different algorithms. A general framework for expressing the results of RIP based analyses as statements in the phase transition framework is presented in [3], where it is also applied to three exemplar greedy algorithms CoSaMP [43], Subspace Pursuit [15], and Iterated Hard Thresholding [4]. Bounds on regions of $StrongEquiv(A, \ell^q)$ for ℓ^q -minimization for $q \in (0, 1]$ implied by the RIP are available in an extended technical report [1], where the effects of noise are also considered. Through these objective measures of comparison we hope to make clear the proven efficacy of sparse approximation algorithms and allow for their transparent comparison.

In this article, we have considered only the case of noiseless measurements, Regions of Strong Equivalence, and a particular result obtained via an eigenvalue analysis and the RIP. We briefly discuss some additional considerations for the phase transition framework.

3.3.1. Phase Transitions with Noisy Measurements. In a practical setting, it is more reasonable to assume that the measurements are corrupted by noise, $y = Ax + e$ for some noise vector e . The RIP has played a vital role in establishing stable signal recovery in the presence of noise for many decoders. When noise is present, the curves $\rho_S(\delta)$ bounding regions of strong equivalence serve as an upper bound to the curves depicting the regions of the phase plane which guarantee stable recovery. The RIP constants also describe how significantly the noise will be amplified by the encoder/decoder pairing, details are available for the Gaussian encoder and ℓ^q -minimization decoder [1] and greedy decoders [3] CoSaMP, Subspace Pursuit, and Iterated Hard Thresholding. Hassibi and Xu have developed a stability analysis of ℓ^1 -minimization from the analysis of convex polytopes [53], establishing substantially larger stability regions than the regions implied by the RIP.

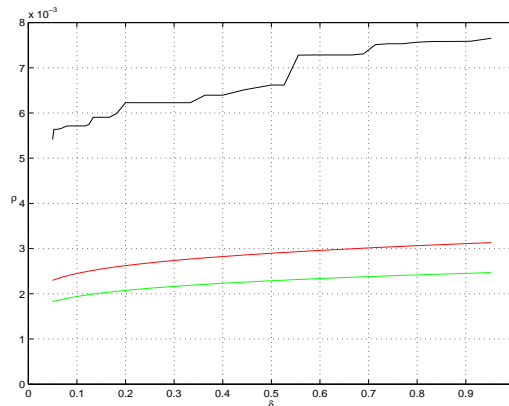


FIG. 3.2. Example Improvements on Bounds on the Strong ℓ^1/ℓ^0 -equivalence phase transition, $\rho_S(\delta)$, for Gaussian random matrices: ($\rho_S^C(\delta)$, green), Theorem 3.5 ($\rho_S^{FL}(\delta)$, red), and ($\rho_S^{emp}(\delta)$, black).

3.3.2. Regions of Weak Equivalence and Average Case Performance. In many applications, it may not be imperative that the decoder is able to reconstruct every k -sparse vector. Instead, one may be willing to lose a small fraction of all possible k -sparse signals. This is the behavior observed when a decoder is tested on k -sparse vectors whose support sets are drawn uniformly at random. Large scale empirical testing of CoSaMP, Subspace Pursuit and Iterated Hard Thresholding were compiled by Donoho and Maleki [19]. Most sparse approximation algorithms do not have a theoretical average case analysis. The polytope analysis of Donoho and Tanner allows for analytical arguments providing a *Region of Weak Equivalence* where recovery is guaranteed for all but a but a small fraction of k -sparse signals. An average case variant of the RIP is being developed, see [49].

3.3.3. Improving the RIP Phase Transition. It is possible that Thm. 3.3 could be improved with alternative methods of analysis. For example, Thm. 3.3 built off the work of Candès, Romberg, and Tao [7, 9, 10]. In [7], Candès proved that if $R(2k, n, N) < \sqrt{2}-1$, then ℓ^1 -minimization will successfully recover every k -sparse vector. An asymmetric analysis and translation into the Strong Equivalence terminology of Sec. 3.1 produces a function $\rho_S^C(\delta)$ which bounds a region of strong equivalence. The alternative methods of Foucart and Lai leading to Thm. 3.3 provided a larger region of strong equivalence. See Fig. 3.2.

Alternatively, the region of strong equivalence might be increased by improving the bounds on the RIP constants, $\mathcal{L}(\delta, \rho), \mathcal{U}(\delta, \rho)$. If the method of analysis remained the same, we can explore the effects of improved bounds by examining the statements with empirically observed *lower bounds* on RIP constants for Gaussian matrices. As detailed in Table 2.1, the current bounds from Thm. 2.3 are no more than twice the empirical RIP constants. Replacing the RIP constants with empirically observed lower bounds of the RIP constants (for $n = 800$) in $\mu^{FL}(k, n, N)$ gives us a function $\rho_S^{emp}(\delta)$, see Figure 3.2, which is an upper bound on the region of strong equivalence implied by Thm. 3.3; this improvement is no more than 2.5 times $\rho_S^{FL}(\delta)$ for $\delta \in [1/20, 20/21]$.

Acknowledgements. The authors would like to thank the editor and the referees for their useful suggestions that have greatly improved the manuscript.

REFERENCES

- [1] J. D. Blanchard, C. Cartis, and J. Tanner. Compressed sensing: how sharp is the restricted isometry property. extended technical report, 2009.
- [2] J. D. Blanchard, C. Cartis, and J. Tanner. Decay properties for restricted isometry constants. *IEEE Signal Proc. Letters*, 16(7):572–575, 2009.

- [3] J. D. Blanchard, C. Cartis, J. Tanner, and A. Thompson. Phase transitions for greedy sparse approximation algorithms. submitted, 2009.
- [4] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comp. Harm. Anal.*, 27(3):265–274, 2009.
- [5] A. M. Bruckstein, David L. Donoho, and Michael Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, 2009.
- [6] E. J. Candès. Compressive sampling. In *International Congress of Mathematicians. Vol. III*, pages 1433–1452. Eur. Math. Soc., Zürich, 2006.
- [7] E. J. Candès. The restricted isometry property and its implications for compressed sensing. *C. R. Math. Acad. Sci. Paris*, 346(9-10):589–592, 2008.
- [8] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.
- [9] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223, 2006.
- [10] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51(12):4203–4215, 2005.
- [11] E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, 2006.
- [12] R. Chartrand. Exact reconstructions of sparse signals via nonconvex minimization. *IEEE Signal Process. Lett.*, 14:707–710, 2007.
- [13] R. Chartrand and V. Staneva. Restricted isometry properties and nonconvex compressive sensing. *Inverse Problems*, 24(035020):1–14, 2008.
- [14] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159 (electronic), 2001. Reprinted from *SIAM J. Sci. Comput.* 20 (1998), no. 1, 33–61.
- [15] W. Dai and O. Milenkovic. Subspace pursuit for compressive sensing signal reconstruction. *IEEE Trans. Inform. Theory*, 55(5):2230–2249, 2009.
- [16] D. L. Donoho. Neighborly polytopes and sparse solution of underdetermined linear equations. Technical Report, Department of Statistics, Stanford University, 2005.
- [17] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
- [18] D. L. Donoho. High-dimensional centrally symmetric polytopes with neighborliness proportional to dimension. *Discrete Comput. Geom.*, 35(4):617–652, 2006.
- [19] D. L. Donoho and A. Maleki. Optimally tuned iterative thresholding algorithms for compressed sensing. *IEEE Sel. Topics Signal Processing*, in press.
- [20] D. L. Donoho and V. Stodden. Breakdown point of model selection when the number of variables exceeds the number of observations. In *Proceedings of the International Joint Conference on Neural Networks*, 2006.
- [21] D. L. Donoho and J. Tanner. Sparse nonnegative solutions of underdetermined linear equations by linear programming. *Proc. Natl. Acad. Sci. USA*, 102(27):9446–9451, 2005.
- [22] D. L. Donoho and J. Tanner. Counting faces of randomly projected polytopes when the projection radically lowers dimension. *J. AMS*, 22(1):1–53, 2009.
- [23] D. L. Donoho and J. Tanner. Exponential bounds implying construction of compressed sensing matrices, error-correcting codes and neighborly polytopes by random sampling. *IEEE Trans. on Information Theory*, 2010. in press.
- [24] D. L. Donoho and Y. Tsaig. Fast solution of l1 minimization problems when the solution may be sparse. *IEEE Trans. Inform. Theory*, 54(11):4789–4812, 2008.
- [25] D. L. Donoho, Y. Tsaig, I. Drori, and J.-L. Stark. Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit. *IEEE Trans. Inform. Theory*, submitted.
- [26] R. Dorfman. The detection of defective members of large populations. *Ann. Math. Statist.*, 14(4):436–440, 1943.
- [27] C. Dossal, G. Peyré, and J. Fadili. A numerical exploration of compressed sampling recovery. *Linear Algebra Appl.*, 432(7):1663–1679, 2010.
- [28] M. F. Duarte, M. A. Davenport, D. Takhar, Laska J. N., T. Sun, K. F. Kelly, and R. G. Baraniuk. Single-pixel imaging via compressed sampling. *IEEE Signal Processing Magazine*, 25(2):83–91, 2008.
- [29] A. Edelman. Eigenvalues and condition numbers of random matrices. *SIAM J. Matrix Anal. Appl.*, 9(4):543–560, 1988.
- [30] A. Edelman and N. R. Rao. Random matrix theory. *Acta Numer.*, 14:233–297, 2005.
- [31] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE J. Sel. Topics Signal Process.*, 1(4):586–597, 2007.
- [32] S. Foucart and M.-J. Lai. Sparsest solutions of underdetermined linear systems via ℓ_q -minimization for $0 < q \leq 1$. *Appl. Comput. Harmon. Anal.*, 26(3):395–407, 2009.
- [33] S. Geman. A limit theorem for the norm of random matrices. *Ann. Probab.*, 8(2):252–261, 1980.
- [34] A. C. Gilbert, M. A. Iwen, and M. J. Strauss. Group testing and sparse signal recovery. In *42nd Asilomar Conference on Signals, Systems, and Computers*, 2008.
- [35] R. Gribonval and M. Nielsen. Sparse representations in unions of bases. *IEEE Trans. Inform. Theory*, 49(12):3320–3325, 2003.
- [36] A. K. Hartmann and H. Rieger. *New Optimization Algorithms in Physics*. Wiley VCH, Cambridge, 2006.

- [37] A. K. Hartmann and M. Weight. *Phase Transitions in Combinatorial Optimization Problems*. Wiley VCH, Cambridge, 2005.
- [38] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *Journal of Machine Learning Research*, 11:451–487, 2010.
- [39] M. Lustig, D. L. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
- [40] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing MRI. *IEEE Signal Processing Magazine*, 25(2):72–82, 2008.
- [41] M. Hedjazi Moghari, M. Ak?akaya, A. O’Connor, P. Hu, V. Tarokh, W. J. Manning, and R. Nezafat. CoSMo: Compressed sensing motion correction for coronary MRI. The Annual Scientific Meeting of International Society for Magnetic Resonance in Medicine (ISMRM), 2010.
- [42] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, 1995.
- [43] D. Needell and J. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Appl. Comp. Harm. Anal.*, 26(3):301–321, 2009.
- [44] D. Needell and R. Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. *Foundations of Comp. Math.*, 9(3):317–334, 2009.
- [45] V. M. Patel, G. R. Easley, D. M. Jr Healy, and R. Chellappa. Compressed synthetic aperture radar. *IEEE Sel. Topics Signal Processing*, in press.
- [46] M. Rudelson and R. Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Comm. Pure Appl. Math.*, 61(8):1025–1045, 2008.
- [47] R. Saab and O. Yilmaz. Sparse recovery by non-convex optimization—instance optimality. *Appl. Comp. Harm. Anal.*, in press.
- [48] J. W. Silverstein. The smallest eigenvalue of a large-dimensional Wishart matrix. *Ann. Probab.*, 13(4):1364–1368, 1985.
- [49] J. Tropp. On the conditioning of random subdictionaries. *Appl. Comp. Harm. Anal.*, 25(1):1–24, 2008.
- [50] J. A. Tropp. Greed is good: algorithmic results for sparse approximation. *IEEE Trans. Inform. Theory*, 50(10):2231–2242, 2004.
- [51] E. van den Berg and M. P. Friedlander. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008.
- [52] E. T. Whittaker and G. N. Watson. *A course of modern analysis*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 1996. An introduction to the general theory of infinite processes and of analytic functions; with an account of the principal transcendental functions, Reprint of the fourth (1927) edition.
- [53] W. Xu and B. Hassibi. Compressed sensing over the grassmann manifold: A unified analytical framework. Forty-Sixth Annual Allerton Conference, 2008.
- [54] W. Yin, S. Osher, D. Goldfarb, and J. Darbon. Bregman iterative algorithms for ℓ^1 -minimization with applications to compressed sensing. *SIAM Journal on Imaging Science*, 1(1):143–168, 2008.