

ON THE CONDITIONING OF RANDOM SUBDICTIONARIES

JOEL A. TROPP

ABSTRACT. An important problem in the theory of sparse approximation is to identify well-conditioned subsets of vectors from a *general* dictionary. In most cases, current results do not apply unless the number of vectors is smaller than the square root of the ambient dimension, so these bounds are too weak for many applications. This paper shatters the square-root bottleneck by focusing on *random* subdictionaries instead of *arbitrary* subdictionaries. It provides explicit bounds on the extreme singular values of random subdictionaries that hold with overwhelming probability. The results are phrased in terms of the coherence and spectral norm of the dictionary, which capture information about its global geometry. The proofs rely on standard tools from the area of Banach space probability.

As an application, the paper shows that the conditioning of a subdictionary is the major obstacle to the uniqueness of sparse representations and the success of ℓ_1 minimization techniques for signal recovery. Indeed, if a fixed subdictionary is well conditioned and its cardinality is slightly smaller than the ambient dimension, then a random signal formed from this subdictionary almost surely has no other representation that is equally sparse. Moreover, with overwhelming probability, the maximally sparse representation can be identified via ℓ_1 minimization.

Note that the results in this paper are not directly comparable with recent work on subdictionaries of *random* dictionaries.

1. INTRODUCTION

To motivate the results of this article, we begin with *sparse representation*, the problem of finding a sparse solution to an underdetermined system of linear equations. Fix a dictionary Φ , which is a $d \times N$ complex matrix whose columns have unit ℓ_2 norm. Suppose we are given a signal \mathbf{s} that is formed as a linear combination of m columns from Φ :

$$\mathbf{s} = \Phi \mathbf{c} \quad \text{where} \quad \|\mathbf{c}\|_0 = m.$$

(The ℓ_0 quasi-norm $\|\cdot\|_0$ counts the number of nonzero components in its argument.) The sparse representation problem asks us to determine the vector \mathbf{c} of coefficients, given only the dictionary Φ and the observed signal \mathbf{s} . In particular, we must locate the nonzero components of \mathbf{c} to determine which columns participate in the signal. Observe that the active columns must form a linearly independent set before this problem is even well posed.

One of the central challenges in sparse representation, therefore, is to identify the linearly independent collections of columns from a dictionary. In fact, many applications demand stronger information on the conditioning of column submatrices in the form of bounds for their extreme singular values. Ideally, these results should depend only on global properties of the dictionary matrix and on the number of columns extracted.

Date: August 2006. Revised: 28 August 2007.

Key words and phrases. Banach space probability, Basis Pursuit, coherence, convex optimization, dictionary, random matrix, regularization, sparsity, uncertainty principle, underdetermined linear system.

JAT is with Applied & Computational Mathematics, MC 217-50, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125-5000. E-mail: jtropp@acm.caltech.edu. This work was supported in part by NSF DMS 0503299 and by DARPA/ONR N660010612011.

The *coherence* of a dictionary encapsulates information about its global geometry, and it can be used to study the conditioning of its column submatrices. The coherence is defined by

$$\mu \stackrel{\text{def}}{=} \max_{j \neq k} |\langle \varphi_j, \varphi_k \rangle|$$

where φ_k denotes the k th column of Φ . In words, μ bounds the cosine of the acute angle between pairs of columns from the dictionary. The following easy result connects the coherence with the extreme singular values of a column submatrix from the dictionary. Here and elsewhere, $\|\cdot\|$ denotes the spectral norm (i.e., the norm on linear maps from ℓ_2 to ℓ_2), and the letter \mathbf{I} represents a conformal identity matrix.

Proposition 1. *Let Φ be a dictionary with coherence μ , and let \mathbf{A} be an arbitrary m -column submatrix of Φ . Then*

$$\|\mathbf{A}^* \mathbf{A} - \mathbf{I}\| \leq (m - 1)\mu$$

In particular, every collection of m columns is linearly independent when $(m - 1)\mu < 1$.

To interpret this theorem, one should note that

$$\|\mathbf{A}^* \mathbf{A} - \mathbf{I}\| = \max \{ \sigma_{\max}^2(\mathbf{A}) - 1, 1 - \sigma_{\min}^2(\mathbf{A}) \},$$

where $\sigma_{\max}^2(\mathbf{A})$ and $\sigma_{\min}^2(\mathbf{A})$ indicate the extreme eigenvalues of $\mathbf{A}^* \mathbf{A}$.

Proof sketch. Apply Gershgorin's Theorem to the Gram matrix $\mathbf{A}^* \mathbf{A}$. See [DE03, Thm. 5] for details. \square

This result is quite attractive, but its utility is limited by the fact that the coherence cannot be very small [SH03, Thm. 2.3]:

$$\mu \geq \sqrt{\frac{N - d}{d(N - 1)}}. \tag{1.1}$$

When $N \geq 2d$, it follows that $\mu \geq (2d - 1)^{-1/2}$. In this common parameter regime, Proposition 1 only yields information on the smallest singular value when $m = O(\sqrt{d})$. It would be valuable to understand larger collections of columns, but deterministic methods cannot provide a much better assessment than Proposition 1. Instead of asking about arbitrary sets of columns, therefore, we study *random* collections of columns from the dictionary. This change of focus allows us to shatter the square-root bottleneck.

1.1. Contributions. This paper provides explicit bounds for the extreme singular values of random collections of columns from a *general dictionary*. The results are phrased in terms of the global geometry of the dictionary, and they are vastly superior to the simple bounds of Proposition 1. The proofs involve standard methods from the field of Banach space probability, such as the systematic use of symmetrization and (noncommutative) Khintchine inequalities. Although the approach is not new, this paper provides satisfying answers to fundamental questions, and we believe that both the results and the methods will be valuable to the computational harmonic analysis community.

1.2. Outline. A brief outline is standard at this point. Section 2 offers a snapshot of the most important theorems, and Section 3 compares our findings with earlier work from the fields of sparse approximation and Banach space theory. The key technical tools are introduced without proof in Section 4. Sections 5 and 6 develop the main results upon these foundations. Sections 7 and 8 discuss applications of these results to problems in sparse representation. Finally, Sections 9 and 10 contain complete proofs of the technical theorems, which is valuable because the functional analysis literature tends to be indifferent about the precise values of numerical constants.

2. BACKGROUND AND MAJOR RESULTS

This section presents some background material and sets the notation for the rest of the work. Then it states four representative theorems that summarize the major contributions of the paper.

2.1. Background and notation. For each $p \geq 1$, we write $\|\cdot\|_p$ for the usual ℓ_p vector norm, while the symbol $\|\cdot\|_{p,q}$ denotes the norm on linear maps from ℓ_p to ℓ_q . In this work, two norms are especially important:

- The quantity $\|\mathbf{A}\|_{1,2}$ is the maximum ℓ_2 norm of a column of \mathbf{A} .
- The spectral norm $\|\mathbf{A}\|_{2,2}$ returns the largest singular value of \mathbf{A} . It is always written as $\|\cdot\|$ without any subscripts.

For a natural number N , abbreviate the set $\{1, 2, \dots, N\}$ by the symbol $\llbracket N \rrbracket$. Given a subset Ω of $\llbracket N \rrbracket$, we write \mathbb{C}^Ω for the set of functions mapping Ω to \mathbb{C} , equipped with the usual addition and scalar multiplication to form a linear space.

The symbol $\mathbb{P}\{\cdot\}$ denotes the probability operator, which returns the probability of a given event. The letter \mathbb{E} indicates the expectation operator. We use a special notation for conditional expectation. For example, if Z is a random variable, then \mathbb{E}_Z denotes integration with respect to Z , holding other variables fixed.

The (positive homogeneous) q th moment of a random variable is computed via the expression $(\mathbb{E}|\cdot|^q)^{1/q}$. This expression defines the L_q norm on the space of complex-valued random variables. In particular, it verifies a triangle inequality:

$$(\mathbb{E}|Y + Z|^q)^{1/q} \leq (\mathbb{E}|Y|^q)^{1/q} + (\mathbb{E}|Z|^q)^{1/q}.$$

Throughout this work, the word *random*, without additional qualification, always means “uniformly random over the specified set.” For example, the phrase “a random m -column submatrix” means that each m -column submatrix is drawn with equal probability.

We follow the analysts’ convention that upright Roman letters (c, C, D, \dots) denote absolute constants that may change at each appearance. Finally, be aware that all logarithms in this paper have base e .

2.2. The dictionary. The term *dictionary* refers to a $d \times N$ complex matrix Φ whose columns have unit ℓ_2 norm. Without loss of generality, we may assume that the columns of Φ span \mathbb{C}^d , which implies that $N \geq d$. A column submatrix of Φ is referred to as a *subdictionary*. If T is a subset of $\llbracket N \rrbracket$, we sometimes write Φ_T for the subdictionary consisting of the columns in T .

After the coherence, the spectral norm $\|\Phi\|$ is the most important geometric quantity associated with the dictionary. It measures how much the dictionary matrix can dilate a unit-norm coefficient vector, so it reflects how much the columns of Φ are “spread out.” Using Hölder’s inequality, one may develop a lower bound on the spectral norm of the dictionary:

$$\|\Phi\|^2 = \|\Phi^* \Phi\| \geq d^{-1} \text{trace}(\Phi^* \Phi) = \frac{N}{d}. \quad (2.1)$$

When equality holds in this relation, the dictionary is called a *unit-norm tight frame*. Equivalently, the rows of Φ are mutually orthogonal vectors with equal norms.

2.3. Pairs of orthonormal bases. One frequently encounters dictionaries that are formed as a concatenation of two orthonormal bases:

$$\Phi = [\Phi_1 \quad \Phi_2],$$

where Φ_1 and Φ_2 are both $d \times d$ unitary matrices. The paradigmatic example is the Fourier–Dirac dictionary, in which Φ_1 is the discrete Fourier transform matrix and Φ_2 is the identity matrix. This dictionary is connected with discrete uncertainty principles [DS89, DH01].

The coherence of a pair of orthonormal bases is calculated as

$$\mu = \max_{j,k} \left| \left\langle \varphi_j^{(1)}, \varphi_k^{(2)} \right\rangle \right| \quad \text{for } j, k = 1, 2, \dots, d.$$

In this setting, the coherence satisfies a more stringent bound:

$$\mu \geq d^{-1/2}. \tag{2.2}$$

Note that the Fourier–Dirac dictionary has coherence $\mu = d^{-1/2}$.

For a pair of orthonormal bases, we can strengthen Proposition 1 substantially [DH01, DE03].

Proposition 2. *Let Φ be a pair of orthonormal bases. Suppose the subdictionary \mathbf{A} contains $|\Omega|$ columns from the first basis and $|T|$ columns from the second basis. Then*

$$\|\mathbf{A}^* \mathbf{A} - \mathbf{I}\| \leq \mu \sqrt{|\Omega| |T|}.$$

In particular, \mathbf{A} has linearly independent columns whenever $|\Omega| |T| < \mu^{-2}$.

Proof sketch. Let \mathbf{B} denote the off-diagonal block of $\mathbf{A}^* \mathbf{A}$. Observe that $\|\mathbf{A}^* \mathbf{A} - \mathbf{I}\| = \|\mathbf{B}\|$, and apply the estimate $\|\mathbf{B}\|^2 \leq \|\mathbf{B}\|_{1,1} \|\mathbf{B}\|_{\infty, \infty}$. \square

2.4. Results for incoherent dictionaries. This subsection presents two central results about the conditioning of random subdictionaries drawn from an incoherent dictionary. The proofs appear in Section 6. The first theorem focuses on the special case of a unit-norm tight frame with typical coherence. These hypotheses enable us to offer a transparent statement that illuminates the pessimism of the predictions in Proposition 1.

Theorem A. *Let Φ be a unit-norm tight frame for \mathbb{C}^d with at least $2d$ columns. Suppose that \mathbf{X} is a random m -column subdictionary. Then*

$$\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \leq C \sqrt{\mu^2 m \log(m+1)},$$

provided that the right-hand side is less than one.

It is instructive to compare Theorem A with Proposition 1. Consider the important case where $\mu = d^{-1/2}$. The theorem predicts that most m -column subdictionaries are linearly independent when $m = \text{const} \cdot d / \log d$. In sharp contrast, the proposition allows that all m -column subdictionaries are linearly independent only when $m \leq \sqrt{d}$. The price we pay for the better bound is a weakening from certainty about the conditioning to near certainty. Our arguments do not yield a good estimate for the constant; the proof delivers a value of about 12. As a consequence, the new result does not quantitatively improve on the older result until the dimension d is quite large.

Theorem A is straightforward, but it does not take full advantage of the information available. In particular, our methods yield strong estimates on the probability that a random subdictionary is well conditioned. Here is a more detailed (i.e., complicated) result.

Theorem B. *Let Φ be a dictionary, and let \mathbf{X} be a random m -column subdictionary. The condition*

$$\sqrt{\mu^2 m \log(m+1)} \cdot s + \frac{m}{N} \|\Phi\|^2 \leq c\delta \quad \text{with } s \geq 1$$

implies that

$$\mathbb{P} \{ \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \geq \delta \} \leq m^{-s}.$$

In words, the number δ controls the range of the extreme singular values of a random subdictionary. The probability that a subdictionary is ill conditioned drops exponentially fast as the parameter s increases. This result also exposes the participation of the spectral norm, which was hidden in Theorem A because of the hypotheses.

To appreciate why Theorem B improves on Proposition 1, just compare the term $\mu^2 m \log m$ here against the term μm in the older result. On account of (2.1), we also have the bound

$$\frac{m}{N} \|\Phi\|^2 \geq \frac{m}{d}.$$

This contribution ensures that—no matter how small the coherence—the estimate of the smallest singular value is vacuous when the size of the subdictionary is close to the ambient dimension. Of course, this condition is natural because we cannot construct a linearly independent collection of more than d vectors in d dimensions.

2.5. Results for pairs of orthonormal bases. The special structure of a pair of orthonormal bases allows us to prove theorems that are essentially different, and it also yields improved constants. This section describes two major results that are established in Section 5. The first theorem estimates the expected spread of the extreme singular values of a *partially random* subdictionary.

Theorem C. *Let Φ be a pair of orthonormal bases, and suppose that \mathbf{X} contains $|\Omega|$ arbitrary columns from the first basis and $|T|$ random columns from the second basis. Then*

$$\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \leq C \sqrt{\mu^2 |\Omega| \log(m+1)} + \sqrt{\frac{|T|}{d}}$$

where $m = \min\{|\Omega|, |T|\}$.

This result should be compared with Proposition 2. Consider the case where $\mu = d^{-1/2}$. Theorem C allows the subdictionary to contain $|\Omega| = \text{const} \cdot d / \log d$ arbitrary columns from the first basis and $|T| = \text{const} \cdot d$ random columns from the second basis. In sharp contrast, the proposition requires that the subdictionary contain no more than \sqrt{d} columns from at least one basis.

We would like to emphasize several curious features of Theorem C. First, the coherence and the logarithm only limit the size of the arbitrary basis; the number of randomly chosen columns can be proportional to the ambient dimension. Second, the logarithm only involves the cardinality of the smaller set of columns. Finally, note that it is possible to obtain strong probability estimates, even though this point is not immediately apparent from the statement of the theorem.

Here is another result for pairs of orthonormal bases that provides a probability estimate and is directly comparable with earlier work. In the next section, we weigh this theorem against results of Candès and Romberg [CR06].

Theorem D. *Assume $d \geq 3$. Let Φ be a pair of orthonormal bases, and let \mathbf{X} consist of $|\Omega|$ arbitrary columns from the first basis and $|T|$ random columns from the second basis. If*

$$|\Omega| + |T| \leq \frac{c\mu^{-2}}{s \log d} \quad \text{with } s \geq 1$$

then

$$\mathbb{P} \{\|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \geq 0.5\} \leq d^{-s}.$$

The constant c is no smaller than 0.004212.

2.6. Applications and Extensions. This paper also develops applications of these theorems in the field of sparse representation. Consider a subdictionary that is well conditioned and has cardinality slightly smaller than the ambient dimension. We show that a random signal formed from this subdictionary almost surely has no other representation that is equally sparse. Moreover, with overwhelming probability, the maximally sparse representation can be identified by ℓ_1 minimization. Detailed statements of these results require some background, so we refer the reader to Sections 7 and 8.

It is also important to be aware that the work in this paper is tailored to the case where the dictionary is “uniformly coherent.” That is, many of inner products between columns are close to the coherence bound. When the dictionary has a different type of correlation structure, the

theory here may give poor results. In particular, it is possible to develop much stronger estimates for localized dictionaries—those whose Gram matrix is concentrated near the diagonal. These conclusions can be reached with similar methods, but we have chosen not to present them for several reasons. The theorems are more complicated and less compelling; the proofs are substantially longer; and the constants are awful. If there is sufficient interest, we may pursue this line of inquiry in a future paper.

3. RELATED WORK

This section describes two strands of related work: one from the field of sparse approximation and the second from the field of Banach space theory. The former provides us our motivations, while the latter equips us with the technical tools to solve the problems.

3.1. Results from sparse approximation. The earliest work on the behavior of subdictionaries can be traced to the paper of Donoho and Stark [DS89] about applications of uncertainty principles in signal processing. Their results include a proof of Proposition 2 for the Fourier–Dirac dictionary. The works [DH01, EB02, GN03] broadened this investigation to cover other pairs of orthonormal bases, resulting in Proposition 1 and the general form of Proposition 2.

Early work of Fuchs [Fuc04] hinted that certain results on sparse representation could be strengthened substantially by shifting to a probabilistic viewpoint. The first explicit results in this direction were established by Candès, Romberg, and Tao. In particular, they developed a probabilistic uncertainty principle for the Fourier–Dirac dictionary that outstrips Proposition 2 [CRT06]. The following result, which is representative, appears in [CR06].

Proposition 3 (Candès–Romberg–Tao). *Assume that $d \geq 512$. Suppose that \mathbf{X} contains $|\Omega|$ arbitrary columns from the Fourier basis and $|T|$ random columns from the Dirac basis. For each $s \geq 1$,*

$$|\Omega| + |T| \leq \frac{cd}{\sqrt{(s+1)\log d}},$$

implies that

$$\mathbb{P}\{\|\mathbf{X}^*\mathbf{X} - \mathbf{I}\| \geq 0.5\} \leq Cd^{-s} \log d.$$

The constant c is no smaller than 0.2660.

This theorem states that almost all collections of columns from the Fourier–Dirac dictionary are well conditioned, provided that the total number of columns is smaller than the dimension d by a factor of $\sqrt{\log d}$. In contrast, Theorem D requires that we back away from the dimension by a factor of $\log d$. The present methods are not perspicuous enough to reproduce this result.

Candès and Romberg also state a result for general pairs of orthonormal bases [CR06].

Proposition 4 (Candès–Romberg). *Suppose that Φ is a pair of orthonormal bases whose coherence satisfies*

$$\mu \leq \frac{1}{\sqrt{2(s+1)\log d}}.$$

Assume that each column of the first basis appears in \mathbf{X} independently with probability $d^{-1}|\Omega|$ and each column of the second basis appears independently with probability $d^{-1}|T|$. If

$$|\Omega| + |T| \leq \frac{c\mu^{-2}}{((s+1)\log d)^{5/2}},$$

then

$$\mathbb{P}\{\|\mathbf{X}^*\mathbf{X} - \mathbf{I}\| \geq 0.5\} \leq Cd^{-s} \log d.$$

Theorem D improves on this proposition in several ways. Our result removes the coherence restriction; it reduces the power on the logarithm; it improves the probability decay; and it allows us to choose an arbitrary set of columns from one basis.

We believe that there are no direct precedents in the sparse approximation literature for the results on general dictionaries, Theorems A and B.

The results of Candès–Romberg–Tao actually come from a different strand of research about the properties of *random dictionaries*. Over the last two years, this area has seen impressive contributions from a variety of researchers. A partial list of important works includes [CT06, RV06, BDDW07, CDD06, DT06]. The results in this paper come from a different tradition that focuses on sparse approximation with respect to a *fixed, deterministic dictionary*. Structured dictionaries appear in many applications, which is why our results are valuable.

3.2. Results from Banach space geometry. The Banach space literature does not directly address the problem we are studying, but Rudelson has completely solved the dual problem [Rud99] by building on earlier work of Bourgain [Bou99]. One version of the result follows.

Proposition 5 (Rudelson). *Let Φ be a unit-norm tight frame. Let \mathbf{X} be a random collection of m columns from Φ . Then*

$$\mathbb{E} \left\| \frac{1}{m} \mathbf{X} \mathbf{X}^* - \frac{1}{d} \mathbf{I} \right\| \leq C \sqrt{\frac{d \log d}{m}},$$

provided that the right-hand side does not exceed one.

Since Φ is a tight frame, its normalized inertia matrix $N^{-1} \Phi \Phi^* = d^{-1} \mathbf{I}$. This proposition states that we can approximate the normalized inertia matrix by drawing $m = \text{const} \cdot d \log d$ random columns and computing a sample average. The result is qualitatively optimal, and it completes a line of inquiry initiated by Kannan et al. [KLS97]. For recent progress on a closely related problem, see [Aub07]. The present work relies strongly on Rudelson’s techniques.

It is also interesting to view our work in the light of the Restricted Invertibility Principle of Bourgain and Tzafriri. They prove that every matrix whose columns are not too small and whose spectral norm is not too large must contain a large column submatrix that is well conditioned [BT87].

Proposition 6 (Restricted Invertibility Principle). *Let Φ be a dictionary. There exists a collection T of columns with cardinality*

$$|T| \geq \frac{cN}{\|\Phi\|^2}$$

whose smallest singular value is bounded away from zero:

$$\|\Phi_T \mathbf{x}\|_2 \geq c \|\mathbf{x}\|_2 \quad \text{for all } \mathbf{x}.$$

This remarkable theorem has applications in harmonic analysis and Banach space geometry. Other authors have extended this work; see [Ver05] and its bibliography for recent progress.

A result with such weak hypotheses can only guarantee the *existence* of a well-conditioned subdictionary. Our goal is to understand when well-conditioned subdictionaries are *prevalent*, and this investigation requires additional hypotheses. To verify this point, consider a dictionary Φ that consists of two copies of the same orthonormal basis for \mathbb{C}^d . It is easy to find orthonormal subdictionaries, but a *random* subdictionary with expected cardinality m usually contains duplicated columns when $2m \geq \sqrt{d}$. One may explain this phenomenon by noting that the coherence of the dictionary is heinous ($\mu = 1$), even though the spectral norm is small ($\|\Phi\| = \sqrt{2}$).

After the present paper was complete, we learned of another paper by Bourgain and Tzafriri [BT91] that bears directly on our work. Here is a simplified version of their result.

Proposition 7 (Bourgain–Tzafriri). *Let \mathbf{A} be an $n \times n$ matrix whose norm $\|\mathbf{A}\| \leq 1$ and whose entries satisfy*

$$|a_{jk}| \leq \frac{1}{\log^2 n}.$$

From \mathbf{A} , draw a random principal submatrix \mathbf{B} with dimensions $cn \times cn$. Then

$$\mathbb{P}\{\|\mathbf{B}\| \geq 0.5\} \leq n^{-c}.$$

By applying this proposition to the hollow Gram matrix of a dictionary, it is possible to obtain estimates for the conditioning of random subdictionaries that are much stronger than the ones presented here. Nevertheless, we feel that our results remain valuable. The primary advantage of our approach is its versatility and accessibility. In fact, we have used the methods of this paper to develop a modern proof of Proposition 7 [Tro06b]. A secondary advantage is that our techniques yield explicit and somewhat reasonable constants.

4. THE TOOLBOX

Our investigation centers around expressions of the form $\|\mathbf{R}\mathbf{A}\mathbf{R}^*\|$ where \mathbf{A} is a fixed matrix and \mathbf{R} is a restriction to a randomly chosen set of coordinates. We must understand the habitat and behavior of this animal. What is its expectation? What is the probability that it achieves an unusually large value? These questions are answered with methods from the field of probability in Banach spaces [LT91]. The technologies are standard, although most of the results have not appeared in the precise form given here. Major sources for this work include [LP86, BT91, LT91, Rud99, Buc01, RV07]; more detailed citations appear throughout the paper.

4.1. Restriction maps. Before we continue, let us introduce a little more notation. Let Ω be a subset of $\llbracket N \rrbracket$. We may define a *restriction map*

$$\mathbf{R}_\Omega : \mathbb{C}^N \rightarrow \mathbb{C}^\Omega \quad \text{via the rule} \quad (\mathbf{R}_\Omega \mathbf{f})(\omega) = f(\omega) \quad \text{for } \omega \in \Omega.$$

That is, the map \mathbf{R}_Ω limits a vector in \mathbb{C}^N to the components in Ω . The adjoint is an *extension map* with the action

$$(\mathbf{R}_\Omega^* \mathbf{f})(k) = \begin{cases} f(k) & \text{when } k \in \Omega \\ 0 & \text{when } k \in \llbracket N \rrbracket \setminus \Omega. \end{cases}$$

In words, the map \mathbf{R}_Ω^* extends a vector from \mathbb{C}^Ω to \mathbb{C}^N by padding it with zeros. We often drop the subscript from the restriction if it is obvious or irrelevant. In this work, the term “restriction” *always* refers to a coordinate restriction.

4.2. Key Technical Results. Our most important theorem is adapted from work of Rudelson and Vershynin [RV07], who build essentially on earlier work of Rudelson [Rud99]. This theorem gives information about the spectral norm of a matrix that has been restricted to a random collection of columns.

Theorem 8 (Spectral Norm of a Random Compression). *Suppose that \mathbf{A} is a matrix with N columns, and let \mathbf{R} be a restriction to m coordinates, chosen at random from $\llbracket N \rrbracket$. Fix $q \geq 1$. For each $p \geq \max\{2, 2 \log(\text{rank } \mathbf{A}\mathbf{R}^*), q/2\}$, it holds that*

$$(\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|^q)^{1/q} \leq 3\sqrt{p} \|\mathbf{A}\|_{1,2} + \sqrt{\frac{m}{N}} \|\mathbf{A}\|.$$

This theorem tells us that the spectral norm of a random submatrix of \mathbf{A} carries its share of the spectral norm of the entire matrix, plus an additional component that depends on the size of the columns of \mathbf{A} . As we will see, the \sqrt{p} can be converted into information about tail probabilities. Section 9 contains a complete proof of Theorem 8.

The next major theorem is a decoupling result. It is challenging to work directly with expressions like $\|\mathbf{R}\mathbf{A}\mathbf{R}^*\|$ because the random restriction appears twice. Decoupling allows us to move to an expression that involves (conditionally) independent restrictions. This alteration permits us to apply tools such as Theorem 8.

Theorem 9 (Decoupling). *Let \mathbf{A} be a $2N \times 2N$ Hermitian matrix with a zero diagonal, and let \mathbf{R} be a restriction to m random coordinates. For each $q \geq 1$, there exists a partition of $\llbracket 2N \rrbracket$ into two blocks T_1 and T_2 with N elements each so that*

$$(\mathbb{E} \|\mathbf{R}\mathbf{A}\mathbf{R}^*\|^q)^{1/q} < 2 \max_{m_1+m_2=m} (\mathbb{E} \|\mathbf{R}_1 \mathbf{A}_{T_1 \times T_2} \mathbf{R}_2^*\|^q)^{1/q}$$

where

- the maximum occurs over integers $m_1, m_2 \in [0, N]$,
- the symbol $\mathbf{A}_{T_1 \times T_2}$ denotes the submatrix of \mathbf{A} indexed by $T_1 \times T_2$, and
- the maps \mathbf{R}_i are independent restrictions to m_i random coordinates from T_i for $i = 1, 2$.

When \mathbf{A} has odd order $(2N + 1) \times (2N + 1)$, an analogous result holds for a partition into blocks of size N and $N + 1$.

The proof of this theorem is the subject of Section 10. The argument is a variation on a classical technique [LT91, Sec. 4.4]

Finally, we need to understand the relationship between the moments of a random variable and its tail behavior. In particular, a random variable whose L_q norm is proportional to \sqrt{q} exhibits subgaussian decay. The following result is essentially Lemma 4.10 from [LT91].

Proposition 10 (Subgaussian Tail Bound). *Let Z be a nonnegative random variable whose moments satisfy*

$$(\mathbb{E}(Z^q))^{1/q} \leq \alpha\sqrt{q} + \beta \quad \text{for all } q \geq Q.$$

For all $u \geq \sqrt{Q}$, one has the tail bound

$$\mathbb{P} \left\{ Z \geq e^{1/4}(\alpha u + \beta) \right\} \leq e^{-u^2/4}.$$

Proof. Let κ be a nonnegative number. Combining Markov's inequality with the moment bound, we have

$$\mathbb{P} \left\{ Z \geq e^\kappa(\alpha u + \beta) \right\} \leq \frac{\mathbb{E} Z^q}{(e^\kappa(\alpha u + \beta))^q} \leq \left[\frac{\alpha\sqrt{q} + \beta}{e^\kappa(\alpha u + \beta)} \right]^q.$$

Choose $q = u^2$ to make the bracket equal $e^{-\kappa}$. Then select $\kappa = 0.25$. \square

5. PAIRS OF ORTHONORMAL BASES

We begin with results for subdictionaries drawn from a pair of orthonormal bases. The special structure of these dictionaries allows us to simplify the arguments substantially.

Fix the ambient dimension d . Suppose that the dictionary $\Phi = [\Phi_1 \ \Phi_2]$, where the blocks Φ_1 and Φ_2 are unitary matrices of dimension d . In this case, the number of columns $N = 2d$. The Gram matrix of the dictionary has the form

$$\Phi^* \Phi = \begin{bmatrix} \mathbf{I} & \mathbf{F} \\ \mathbf{F}^* & \mathbf{I} \end{bmatrix}$$

where $\mathbf{F} = \Phi_1^* \Phi_2$ is a unitary matrix of dimension d . It is easy to check that the spectral norm of the dictionary always satisfies $\|\Phi\| = \sqrt{2}$. The coherence μ of the dictionary can be calculated as

$$\mu = \max_{j,k} |f_{jk}| = \max_{j,k} \left| \left\langle \varphi_j^{(1)}, \varphi_k^{(2)} \right\rangle \right| \quad \text{for } j, k = 1, 2, \dots, d$$

where $\varphi_j^{(i)}$ is the j th column from basis i . Recall from (2.2) that $\mu \geq d^{-1/2}$.

5.1. Random Subdictionaries. Now, we describe our model for choosing a random subdictionary. Suppose that Ω indexes a fixed set of columns from the first basis, and let us draw a random set T of columns from the second basis. Define $m = \min\{|\Omega|, |T|\}$. These choices lead to the subdictionary

$$\mathbf{X} = [\Phi_1 \mathbf{R}_\Omega^* \quad \Phi_2 \mathbf{R}_T^*].$$

Its Gram matrix has the form

$$\mathbf{X}^* \mathbf{X} = \begin{bmatrix} \mathbf{I} & \mathbf{R}_\Omega \mathbf{F} \mathbf{R}_T^* \\ (\mathbf{R}_\Omega \mathbf{F} \mathbf{R}_T^*)^* & \mathbf{I} \end{bmatrix}.$$

Using a standard identity, we find that

$$\|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| = \|\mathbf{R}_\Omega \mathbf{F} \mathbf{R}_T^*\|. \quad (5.1)$$

We must estimate the moments of this random variable.

This calculation is an easy consequence of Theorem 8. It is clear that $\text{rank}(\mathbf{R}_\Omega \mathbf{F} \mathbf{R}_T^*) \leq m$. Therefore, given $q \geq 1$, we may select $p = \max\{2, 2 \log(m+1), q/2\}$. Invoking the theorem to draw $|T|$ random columns from $\mathbf{R}_\Omega \mathbf{F}$, we find that

$$(\mathbb{E}_T \|\mathbf{R}_\Omega \mathbf{F} \mathbf{R}_T^*\|^q)^{1/q} \leq 3 \max\left\{\sqrt{2}, \sqrt{2 \log(m+1)}, \sqrt{q/2}\right\} \|\mathbf{R}_\Omega \mathbf{F}\|_{1,2} + \sqrt{\frac{|T|}{d}} \|\mathbf{R}_\Omega \mathbf{F}\|.$$

Since the entries of \mathbf{F} have magnitude no greater than μ and $\mathbf{R}_\Omega \mathbf{F}$ has $|\Omega|$ rows,

$$\|\mathbf{R}_\Omega \mathbf{F}\|_{1,2} \leq \mu \sqrt{|\Omega|}.$$

It also holds that $\|\mathbf{R}_\Omega \mathbf{F}\| \leq \|\mathbf{F}\| = 1$. Recalling (5.1), we obtain

$$(\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\|^q)^{1/q} \leq \sqrt{18\mu^2 |\Omega|} \max\left\{1, \sqrt{\log(m+1)}, \sqrt{q/4}\right\} + \sqrt{\frac{|T|}{d}}. \quad (5.2)$$

5.2. Results. With the moments at hand, we can immediately prove several different results. The statement of Theorem C is furnished by the choice $q = 1$ in (5.2).

Next, we prove a more detailed result. Let us assume $m \geq 2$ so that $\log(m+1) > 1$. (In case $m = 1$, the theorem is vacuous.) Select $q \geq 4 \log(m+1)$ in (5.2) to obtain

$$(\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\|^q)^{1/q} \leq \sqrt{18\mu^2 |\Omega|} \sqrt{q} + \sqrt{\frac{|T|}{d}}. \quad (5.3)$$

Choose $s \geq 1$ and muster Proposition 10 with the parameter $u = \sqrt{4s \log(m+1)}$ to reach the following theorem.

Theorem 11. *Let Φ be a pair of orthonormal bases. Let \mathbf{X} consist of $|\Omega|$ arbitrary columns from the first basis and $|T|$ random columns from the second basis. Suppose that*

$$\sqrt{18\mu^2 |\Omega| \log(m+1)} \cdot s + \sqrt{\frac{|T|}{d}} \leq e^{-1/4} \delta$$

where $m = \min\{|\Omega|, |T|\}$ and $s \geq 1$. Then

$$\mathbb{P}\{\|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \geq \delta\} \leq m^{-s}.$$

Finally, let us explain how to adduce Theorem D, which is really just a simplified version of the last result. Assume that $d \geq 3$. Continuing from (5.3), recall that $\mu \geq d^{-1/2}$ to advance to

$$(\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\|^q)^{1/q} \leq \sqrt{4.5\mu^2} \left(\sqrt{|\Omega|} + \sqrt{|T|}\right) \sqrt{q} \quad \text{for } q \geq 4 \log d.$$

The square-root function is concave, whence

$$(\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\|^q)^{1/q} \leq \sqrt{9\mu^2 (|\Omega| + |T|)} \sqrt{q}.$$

Fix $s \geq 1$. Calling on Proposition 10 with parameter $u = \sqrt{4s \log d}$, we find that the condition

$$\sqrt{36\mu^2(|\Omega| + |T|) \log d \cdot s} \leq 0.5e^{-0.25}$$

ensures that the random variable exceeds 0.5 with probability less than d^{-s} . Perform some algebraic manipulations to isolate $|\Omega| + |T|$:

$$|\Omega| + |T| \leq \frac{c\mu^{-2}}{s \log d}$$

where $c \geq 0.004212$.

6. INCOHERENT DICTIONARIES

In this section, we develop results for subdictionaries of a general dictionary. To that end, suppose that Φ is a $d \times N$ matrix with unit-norm columns and coherence μ .

6.1. Moment bounds. Let \mathbf{X} be an m -column submatrix of Φ , drawn at random. In this section, we calculate for $q \geq 1$ that

$$(\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\|^q)^{1/q} \leq \sqrt{144\mu^2 m} \max \left\{ 1, \sqrt{\log(m/2 + 1)}, \sqrt{q/4} \right\} + \frac{2m}{N} \|\Phi\|^2 \quad (6.1)$$

In the sequel, we assume that N , the number of columns in the dictionary, is even; the proof in the odd case is essentially the same.

Define the hollow Gram matrix of the dictionary: $\mathbf{H} = \Phi^* \Phi - \mathbf{I}$. This matrix has a zero diagonal since the columns of the dictionary have unit norm. Let \mathbf{R} be a random restriction onto m coordinates from $[N]$. With this new notation, the matrix of interest may be viewed as a compression of the hollow Gram matrix:

$$\mathbf{X}^* \mathbf{X} - \mathbf{I} = \mathbf{R} \mathbf{H} \mathbf{R}^*.$$

To begin our calculation, we invoke the decoupling result, Theorem 9, to see that

$$(\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\|^q)^{1/q} \leq 2 \max_{m_1 + m_2 = m} \left(\mathbb{E} \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q \right)^{1/q} \quad (6.2)$$

where $\widehat{\mathbf{H}}$ is a submatrix of \mathbf{H} of dimension $N/2 \times N/2$ and the matrices \mathbf{R}_i are independent restrictions to m_i random coordinates for each $i = 1, 2$. In the sequel, the symbol \mathbb{E}_i will denote expectation with respect to \mathbf{R}_i , holding the other random restriction fixed.

We may express

$$\left(\mathbb{E} \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q \right)^{1/q} = \left(\mathbb{E}_1 \left(\mathbb{E}_2 \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q \right)^{q/q} \right)^{1/q}.$$

It is evident that $\text{rank}(\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*) \leq m/2$ since one of the numbers m_1 or m_2 is less than or equal to $m/2$. Therefore, we may choose

$$p = \max\{2, 2 \log(m/2 + 1), q/2\}. \quad (6.3)$$

Apply the random compression bound, Theorem 8, to the inner expectation to select m_2 of the $N/2$ columns from $\widehat{\mathbf{H}}$. This step results in

$$\left(\mathbb{E}_2 \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q \right)^{1/q} \leq 3\sqrt{p} \|\mathbf{R}_1 \widehat{\mathbf{H}}\|_{1,2} + \sqrt{\frac{m_2}{N/2}} \|\mathbf{R}_1 \widehat{\mathbf{H}}\|.$$

Observe that $\mathbf{R}_1 \widehat{\mathbf{H}}$ is a submatrix of \mathbf{H} with m_1 rows. Therefore, none of its columns has ℓ_2 norm greater than $\mu\sqrt{m_1}$. Combining these bounds,

$$\left(\mathbb{E} \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q \right)^{1/q} \leq \left(\mathbb{E}_1 \left[3\mu\sqrt{m_1}\sqrt{p} + \sqrt{\frac{m_2}{N/2}} \|\mathbf{R}_1 \widehat{\mathbf{H}}\| \right]^q \right)^{1/q}.$$

Apply the triangle inequality and the homogeneity of the L_q norm to reach

$$\left(\mathbb{E} \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q\right)^{1/q} \leq 3\mu\sqrt{m_1}\sqrt{p} + \sqrt{\frac{m_2}{N/2}} \left(\mathbb{E}_1 \|\mathbf{R}_1 \widehat{\mathbf{H}}\|^q\right)^{1/q}. \quad (6.4)$$

Next, we examine the remaining expectation. The spectral norm is invariant under conjugate transposition, so we apply the compression theorem to $\widehat{\mathbf{H}}^* \mathbf{R}_1^*$ to select m_1 of the $N/2$ columns. With the same choice of p as in (6.3),

$$\left(\mathbb{E}_1 \|\mathbf{R}_1 \widehat{\mathbf{H}}\|^q\right)^{1/q} \leq 3\sqrt{p} \|\widehat{\mathbf{H}}^*\|_{1,2} + \sqrt{\frac{m_1}{N/2}} \|\widehat{\mathbf{H}}^*\|.$$

We may bound the two norms above in terms of properties of the dictionary. The entries of the matrix $\widehat{\mathbf{H}}$ are inner products between distinct columns of Φ , so the coherence controls the $(1, 2)$ norm:

$$\|\widehat{\mathbf{H}}^*\|_{1,2} \leq \mu\sqrt{N/2}.$$

The spectral norm of the dictionary controls the spectral norm of \widehat{H} :

$$\|\widehat{\mathbf{H}}^*\| \leq \|\mathbf{H}^*\| = \|\Phi^* \Phi - \mathbf{I}\| = \max\{1, \|\Phi\|^2 - 1\} \leq \|\Phi\|^2.$$

Introduce the last three estimates into (6.4) to discover that

$$\left(\mathbb{E} \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q\right)^{1/q} \leq 3\mu(\sqrt{m_1} + \sqrt{m_2})\sqrt{p} + \frac{\sqrt{4m_1 m_2}}{N} \|\Phi\|^2. \quad (6.5)$$

Let us maximize this inequality over parameters $m_1 + m_2 = m$, subject to $0 \leq m_i \leq N/2$. First, notice that $\sqrt{4m_1 m_2} \leq m$. The square-root function is concave, so

$$\sqrt{m_1} + \sqrt{m_2} \leq 2\sqrt{(m_1 + m_2)/2} = \sqrt{2m}.$$

Introducing the last two bounds into (6.5), we reach

$$\left(\mathbb{E} \|\mathbf{R}_1 \widehat{\mathbf{H}} \mathbf{R}_2^*\|^q\right)^{1/q} \leq \sqrt{18\mu^2 m} \sqrt{p} + \frac{m}{N} \|\Phi\|^2.$$

Finally, recall the value of p from (6.3) and substitute the most recent bound into (6.2) to reach the announced inequality (6.1).

6.2. Results. First, we show how to reach Theorem A of Section 2. This result frames the hypotheses that $N \geq 2d$ and that Φ is a unit-norm tight frame. For simplicity, assume that $m \geq 6$ so that $\log(m/2 + 1) > 1$. The moment bound (6.1) with $q = 1$ implies that

$$\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \leq \sqrt{144\mu^2 m \log(m/2 + 1)} + \frac{2m}{N} \|\Phi\|^2.$$

Since Φ is a tight frame, $\|\Phi\|^2 = N/d$. Thus

$$\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \leq \sqrt{144\mu^2 m \log(m/2 + 1)} + \frac{2m}{d}.$$

The assumption that $N \geq 2d$ implies that the coherence $\mu > (2d)^{-1/2}$ on account of inequality (1.1). When the first term is less than one, it follows that

$$\frac{2m}{d} < 4\mu^2 m \ll 4\sqrt{\mu^2 m}.$$

Therefore, we can absorb the first term into the second term by adjusting the constant. Increasing $\log(m/2 + 1)$ to $\log(m + 1)$ completes the argument. When $m < 6$, the same result clearly holds with a (potentially) larger constant.

Our major result contains more detailed information about the role of the dictionary's spectral norm and the probability decay. In case $q \geq 4 \log(m/2 + 1) \geq 4$, the moment bound (6.1) becomes

$$(\mathbb{E} \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\|^q)^{1/q} \leq \sqrt{36\mu^2 m} \sqrt{q} + \frac{2m}{N} \|\Phi\|^2.$$

Fix $s \geq 1$, and invoke Proposition 10 with parameter $u = \sqrt{4s \log(m/2 + 1)}$.

Theorem 12. *Let Φ be a dictionary, and let \mathbf{X} be a random m -column subdictionary where $m \geq 4$. Suppose that*

$$\sqrt{144\mu^2 m \log(m/2 + 1)} \cdot s + \frac{2m}{N} \|\Phi\|^2 \leq e^{-1/4} \delta \quad \text{with } s \geq 1.$$

Then

$$\mathbb{P} \{ \|\mathbf{X}^* \mathbf{X} - \mathbf{I}\| \geq \delta \} \leq (m/2)^{-s}.$$

Theorem B of Section 2 follows from a similar argument, along with some algebraic simplifications.

7. APPLICATION: SPARSE REPRESENTATION

Our previous results have several interesting applications. The first one concerns uniqueness of sparse representations, the problem that was raised in the introduction. We draw a random sparse signal according to the following model and ask about the probability that this signal has another representation that is equally sparse.

Model (M0) for a random signal $\mathbf{s} = \Phi \mathbf{z}$	
The dictionary:	Φ has coherence μ .
The subdictionary:	Φ_T has least singular value $\sigma_{\min}(\Phi_T) \geq 2^{-1/2}$, and its cardinality satisfies $\mu^2 T < 2^{-1/2}$.
The coefficients:	\mathbf{z} is supported on T , and $\mathbf{R}_T \mathbf{z}$ is continuously distributed.

Observe that our previous results (such as Theorems 11 and 12) allow us to determine the probability that a randomly chosen subdictionary fits the requirements of this model.

Theorem 13. *Suppose that $\mathbf{s} = \Phi \mathbf{z}$ is a random signal drawn from Model (M0). Then \mathbf{z} is almost surely the unique vector that satisfies the constraints*

$$\Phi \mathbf{c} = \mathbf{s} \quad \text{and} \quad \|\mathbf{c}\|_0 \leq m.$$

This theorem is closely related to results of Candès and Romberg about the uniqueness of signals that are sparse with respect to a pair of orthonormal bases [CR06, Thms. 4.1, 5.2]. The extension to other types of dictionaries is new.

Proof. We make the abbreviation $m = |T|$. Since Φ_T has full rank, the signal $\mathbf{s} = \Phi \mathbf{z}$ has a continuous distribution on the range of Φ_T . Let Ω be an arbitrary set of m indices or fewer. We claim that

$$\text{range}(\Phi_\Omega) = \text{range}(\Phi_T) \quad \implies \quad \Omega = T.$$

In particular,

$$\Omega \neq T \quad \implies \quad \dim(\text{range}(\Phi_\Omega) \cap \text{range}(\Phi_T)) < m.$$

To this how this fact implies the result, consider the set of signals in $\text{range}(\Phi_T)$ that can be represented using a different set of m columns. This set can be written as a finite union of subspaces with dimension strictly less than m . Therefore, it has zero volume with respect to any nonatomic

measure. In other words, it is almost surely the case that the signal has no other representation as a linear combination of m columns from Φ . The theorem follows.

Let us turn to the claim. We may assume that $|\Omega| = m$, or else the range of Φ_Ω cannot exhaust the range of Φ_T . Since the dimensions match, it suffices to find one vector in $\text{range}(\Phi_\Omega)$ that does not lie in $\text{range}(\Phi_T)$. We check that, for any $\omega \in \Omega \setminus T$,

$$\|P_T \varphi_\omega\|_2 < 1,$$

where P_T is the orthogonal projector onto $\text{range}(\Phi_T)$. Since Ω is distinct from T , it follows that $\Omega \setminus T$ is nonempty, which supplies us with the coveted vector.

Observe that the projector can be written as

$$P_T = (\Phi_T^\dagger)^* \Phi_T^*$$

where the dagger \dagger indicates the Moore–Penrose pseudoinverse. The usual norm estimate implies that

$$\|P_T \varphi_\omega\|_2 \leq \|\Phi_T^\dagger\| \|\Phi_T^* \varphi_\omega\|_2.$$

Since $\sigma_{\min}(\Phi_T) \geq 2^{-1/2}$, the first term is no greater than $\sqrt{2}$. Since $\omega \notin T$, each entry of the vector $\Phi_T^* \varphi_\omega$ is bounded in magnitude by the coherence μ . Therefore, $\|\Phi_T^* \varphi_\omega\|_2 \leq \mu\sqrt{m}$. We conclude that

$$\|P_T \varphi_\omega\|_2 \leq \sqrt{2} \cdot \mu\sqrt{m} < 1$$

owing to the assumptions in Model (M0). □

8. APPLICATION: SPARSE RECOVERY

In this section, we consider an algorithmic method for identifying representations of a sparse signal. Suppose that Φ is a dictionary. Let \mathbf{s} be a signal that has a sparse representation with respect to the dictionary:

$$\mathbf{s} = \Phi \mathbf{c}_{\text{opt}} \quad \text{where } \|\mathbf{c}_{\text{opt}}\|_0 \leq m.$$

In general, it is NP-hard to identify the coefficient vector \mathbf{c}_{opt} [Nat95, DMA97], so various heuristic methods have been developed. One approach, especially popular with mathematicians, is to solve

$$\min \|\mathbf{c}\|_1 \quad \text{subject to } \Phi \mathbf{c} = \mathbf{s}. \tag{P1}$$

The idea is that the ℓ_1 norm is a convex relaxation of the ℓ_0 quasi-norm, so one hopes that the solutions will coincide. The problem (P1) can be cast as a second-order cone program, which means that the optimization can be completed in polynomial time. The paper [CDS99] was the first to advocate this heuristic for sparse recovery. Some other references include [DH01, Tro06a, CR06].

Before we present the signal model, a little background is necessary. Define the *signum function*

$$\text{sgn}(re^{i\theta}) \stackrel{\text{def}}{=} \begin{cases} e^{i\theta} & \text{when } r > 0 \\ 0 & \text{when } r = 0. \end{cases}$$

Extend this function to vectors by applying it to each component. A *Steinhaus random variable* is a complex random variable that is distributed uniformly on the unit circle $\{w \in \mathbb{C} : |w| = 1\}$. A *Steinhaus sequence* is a (countable) collection of independent Steinhaus random variables.

We draw a random sparse signal according to the following model and ask about the probability that this representation can be identified by solving (P1).

Model (M1) for a random signal $\mathbf{s} = \Phi \mathbf{z}$		
The dictionary:		has coherence μ .
The subdictionary:	Φ_T	has least singular value $\sigma_{\min}(\Phi_T) \geq 2^{-1/2}$, and its cardinality satisfies $8\mu^2 T \leq \log(N/\delta)$.
The coefficients:	\mathbf{z}	is supported on T , and $\text{sgn}(\mathbf{R}_T \mathbf{z})$ forms a Steinhaus sequence.

In words, the magnitudes of the nonzero entries of the coefficient vector are completely arbitrary, but the phases must be independent and uniformly distributed on the circle. If in addition the distribution of the nonzero coefficients is continuous, then the signal also satisfies the requirements of Model (M0).

Theorem 14. *Suppose that $\mathbf{s} = \Phi \mathbf{z}$ is a random signal drawn according to Model (M1). Then \mathbf{z} is the unique solution to (P1), except with probability 2δ .*

If the signal satisfies the requirements of both Model (M0) and (M1), then Theorem 13 shows that the sparsest representation of the random signal is almost surely unique and Theorem 14 shows that (P1) identifies this maximally sparse representation with overwhelming probability.

The proof of Theorem 14 is an application of a result due to Fuchs [Fuc04] and the present author [Tro05] combined with a basic large deviation inequality. Candès and Romberg have established an analogous result for pairs of orthonormal bases [CR06]. The extension to other types of dictionaries is novel.

Proposition 15 (Fuchs, Tropp). *Suppose that $\mathbf{s} = \Phi \mathbf{c}_{\text{opt}}$, and write $T = \text{supp}(\mathbf{c}_{\text{opt}})$. If*

$$\left| \left\langle \Phi_T^\dagger \varphi_k, \mathbf{R}_T \text{sgn}(\mathbf{c}_{\text{opt}}) \right\rangle \right| < 1 \quad \text{for all } k \notin T$$

then \mathbf{c}_{opt} is the unique solution to (P1).

As before, the dagger \dagger denotes the Moore–Penrose pseudoinverse.

Proposition 16 (Complex Bernstein). *Let \mathbf{a} be a complex vector, and let ε be a Steinhaus sequence. For all $u \geq 0$ and all $\kappa \in (0, 1)$,*

$$\mathbb{P} \left\{ \left| \sum_j \varepsilon_j a_j \right| \geq u \|\mathbf{a}\|_2 \right\} \leq \frac{e^{-\kappa u^2}}{1 - \kappa}.$$

In particular, we may take the right-hand side to be $2e^{-u^2/2}$.

Proof sketch. Let $Z = \left| \sum_j \varepsilon_j a_j \right|$. The Khintchine inequality for Steinhaus sequences yields a sharp estimate for the even moments of this random variable [PS95, Sec. 9.3]:

$$\mathbb{E} Z^{2p} \leq p! \|\mathbf{a}\|_2^{2p}.$$

Apply the Laplace transform method to Z^2 to obtain the tail bound. □

Proof of Theorem 14. We abbreviate $m = |T|$. To ensure that the unique solution to (P1) equals \mathbf{z} , Proposition 15 asks us to verify that

$$\left| \left\langle \Phi_T^\dagger \varphi_k, \mathbf{R}_T(\text{sgn } \mathbf{z}) \right\rangle \right| < 1 \quad \text{for all } k \notin T.$$

First, let us estimate the norm of the left-hand member of the inner product. Expanding the pseudoinverse, we obtain

$$\Phi_T^\dagger \varphi_k = (\Phi_T^* \Phi_T)^{-1} \Phi_T^* \varphi_k.$$

Taking norms and applying the familiar operator norm bound,

$$\|\Phi_T^\dagger \varphi_k\|_2 \leq \|(\Phi_T^* \Phi_T)^{-1}\| \|\Phi_T^* \varphi_k\|_2.$$

Since $\sigma_{\min}(\Phi_T) \geq 2^{-1/2}$, the spectral norm of the inverse matrix does not exceed two. Since $k \notin T$, each entry of the vector $\Phi_T^* \varphi_k$ is bounded in magnitude by the coherence parameter μ , so $\|\Phi_T^* \varphi_k\|_2 \leq \mu\sqrt{m}$. Therefore,

$$\|\Phi_T^\dagger \varphi_k\|_2 \leq 2\mu\sqrt{m} \quad \text{for all } k \notin T.$$

The inner product can be rewritten as

$$\langle \Phi_T^\dagger \varphi_k, \mathbf{R}_T(\text{sgn } z) \rangle = \sum_{j \in T} (\Phi_T^\dagger \varphi_k)_j \overline{\text{sgn } z_j}.$$

Since $\{\overline{\text{sgn } z_j} : j \in T\}$ is a Steinhaus sequence, the complex Bernstein inequality results in

$$\mathbb{P} \left\{ \left| \langle \Phi_T^\dagger \varphi_k, \mathbf{R}_T(\text{sgn } z) \rangle \right| \geq 1 \right\} \leq 2e^{-1/8\mu^2 m}.$$

Invoking the union bound over at most N choices of φ_k , we reach

$$\mathbb{P} \left\{ \max_{k \notin T} \left| \langle \Phi_T^\dagger \varphi_k, \mathbf{R}_T(\text{sgn } z) \rangle \right| \geq 1 \right\} \leq 2Ne^{-1/8\mu^2 m}.$$

Owing to the choice of m in Model (M1), the right-hand side does not exceed 2δ . \square

9. THE SPECTRAL NORM OF A RANDOM COMPRESSION

The goal of this section is to study the moments of the random variable $\|\mathbf{A}\mathbf{R}_\Omega^*\|$ where the matrix \mathbf{A} is arbitrary and \mathbf{R}_Ω is a restriction to m random coordinates. If we treat the columns of \mathbf{A} as separate points in a linear space of matrices, the restriction map effectively creates the sum of a random subset of these columns:

$$\mathbf{A}\mathbf{R}_\Omega^* = \sum_{k \in \Omega} \mathbf{a}_k \mathbf{e}_k^*$$

where \mathbf{e}_k is the k th standard basis vector. To estimate the spectral norm of this sum, we first introduce additional randomness of a type that is easier to understand. Then we can study the sum by conditioning on the subset Ω and applying methods that are adapted for the simple random variables.

9.1. Rademacher Sums. The simplest nontrivial random variable is the *Rademacher random variable*, which takes the values ± 1 with equal probability. It is traditionally denoted by the letter ε . A *Rademacher sequence* is a (countable) collection of independent Rademacher random variables.

Let $\{x_1, x_2, \dots, x_N\}$ be a sequence of points in a normed linear space X . The associated *Rademacher sum* is the random linear combination

$$\sum_{k=1}^N \varepsilon_k x_k$$

where $\{\varepsilon_k\}$ is a Rademacher sequence. As we will see, Rademacher sums play a central role in our investigation because there are powerful techniques for estimating their norms.

9.2. Symmetrization of Random Subset Sums. A fundamental method for studying the sum of independent random variables in a normed space is to compare the sum against a related Rademacher sum. This process is known as *symmetrization*, and the key result is Lemma 6.3 of [LT91]. This lemma does not apply to random subset sums, but it is possible to establish an analogous bound. In words, a random subset sum carries its share of the norm plus an additional component that is dominated by a Rademacher sum.

Lemma 17. *Let $\{x_1, x_2, \dots, x_N\}$ be a sequence of points in a normed linear space X , and let Ω be a random subset of $\llbracket N \rrbracket$ with cardinality m . For each $q \geq 1$,*

$$\left(\mathbb{E}_\Omega \left\| \sum_{\Omega} x_k \right\|_X^q \right)^{1/q} \leq 2 \left(\mathbb{E}_\Omega \mathbb{E}_\varepsilon \left\| \sum_{\Omega} \varepsilon_k x_k \right\|_X^q \right)^{1/q} + \frac{m}{N} \left\| \sum_{k=1}^N x_k \right\|_X.$$

where ε is a Rademacher sequence independent from Ω .

This result may be known, but we do not have a reference. (See the proof of Theorem 3.4 in [RV06], which mistakenly cites Lemma 6.3 of [LT91].) Before we begin, let us make another definition. A *balanced Rademacher sequence* of length $2K$ is a random vector $\xi \in \{\pm 1\}^{2K}$ whose entries sum to zero, with all such choices equally likely.

Proof. We prove the result for $q = 1$, since the other cases are substantially identical. Let $\delta \in \{0, 1\}^N$ be a random vector with exactly m entries equal to one. The variables $\{\delta_k\}$ will select points to participate in the sum. That is,

$$E \stackrel{\text{def}}{=} \mathbb{E}_\Omega \left\| \sum_{\Omega} x_k \right\|_X = \mathbb{E}_\delta \left\| \sum_{k=1}^N \delta_k x_k \right\|_X.$$

Define $\bar{\delta} = m/N$ to be the (common) expectation of the entries of δ . Center the selectors by subtracting their mean and applying the triangle inequality:

$$E \leq \mathbb{E}_\delta \left\| \sum_{k=1}^N (\delta_k - \bar{\delta}) x_k \right\|_X + \bar{\delta} \left\| \sum_{k=1}^N x_k \right\|_X.$$

(When $q > 1$, we must apply the triangle inequality a second time to draw the second summand out of the L_q norm.) Let the random vector δ' be an independent copy of δ . Then,

$$\begin{aligned} E &\leq \mathbb{E}_\delta \left\| \mathbb{E}_{\delta'} \sum_{k=1}^N (\delta_k - \delta'_k) x_k \right\|_X + \bar{\delta} \left\| \sum_{k=1}^N x_k \right\|_X \\ &\leq \mathbb{E}_{\delta, \delta'} \left\| \sum_{k=1}^N (\delta_k - \delta'_k) x_k \right\|_X + \bar{\delta} \left\| \sum_{k=1}^N x_k \right\|_X, \end{aligned}$$

where the second relation follows from Jensen's inequality.

The rest of the proof focuses on the quantity

$$F \stackrel{\text{def}}{=} \mathbb{E}_{\delta, \delta'} \left\| \sum_{k=1}^N (\delta_k - \delta'_k) x_k \right\|_X.$$

Consider the set of places where exactly one of the two selectors is active:

$$T \stackrel{\text{def}}{=} \{k \in \llbracket N \rrbracket : \delta_k + \delta'_k = 1\}.$$

Note the following properties of this set:

- (1) The cardinality of T is an even number in the range $[0, 2m]$.
- (2) Conditional on $|T|$, the set T is a uniformly random subset of $\llbracket N \rrbracket$.
- (3) Exactly $\frac{1}{2}|T|$ of the components k listed in T have $\delta_k = 1$. These locations form a uniformly random subset of T .

If $\xi \in \{\pm 1\}^T$ is a balanced Rademacher sequence of length $|T|$, independent from everything, then our expression for F becomes

$$F = \mathbb{E}_T \mathbb{E}_\xi \left\| \sum_T \xi_k x_k \right\|_X.$$

We estimate the balanced Rademacher sum using Lemma 18 of the sequel, which furnishes

$$F \leq 2 \mathbb{E}_T \mathbb{E}_U \mathbb{E}_\varepsilon \left\| \sum_U \varepsilon_k x_k \right\|_X$$

where U is a uniformly random subset of T with $\frac{1}{2}|T|$ entries and ε is a standard Rademacher sequence, independent from everything. The distribution induced on U has two notable properties.

- (1) The cardinality $|U|$ is an integer random variable ranging over $[0, m]$.
- (2) Conditional on $|U|$, the set U is a uniformly random subset of $\llbracket N \rrbracket$.

Forgetting about the subset T , we may write

$$\begin{aligned} F &\leq 2 \mathbb{E}_{|U|} \mathbb{E}_U \mathbb{E}_\varepsilon \left\| \sum_U \varepsilon_k x_k \right\|_X \\ &\leq 2 \max_{|U|=0,1,\dots,m} \mathbb{E}_U \mathbb{E}_\varepsilon \left\| \sum_U \varepsilon_k x_k \right\|_X. \end{aligned}$$

Using Jensen's inequality, one easily checks that Rademacher series are monotonic in the sense that

$$\mathbb{E}_\varepsilon \left\| \sum_{k=1}^K \varepsilon_k y_k \right\|_X \leq \mathbb{E}_\varepsilon \left\| \sum_{k=1}^{K+1} \varepsilon_k y_k \right\|_X \quad \text{for each sequence } \{y_k\} \subset X.$$

It follows that the maximum is attained when $|U| = m$. We conclude that

$$F \leq 2 \mathbb{E}_\Omega \mathbb{E}_\varepsilon \left\| \sum_\Omega \varepsilon_k x_k \right\|_X,$$

where Ω is a uniformly random subset of $\llbracket N \rrbracket$ with cardinality m . \square

The next lemma shows that a balanced Rademacher sum is dominated by a certain standard Rademacher sum.

Lemma 18. *Let $\{x_1, x_2, \dots, x_{2K}\}$ be a sequence of points in a normed linear space X , and suppose that ξ is a balanced Rademacher sequence of length $2K$. For each $q \geq 1$,*

$$\left(\mathbb{E}_\xi \left\| \sum_{k=1}^{2K} \xi_k x_k \right\|_X^q \right)^{1/q} \leq 2 \left(\mathbb{E}_U \mathbb{E}_\varepsilon \left\| \sum_U \varepsilon_k x_k \right\|_X^q \right)^{1/q}$$

where U is a random set of K numbers from $\llbracket 2K \rrbracket$ and ε is a standard Rademacher sequence, independent from U .

Proof. We establish the result for $q = 1$ since the proof is the same for the other cases. The basic strategy is to pair positive and negative entries of the balanced Rademacher sequence at random and then to turn off one element of each pair, also at random.

Let π be a random permutation of $\llbracket 2K \rrbracket$. For each π , define $\pi_+(k) = \pi(k)$ and $\pi_-(k) = \pi(K+k)$ for each $k = 1, 2, \dots, K$. Then

$$E \stackrel{\text{def}}{=} \mathbb{E}_\xi \left\| \sum_{k=1}^{2K} \xi_k x_k \right\|_X = \mathbb{E}_\pi \left\| \sum_{k=1}^K [x_{\pi_+(k)} - x_{\pi_-(k)}] \right\|_X.$$

Independent from π , draw a random vector δ from $\{0, 1\}^K$, and observe that $\mathbb{E} \delta_k = \frac{1}{2}$ for each component k . Using Jensen's inequality and independence,

$$\begin{aligned} E &= 2 \mathbb{E}_\pi \left\| \mathbb{E}_\delta \sum_{k=1}^K [\delta_k x_{\pi_+(k)} - (1 - \delta_k) x_{\pi_-(k)}] \right\|_X \\ &\leq 2 \mathbb{E}_\delta \mathbb{E}_\pi \left\| \sum_{k=1}^K [\delta_k x_{\pi_+(k)} - (1 - \delta_k) x_{\pi_-(k)}] \right\|_X \\ &= 2 \mathbb{E}_\varepsilon \mathbb{E}_\pi \left\| \sum_{k=1}^K \varepsilon_k x_{\pi_{\varepsilon_k}(k)} \right\|_X \end{aligned}$$

where ε is a standard Rademacher sequence, independent from everything. For each fixed ε ,

$$\{\pi_{\varepsilon_1}(1), \pi_{\varepsilon_2}(2), \dots, \pi_{\varepsilon_K}(K)\}$$

is a uniformly random sequence of K distinct numbers from $\llbracket 2K \rrbracket$ independent from ε , so we can write

$$E \leq 2 \mathbb{E}_\varepsilon \mathbb{E}_\pi \left\| \sum_{k=1}^K \varepsilon_k x_{\pi(k)} \right\|_X.$$

Interchange the order of the expectations and invert the permutation:

$$E \leq 2 \mathbb{E}_\pi \mathbb{E}_\epsilon \left\| \sum_U \epsilon_{\pi^{-1}(k)} x_k \right\|_X$$

where $U = \pi^{-1}([K])$ is a uniformly random set of K numbers from $[2K]$. Rademacher sequences are exchangeable, so we may remove the inverse permutation to complete the proof. \square

9.3. Schatten Norms. Next, we present another piece of background. To each matrix \mathbf{A} , one may associate the vector $\boldsymbol{\sigma}(\mathbf{A})$ of singular values. Given a parameter $1 \leq p \leq \infty$, the Schatten p -norm is defined as

$$\|\mathbf{A}\|_{S_p} \stackrel{\text{def}}{=} \|\boldsymbol{\sigma}(\mathbf{A})\|_p,$$

where $\|\cdot\|_p$ is the usual ℓ_p vector norm. Note that the choices $p = 2, \infty$ lead respectively to the Frobenius and spectral norms. The Schatten norms inherit many properties from the ℓ_p norms. In particular,

- When $q \leq p$, it holds that $\|\mathbf{A}\|_{S_p} \leq \|\mathbf{A}\|_{S_q}$.
- When $1 \leq p$, it holds that $\|\mathbf{A}\|_{S_p} \leq \text{rank}(\mathbf{A})^{1/p} \|\mathbf{A}\|$.

Moreover, the Schatten spaces form a scale, so we can interpolate between them. The books [TJ89, Sim05] contain a detailed study of the Schatten classes.

9.4. Khintchine Inequalities. The scalar Khintchine inequality, which dates to 1923, gives detailed information about the moments of a Rademacher sum of real numbers. More recently, it was discovered that an analogous inequality holds for Rademacher sums of matrices. This astonishing fact is due to Lust-Picquard [LP86], but the version here depends on an argument of Buchholz [Buc01] that extends the original method of Khintchine [PS95, Sec. 2].

Proposition 19 (Noncommutative Khintchine Inequality). *Let $\{\mathbf{A}_k\}$ be a finite sequence of matrices of the same dimension, and let γ be a sequence of independent, standard Gaussian variables. For each even number $p \geq 2$,*

$$\left[\mathbb{E} \left\| \sum_k \gamma_k \mathbf{A}_k \right\|_{S_p}^p \right]^{1/p} \leq C_p \max \left\{ \left\| \left(\sum_k \mathbf{A}_k \mathbf{A}_k^* \right)^{1/2} \right\|_{S_p}, \left\| \left(\sum_k \mathbf{A}_k^* \mathbf{A}_k \right)^{1/2} \right\|_{S_p} \right\}, \quad (9.1)$$

where the optimal constant

$$C_p = \left(\frac{p!}{2^{p/2} (p/2)!} \right)^{1/p}.$$

This result is a direct application of Theorem 5 from [Buc01] with the expectation playing the role of the linear functional ψ . It is elementary to check that Gaussian variables have “mixed moments defined by pairings” with the pairing function m_n identically equal to one. Note that the optimal value of the constant matches the scalar case [Haa82].

The following corollary extends Proposition 19 to other values of p , although it is unable to locate the best constant.

Corollary 20. *When $2 \leq p < \infty$, the Khintchine inequality (9.1) holds with a constant that satisfies*

$$C_p \leq 2^{0.25} e^{-1/2} \sqrt{p}.$$

Proof sketch. Choose $n \in \mathbb{N}$ and $\theta \in (0, 1)$. To obtain the Khintchine inequality for $S_{2n+2\theta}$, invoke the Riesz–Thorin theorem [Zyg02, Chap. XII] to interpolate between the inequalities for S_{2n} and S_{2n+2} . The constant satisfies

$$C_{2n+2\theta} \leq C_{2n}^{1-\theta} C_{2n+2}^\theta.$$

A routine estimate (using, for example, Stirling’s approximation) shows that

$$C_{2n} \leq 2^{1/4n} e^{-1/2} \sqrt{2n}.$$

Introducing the latter bound into the former yields the conclusion. \square

It is an easy consequence of Jensen's inequality that the noncommutative Khintchine inequality also holds for Rademacher sums of matrices.

Corollary 21. *Let $\{\mathbf{A}_k\}$ be a finite sequence of matrices of the same dimension, and let ε be a Rademacher sequence. For each $p \geq 2$,*

$$\left[\mathbb{E} \left\| \sum_k \varepsilon_k \mathbf{A}_k \right\|_{S_p}^p \right]^{1/p} \leq D_p \max \left\{ \left\| \left(\sum_k \mathbf{A}_k \mathbf{A}_k^* \right)^{1/2} \right\|_{S_p}, \left\| \left(\sum_k \mathbf{A}_k^* \mathbf{A}_k \right)^{1/2} \right\|_{S_p} \right\}, \quad (9.2)$$

where the constant

$$D_p \leq \sqrt{\frac{\pi}{2}} C_p \leq 2^{-0.25} \sqrt{\frac{\pi}{e}} \sqrt{p}.$$

The optimal value of the Rademacher constant D_p is not available except in the simplest case $p = 2$. The Central Limit Theorem implies that $D_p \geq C_p$. It seems possible that the Rademacher and Gaussian constants are identical, just as they are in the scalar Khintchine inequality, but the question remains open.

9.5. Rudelson's Lemma. The noncommutative Khintchine inequality does not apply directly to the spectral norm, but it can be used to obtain information about the spectral norm of a Rademacher sum. The following estimate, due to Rudelson [Rud99], shows how to accomplish this.

Lemma 22 (Rudelson). *Suppose that $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m$ are the columns of a matrix \mathbf{A} , and fix $q > 0$. For any $p \geq \max\{2, 2 \log(\text{rank } \mathbf{A}), q\}$, it holds that*

$$\left(\mathbb{E} \left\| \sum_{k=1}^m \varepsilon_k \mathbf{a}_k \mathbf{a}_k^* \right\|^q \right)^{1/q} < 1.5 \sqrt{p} \|\mathbf{A}\|_{1,2} \|\mathbf{A}\|,$$

where ε is a Rademacher sequence.

Proof. Choose a value of p subject to the limitations described above. To begin the estimate, we bound the spectral norm by the Schatten p -norm and use Hölder's inequality to move from the L_q norm to the L_p norm.

$$\begin{aligned} E &\stackrel{\text{def}}{=} \left(\mathbb{E} \left\| \sum_{k=1}^m \varepsilon_k \mathbf{a}_k \mathbf{a}_k^* \right\|^q \right)^{1/q} \\ &\leq \left(\mathbb{E} \left\| \sum_{k=1}^m \varepsilon_k \mathbf{a}_k \mathbf{a}_k^* \right\|_{S_p}^q \right)^{1/q} \\ &\leq \left(\mathbb{E} \left\| \sum_{k=1}^m \varepsilon_k \mathbf{a}_k \mathbf{a}_k^* \right\|_{S_p}^p \right)^{1/p}. \end{aligned}$$

Apply the noncommutative Khintchine inequality to obtain

$$E \leq D_p \left\| \left(\sum_{k=1}^m \|\mathbf{a}_k\|_2^2 \mathbf{a}_k \mathbf{a}_k^* \right)^{1/2} \right\|_{S_p}.$$

The rank of matrix inside the norm does not exceed $r = \text{rank } \mathbf{A}$. Thus, we may bound the Schatten p -norm by the spectral norm if we pay a factor of $r^{1/p}$, which does not exceed \sqrt{e} by our choice of p . Afterward, we draw the square root out from the norm to reach

$$E \leq D_p \sqrt{e} \left\| \sum_{k=1}^m \|\mathbf{a}_k\|_2^2 \mathbf{a}_k \mathbf{a}_k^* \right\|^{1/2}.$$

The summands are all positive semidefinite, so the spectral norm of the sum increases monotonically with each scalar coefficient. Therefore, we may replace each coefficient by $\max_k \|\mathbf{a}_k\|_2^2$ and use the

homogeneity of the norm to obtain

$$E \leq D_p \sqrt{e} \cdot \max_k \|\mathbf{a}_k\|_2 \cdot \left\| \sum_{k=1}^m \mathbf{a}_k \mathbf{a}_k^* \right\|^{1/2}.$$

The maximum can be rewritten as $\|\mathbf{A}\|_{1,2}$, and the spectral norm can be expressed as

$$\left\| \sum_{k=1}^m \mathbf{a}_k \mathbf{a}_k^* \right\|^{1/2} = \|\mathbf{A}\mathbf{A}^*\|^{1/2} = \|\mathbf{A}\|.$$

Recall that $D_p \leq 2^{-0.25} \sqrt{\pi/e} \sqrt{p}$, then calculate the leading constant numerically to complete the proof. \square

9.6. Random Compressions. We are now prepared to prove the main theorem. This argument is essentially due to Rudelson [Rud99]. Some related work appears in [MP06, RV07]. The only significant change in the proof is the application of the new symmetrization result, Lemma 17, in place of Lemma 6.3 of [LT91]. We have also made some minor modifications to reduce the size of the constants.

Theorem 23 (Spectral Norm of a Random Compression). *Suppose that \mathbf{A} is a matrix with N columns, and let \mathbf{R} be a random restriction to m coordinates, chosen at random from $\llbracket N \rrbracket$. Fix $q \geq 2$. For each $p \geq \max\{2, 2 \log(\text{rank } \mathbf{A}\mathbf{R}^*), q/2\}$, it holds that*

$$(\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|^q)^{1/q} \leq 3\sqrt{p} \left(\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|_{1,2}^q \right)^{1/q} + \sqrt{\frac{m}{N}} \|\mathbf{A}\|.$$

Proof. Let us begin with an overview of the proof. First, we express the random compression as a subset sum. Then, we symmetrize the subset sum and apply Rudelson's lemma to obtain an upper bound involving the value we are trying to estimate. Finally, we solve an algebraic relation to obtain an explicit estimate.

The object of interest is the quantity

$$E \stackrel{\text{def}}{=} (\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|^q)^{1/q}.$$

First, observe that

$$E^2 = \left(\mathbb{E} \|\mathbf{A}\mathbf{R}^* \mathbf{R} \mathbf{A}^*\|^{q/2} \right)^{2/q} = \left(\mathbb{E} \left\| \sum_{\Omega} \mathbf{a}_k \mathbf{a}_k^* \right\|^{q/2} \right)^{2/q}$$

where Ω is the random set of coordinates selected by the restriction \mathbf{R} . Since $q/2 \geq 1$, the symmetrization result, Lemma 17, shows that

$$\begin{aligned} E^2 &\leq 2 \left(\mathbb{E}_{\Omega} \mathbb{E}_{\varepsilon} \left\| \sum_{\Omega} \varepsilon_k \mathbf{a}_k \mathbf{a}_k^* \right\|^{q/2} \right)^{2/q} + \frac{m}{N} \left\| \sum_{k=1}^N \mathbf{a}_k \mathbf{a}_k^* \right\| \\ &= 2 \left[\mathbb{E}_{\Omega} \left(\mathbb{E}_{\varepsilon} \left\| \sum_{\Omega} \varepsilon_k \mathbf{a}_k \mathbf{a}_k^* \right\|^{q/2} \right)^{(2/q)(q/2)} \right]^{2/q} + \frac{m}{N} \|\mathbf{A}\|^2. \end{aligned}$$

Choose p subject to the limitations described. To estimate the parenthesis, we invoke Rudelson's lemma, conditional on the choice of Ω . The matrix in the statement of the lemma is $\mathbf{A}\mathbf{R}^*$, resulting in

$$E^2 \leq 3\sqrt{p} \left[\mathbb{E} \left(\|\mathbf{A}\mathbf{R}^*\|_{1,2} \|\mathbf{A}\mathbf{R}^*\| \right)^{q/2} \right]^{2/q} + \frac{m}{N} \|\mathbf{A}\|^2.$$

Invoke the Cauchy–Schwarz inequality to find that

$$E^2 \leq 3\sqrt{p} \left(\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|_{1,2}^q \right)^{1/q} (\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|^q)^{1/q} + \frac{m}{N} \|\mathbf{A}\|^2.$$

Observe that we have obtained a copy of E on the right-hand side of the relation.

This inequality takes the form $E^2 \leq bE + c$. One obtains an upper bound on E by choosing the larger root of the quadratic and invoking the subadditivity of the square root:

$$E \leq \frac{b + \sqrt{b^2 + 4c}}{2} \leq b + \sqrt{c}.$$

We discover that

$$E \leq 3\sqrt{p} \left(\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|_{1,2}^q \right)^{1/q} + \sqrt{\frac{m}{N}} \|\mathbf{A}\|.$$

This is the conclusion of the theorem. \square

There is some interesting information in the statement of Theorem 23 about the precise role of the column norms. Adapting an observation of Rudelson and Vershynin [RV07, Lemma 4.1], we see that

$$\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|_{1,2}^2 \leq 2 \max_{|K|=N/m} \frac{m}{N} \sum_{k \in K} \|\mathbf{a}_k\|_2^2.$$

More or less, if we draw m columns at random from a matrix with N columns, we expect the maximum column norm to be controlled by the average of the largest N/m column norms. Therefore, if the columns of \mathbf{A} have wildly disparate ℓ_2 norms, it is valuable to take this fact into account. For clarity of exposition, we use a simpler result that ignores this information.

Corollary 24. *Suppose that \mathbf{A} is a matrix with N columns, and let \mathbf{R} be a restriction to m coordinates, chosen at random from $\llbracket N \rrbracket$. Fix $q > 0$. For each $p \geq \max\{2, 2 \log(\text{rank } \mathbf{A}\mathbf{R}^*), q/2\}$, it holds that*

$$\left(\mathbb{E} \|\mathbf{A}\mathbf{R}^*\|^q \right)^{1/q} \leq 3\sqrt{p} \|\mathbf{A}\|_{1,2} + \sqrt{\frac{m}{N}} \|\mathbf{A}\|.$$

Proof. For any coordinate restriction \mathbf{R} , it is true that $\|\mathbf{A}\mathbf{R}^*\|_{1,2} \leq \|\mathbf{A}\|_{1,2}$, which yields the result for $q \geq 2$. The result for $0 < q < 2$ follows from Hölder's inequality. \square

10. DECOUPLING UNDER THE SPECTRAL NORM

Decoupling is another important method from probability theory that allows one to replace dependent random variables by independent random variables. The following theorem offers a new twist on a classical result from harmonic analysis. See [BT87, Prop. 1.9] and [LT91, Sec. 4.4].

Theorem 25 (Decoupling in the Spectral Norm). *Let \mathbf{A} be a $2N \times 2N$ Hermitian matrix with a zero diagonal, and let \mathbf{R} be a restriction to m random coordinates. Fix $q \geq 1$. There exists a partition of $\llbracket 2N \rrbracket$ into two blocks T_1 and T_2 with N elements each so that*

$$\left(\mathbb{E} \|\mathbf{R}\mathbf{A}\mathbf{R}^*\|^q \right)^{1/q} < 2 \max_{m_1+m_2=m} \left(\mathbb{E} \|\mathbf{R}_1 \mathbf{A}_{T_1 \times T_2} \mathbf{R}_2^*\|^q \right)^{1/q}$$

where

- the maximum occurs over integers $m_1, m_2 \in [0, N]$,
- the symbol $\mathbf{A}_{T_1 \times T_2}$ denotes the submatrix of \mathbf{A} indexed by $T_1 \times T_2$, and
- the matrices \mathbf{R}_i are independent restrictions to m_i random coordinates from T_i for $i = 1, 2$.

When \mathbf{A} has odd order $(2N + 1) \times (2N + 1)$, an analogous result holds for a partition into blocks of size N and $N + 1$.

This theorem offers several important advantages over the classical one. First, the classical result only applies when the restriction selects each coordinate independent from the others, whereas the new result allows us to select exactly m coordinates. In consequence, the conclusions are slightly weaker: we must pass to an unknown submatrix and the restrictions are only conditionally independent. Second, we have exploited properties of the spectral norm to reduce the constant by an order of magnitude.

Proof. We establish the result for the case where the matrix has even order; the other case is essentially the same. For notational simplicity, we take $q = 1$; the other cases are similar.

Define the matrices $\mathbf{B}_{jk} = a_{jk} \mathbf{e}_j \mathbf{e}_k^*$, and let $\boldsymbol{\delta} \in \{0, 1\}^{2N}$ be a random vector with exactly m components equal to one. Then one can express

$$\mathbf{R}\mathbf{A}\mathbf{R}^* = \sum_{j \neq k} \delta_j \delta_k \mathbf{B}_{jk}$$

Our goal is to bound the expectation

$$E \stackrel{\text{def}}{=} \mathbb{E}_{\boldsymbol{\delta}} \left\| \sum_{j \neq k} \delta_j \delta_k \mathbf{B}_{jk} \right\|.$$

Let $\boldsymbol{\eta} \in \{0, 1\}^{2N}$ be a random vector with N components equal to one. When $j \neq k$, it is elementary that

$$\mathbb{E}_{\boldsymbol{\eta}} [\eta_j (1 - \eta_k) + (1 - \eta_j) \eta_k] = \frac{N}{2N - 1}.$$

We insinuate this quantity into the norm, use Jensen's inequality to draw out the expectation, and exchange the indices in the second copy of the sum:

$$\begin{aligned} E &= \frac{2N - 1}{N} \mathbb{E}_{\boldsymbol{\delta}} \left\| \sum_{j \neq k} \mathbb{E}_{\boldsymbol{\eta}} [\eta_j (1 - \eta_k) + \eta_k (1 - \eta_j)] \delta_j \delta_k \mathbf{B}_{jk} \right\| \\ &< 2 \mathbb{E}_{\boldsymbol{\delta}} \mathbb{E}_{\boldsymbol{\eta}} \left\| \sum_{j \neq k} \eta_j \delta_j \cdot (1 - \eta_k) \delta_k \mathbf{B}_{jk} + \sum_{j \neq k} (1 - \eta_j) \delta_j \cdot \eta_k \delta_k \mathbf{B}_{jk} \right\| \\ &= 2 \mathbb{E}_{\boldsymbol{\eta}} \mathbb{E}_{\boldsymbol{\delta}} \left\| \sum_{j \neq k} \eta_j \delta_j \cdot (1 - \eta_k) \delta_k (\mathbf{B}_{jk} + \mathbf{B}_{kj}) \right\|. \end{aligned}$$

There must exist some vector $\boldsymbol{\eta}^*$ for which the inner expectation is no smaller than its average over $\boldsymbol{\eta}$. Define $T_1 = \{j : \eta_j^* = 1\}$ and $T_2 = \{k : \eta_k^* = 0\}$. Note that these sets partition $[2N]$, and each contains N elements. It holds that

$$E < 2 \mathbb{E}_{\boldsymbol{\delta}} \left\| \sum_{\substack{j \in T_1 \\ k \in T_2}} \delta_j \delta_k (\mathbf{B}_{jk} + \mathbf{B}_{kj}) \right\|.$$

Since $\mathbf{B}_{kj} = \mathbf{B}_{jk}^*$, the matrix inside the norm is Hermitian. Since T_1 and T_2 are disjoint, the matrix is also block counter-diagonal. To complete the argument, we must repackage the matrix as a submatrix of \mathbf{A} .

Let the random variable m_1 denote the number of selectors that are active on the set T_1 , and write $m_2 = m - m_1$. Observe that the sequences $\{\delta_j : j \in T_1\}$ and $\{\delta_k : k \in T_2\}$ are conditionally independent, given the numbers m_1 and m_2 . Moreover, under this conditioning, all choices of m_i entries from T_i are equally likely for $i = 1, 2$. If we write \mathbf{R}_i for a random restriction to m_i coordinates from T_i , then we have

$$\begin{aligned} E &< 2 \mathbb{E}_{m_i} \mathbb{E}_{\mathbf{R}_i} \left\| \mathbf{R}_1 \mathbf{A}_{T_1 \times T_2} \mathbf{R}_2^* + \mathbf{R}_2 \mathbf{A}_{T_2 \times T_1} \mathbf{R}_1^* \right\| \\ &= 2 \mathbb{E}_{m_i} \mathbb{E}_{\mathbf{R}_i} \left\| \mathbf{R}_1 \mathbf{A}_{T_1 \times T_2} \mathbf{R}_2^* \right\|, \end{aligned}$$

where the equality holds on account of the identity

$$\left\| \begin{bmatrix} \mathbf{0} & \mathbf{B} \\ \mathbf{B}^* & \mathbf{0} \end{bmatrix} \right\| = \|\mathbf{B}\|.$$

The expectation with respect to the random variable m_1 is certainly less than the maximum over all choices $m_1 + m_2 = m$ with the caveat $0 \leq m_i \leq N$. Hence,

$$E < 2 \max_{m_1 + m_2 = m} \mathbb{E} \left\| \mathbf{R}_1 \mathbf{A}_{T_1 \times T_2} \mathbf{R}_2^* \right\|.$$

This is the advertised conclusion. \square

ACKNOWLEDGMENTS

Most of the work in this paper was completed during visits with NuHAG at the University of Vienna and with the Signal Processing Institute at EPFL. The author warmly acknowledges their hospitality. The author also wishes to thank Holger Rauhut for identifying an error in the discussion of the noncommutative Khintchine inequality.

REFERENCES

- [Aub07] G. Aubrun. Sampling convex bodies: A random matrix approach. *Proc. Amer. Math. Soc.*, 135:1293–1303, 2007. To appear.
- [BDDW07] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. To appear, *Constr. Approx.*, 2007.
- [Bou99] J. Bourgain. Random points in isotropic convex sets. In *Convex geometric analysis (Berkeley, CA, 1996)*, volume 34 of *Math. Sci. Res. Inst. Publ.*, pages 53–58. Cambridge Univ. Press, Cambridge, 1999.
- [BT87] J. Bourgain and L. Tzafriri. Invertibility of “large” submatrices with applications to the geometry of Banach spaces and harmonic analysis. *Israel J. Math.*, 57(2):137–224, 1987.
- [BT91] J. Bourgain and L. Tzafriri. On a problem of Kadison and Singer. *J. reine angew. Math.*, 420:1–43, 1991.
- [Buc01] A. Buchholz. Operator Khintchine inequality in non-commutative probability. *Math. Annalen*, 319:1–16, 2001.
- [CDD06] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k -term approximation. Submitted for publication, July 2006.
- [CDS99] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by Basis Pursuit. *SIAM J. Sci. Comput.*, 20(1):33–61, 1999.
- [CR06] E. J. Candès and J. Romberg. Quantitative robust uncertainty principles and optimally sparse decompositions. *Foundations of Comput. Math.*, 2006. To appear.
- [CRT06] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete Fourier information. *IEEE Trans. Info. Theory*, 52(2):489–509, Feb. 2006.
- [CT06] E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Info. Theory*, 52(12):5406–5425, Dec. 2006.
- [DE03] D. L. Donoho and M. Elad. Maximal sparsity representation via ℓ_1 minimization. *Proc. Natl. Acad. Sci.*, 100:2197–2202, March 2003.
- [DH01] D. L. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Trans. Info. Theory*, 47:2845–2862, Nov. 2001.
- [DMA97] G. Davis, S. Mallat, and M. Avellaneda. Greedy adaptive approximation. *Constr. Approx.*, 13:57–98, 1997.
- [DS89] D. L. Donoho and P. B. Stark. Uncertainty principles and signal recovery. *SIAM J. Appl. Math.*, 49(3):906–931, June 1989.
- [DT06] D. L. Donoho and J. Tanner. Counting faces of randomly projected polytopes when the projection radically lowers dimension. Dept. of Statistics Technical Report 2006-11, Stanford Univ., July 2006.
- [EB02] M. Elad and A. M. Bruckstein. A generalized uncertainty principle and sparse representation in pairs of bases. *IEEE Trans. Info. Theory*, 48(9):2558–2567, 2002.
- [Fuc04] J.-J. Fuchs. On sparse representations in arbitrary redundant bases. *IEEE Trans. Info. Theory*, 50(6):1341–1344, June 2004.
- [GN03] R. Gribonval and M. Nielsen. Sparse representations in unions of bases. *IEEE Trans. Info. Theory*, 49(12):3320–3325, Dec. 2003.
- [Haa82] U. Haagerup. The best constants in the Khintchine inequality. *Studia Math.*, 70:231–283, 1982.
- [KLS97] R. Kannan, L. Lovász, and M. Simonovits. Random walks and $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures Algorithms II*, 13:1–50, 1997.
- [LP86] F. Lust-Picquard. Inégalités de Khintchine dans C_p ($1 < p < \infty$). *Comptes Rendus Acad. Sci. Paris, Série I*, 303(7):289–292, 1986.
- [LT91] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, 1991.
- [MP06] S. Mendelson and A. Pajor. On singular values of matrices with independent rows. *Bernoulli*, 12:761–773, 2006.
- [Nat95] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24:227–234, 1995.
- [PS95] G. Peshkir and A. N. Shiryayev. The Khintchine inequalities and martingale expanding sphere of their action. *Russ. Math. Surveys*, 50(5):849–904, 1995.
- [Rud99] M. Rudelson. Random vectors in the isotropic position. *J. Functional Anal.*, 164:60–72, 1999.

- [RV06] M. Rudelson and R. Vershynin. Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements. In *Proc. 40th Annual Conference on Information Sciences and Systems*, Princeton, Mar. 2006.
- [RV07] M. Rudelson and R. Vershynin. Sampling from large matrices: An approach through geometric functional analysis. To appear, *J. Assoc. Comput. Mach.*, 2007.
- [SH03] T. Strohmer and R. W. Heath. Grassmannian frames with applications to coding and communication. *Appl. Comp. Harmonic Anal.*, 14(3):257–275, May 2003.
- [Sim05] B. Simon. *Trace Ideals and their Applications*. American Mathematical Society, 2nd edition, 2005.
- [TJ89] N. Tomczak-Jaegermann. *Banach–Mazur distances and finite-dimensional operator ideals*. Number 38 in Pitman monographs and surveys in pure and applied mathematics. Longman, 1989.
- [Tro05] J. A. Tropp. Recovery of short, complex linear combinations via ℓ_1 minimization. *IEEE Trans. Info. Theory*, 51(4):1568–1570, April 2005.
- [Tro06a] J. A. Tropp. Just relax: Convex programming methods for subset selection and sparse approximation. *IEEE Trans. Info. Theory*, 52(3):1030–1051, Mar. 2006.
- [Tro06b] J. A. Tropp. The random paving property for uniformly bounded matrices. Submitted for publication, Dec. 2006.
- [Ver05] R. Vershynin. Random sets of isomorphism of linear operators on Hilbert space. In *High Dimensional Probability*, volume 51 of *Lecture Notes — Monograph Series*, pages 148–154. IMS, Dec. 2005.
- [Zyg02] A. Zygmund. *Trigonometric Series*. Cambridge University Press, 3rd edition, 2002.