

Near Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?

Emmanuel Candes[†] and Terence Tao[‡]

[†] Applied and Computational Mathematics, Caltech, Pasadena, CA 91125

[‡] Department of Mathematics, University of California, Los Angeles, CA 90095

October 2004; Revised March 2006

Abstract

Suppose we are given a vector f in a class $\mathcal{F} \subset \mathbb{R}^N$, e.g. a class of digital signals or digital images. How many linear measurements do we need to make about f to be able to recover f to within precision ϵ in the Euclidean (ℓ_2) metric?

This paper shows that if the objects of interest are sparse in a fixed basis or compressible, then it is possible to reconstruct f to within very high accuracy from a small number of random measurements by solving a simple linear program. More precisely, suppose that the n th largest entry of the vector $|f|$ (or of its coefficients in a fixed basis) obeys $|f|_{(n)} \leq R \cdot n^{-1/p}$, where $R > 0$ and $p > 0$. Suppose that we take measurements $y_k = \langle f, X_k \rangle$, $k = 1, \dots, K$, where the X_k are N -dimensional Gaussian vectors with independent standard normal entries. Then for each f obeying the decay estimate above for some $0 < p < 1$ and with overwhelming probability, our reconstruction f^\sharp , defined as the solution to the constraints $y_k = \langle f^\sharp, X_k \rangle$ with minimal ℓ_1 norm, obeys

$$\|f - f^\sharp\|_{\ell_2} \leq C_p \cdot R \cdot (K/\log N)^{-r}, \quad r = 1/p - 1/2.$$

There is a sense in which this result is optimal; it is generally impossible to obtain a higher accuracy from any set of K measurements whatsoever. The methodology extends to various other random measurement ensembles; for example, we show that similar results hold if one observes few randomly sampled Fourier coefficients of f . In fact, the results are quite general and require only two hypotheses on the measurement ensemble which are detailed.

Keywords. Random matrices, singular values of random matrices, signal recovery, random projections, concentration of measure, sparsity, trigonometric expansions, uncertainty principle, convex optimization, duality in optimization, linear programming.

Acknowledgments. E. C. is partially supported by National Science Foundation grants DMS 01-40698 (FRG) and ACI-0204932 (ITR), and by an Alfred P. Sloan Fellowship. T. T. is supported in part by a grant from the Packard Foundation. E. C. would like to thank Justin Romberg and Houman Ohwadi for conversations related to this project, and Nouredine El Karoui for sending us an early draft of his Ph. D. thesis. E. C. would also like to thank Joel Tropp for sharing some notes. T. T. thanks Wilhelm Schlag and Gerd Mockenhaupt for pointing out the relevance of Bourgain's work [6], [7]. Finally, we would like to thank the anonymous referees for providing useful references (see Section 1.4) and for their role in improving the original manuscript.

1 Introduction and Overview of the Main Results

This paper considers the fundamental problem of recovering a finite signal $f \in \mathbb{R}^N$ from a limited set of measurements. Specifically, given a class of signals $\mathcal{F} \subset \mathbb{R}^N$, one is interested in the minimum number of linear measurements one has to make to be able to reconstruct objects from \mathcal{F} to within a fixed accuracy ϵ , say, in the usual Euclidean ℓ_2 -distance. In other words, how can one specify $K = K(\epsilon)$ linear functionals

$$y_k = \langle f, \psi_k \rangle, \quad k \in \Omega, \quad (1.1)$$

where $(\psi_k)_{k \in \Omega}$ is a set of vectors with cardinality $|\Omega| = K$, so that it is possible to reconstruct an object f^\sharp from the data $(y_k)_{k \in \Omega}$ obeying

$$\|f - f^\sharp\|_{\ell_2} \leq \epsilon, \quad (1.2)$$

for each element f taken from \mathcal{F} ? The primary goal is of course, to find appropriate functionals $(\psi_k)_{k \in \Omega}$ so that the required number K of measurements is as small as possible. In addition, we are also interested in concrete and practical recovery algorithms.

The new results in this paper will address this type of question for signals f whose coefficients with respect to a fixed reference basis obey a power-law type decay condition, and for random measurements $(y_k)_{k \in \Omega}$ sampled from a specified ensemble. However, before we discuss these results, we first recall some earlier results concerning signals of small support. (See also Sections 1.7 and 9.2 for a more extensive discussion of related results.)

1.1 Exact Reconstruction of Sparse Signals

In a previous article [11], the authors together with J. Romberg studied the recovery of *sparse* signals from limited measurements; i.e. of signals which have relatively few nonzero terms or whose coefficients in some fixed basis have relatively few nonzero entries. This paper discussed some surprising phenomena, and we now review a special instance of those. In order to do so, we first need to introduce the discrete Fourier transform which is given by the usual formula¹

$$\hat{f}(k) := \frac{1}{\sqrt{N}} \sum_{t \in \mathbb{Z}_N} f(t) e^{-i2\pi kt/N}, \quad (1.3)$$

where the frequency index k ranges over the set $\mathbb{Z}_N := \{0, 1, \dots, N-1\}$.

Suppose then that we wish to recover a signal $f \in \mathbb{R}^N$ made out of $|T|$ spikes, where the set T denotes the support of the signal

$$T := \{t, f(t) \neq 0\}.$$

We do not know where the spikes are located nor do we know their amplitudes. However, we are given information about f in the form of ‘only’ K randomly sampled Fourier coefficients $F_\Omega f := (\hat{f}(k))_{k \in \Omega}$ where Ω is a random set of K frequencies sampled uniformly at random.

¹Strictly speaking, the Fourier transform is associated to an orthonormal basis in \mathbf{C}^N rather than \mathbb{R}^N . However all of our analysis here extends easily to complex signals instead of real signals (except for some negligible changes in the absolute constants C). For ease of exposition we shall focus primarily on real-valued signals $f \in \mathbb{R}^N$, except when referring to the Fourier basis.

In [11], it was shown that f could be reconstructed *exactly* from these data provided that the expected number of frequency samples obeyed the lower bound

$$|T| \leq \alpha \cdot (K/\log N) \quad (1.4)$$

for all sufficiently small $\alpha > 0$ (i.e. $\alpha \leq \alpha_0$ for some small absolute constant α_0). To recover f from $F_\Omega f$, we simply minimize the ℓ_1 -norm of the reconstructed signal

$$\min \|g\|_{\ell_1} := \sum_{t \in \mathbb{Z}_N} |g(t)|, \quad (1.5)$$

subject to the constraints

$$\hat{g}(k) = \hat{f}(k), \quad \forall k \in \Omega.$$

Moreover, the probability that exact recovery occurs exceeds $1 - O(N^{-\rho/\alpha})$; $\rho > 0$ is here a universal constant and it is worth noting that the aforementioned reference gave explicit values for this constant. The implied constant in the $O()$ notation is allowed to depend on α , but is independent of N . In short, exact recovery may be achieved by solving a simple convex optimization problem—in fact, a linear program for real-valued signals—which is a result of practical significance.

In a following paper, Candès and Romberg [13] extended these results and showed that exact reconstruction phenomena hold for other synthesis/measurement pairs. For clarity of presentation, it will be convenient to introduce some notations that we will use throughout the remainder of the paper. We let F_Ω denote the $|\Omega|$ by N matrix which specifies the set of those $|\Omega|$ linear functionals which describe the measurement process so that the available information y about f is of the form

$$y = F_\Omega f.$$

For instance, in our previous example, F_Ω is the $|\Omega|$ by N partial Fourier matrix whose rows are the sampled sinusoids

$$F_\Omega(k, t) := \frac{1}{\sqrt{N}} e^{-i2\pi kt/N}, \quad k \in \Omega, t \in \mathbb{Z}_N.$$

More generally, suppose that one is given an orthonormal basis Ψ

$$\Psi = (\psi_k(t))_{0 \leq t, k < N},$$

and that one has available partial information about f in the sense that we have knowledge about a randomly selected set $\Omega \subset \{0, \dots, N-1\}$ of coefficients in basis Ψ . For convenience, define Ψ to be the N by N synthesis matrix with entries $\Psi(t, k) := \psi_k(t)$. Then F_Ω is now obtained from Ψ^* by extracting the $|\Omega|$ rows with indices k obeying $k \in \Omega$. Suppose as before that there is another (fixed) orthonormal basis Φ in which the coefficients $\theta(f) = (\theta_t(f))_{1 \leq t \leq N}$ of f in this basis, defined by

$$\theta_t(f) := \langle f, \phi_t \rangle,$$

are sparse in the sense that only few of the entries of $\theta(f)$ are nonzero. Then it was shown in [13] that with overwhelming probability, f is the solution to

$$\min_g \|\theta(g)\|_{\ell_1} \quad \text{subject to} \quad F_\Omega g = F_\Omega f. \quad (1.6)$$

That is, exact reconstruction still occurs; the relationship here between the number of nonzero terms in the basis Φ and the number of observed coefficients $|\Omega|$ depends upon the *incoherence* between the two bases. The more incoherent, the fewer coefficients needed; in the other direction, in the maximally coherent case, e.g. $\Psi = \Phi$, one in fact needs to sample essentially all of the coefficients in order to ensure exact reconstruction (the same holds true if Φ and Ψ share only one element with nonzero inner product with f).

A special instance of these results concerns the case where the set of measurements is generated completely at random; that is, we sample a random orthonormal basis of \mathbb{R}^N and observe only the first K coefficients in that basis (note that there is no advantage in randomizing Ω as in Section 1.1 since the basis is already completely random). As before, we let F_Ω be the submatrix enumerating those sampled vectors and solve (1.6). Then a consequence of the methodology developed in this paper is that exact reconstruction occurs with probability at least $1 - O(N^{-\rho/\alpha})$ (for a different value of ρ) provided that

$$\|\theta(f)\|_{\ell_0} \leq \alpha \cdot (K/\log N), \quad (1.7)$$

where $\alpha > 0$ is sufficiently small, and the ℓ_0 -norm is of course the size of the support of the vector θ

$$\|\theta\|_{\ell_0} := |\{t : \theta_t \neq 0\}|,$$

see [23] for sharper results. In summary, ℓ_1 seems to recover sparse unknown signals in a variety of different situations. The number of measurements simply needs to exceed the number of unknown nonzero coefficients by a proper amount.

Observe that a nice feature of the random basis discussed above is its statistical invariance by rotation. Let Φ be any basis so that $\theta(f)$ are the coefficients of f in that basis: $\theta(f) := \Phi^* f$. The constraints in (1.6) impose

$$F_\Omega \Phi \theta(g) = F_\Omega \Phi \theta(f)$$

and since the distribution of $F_\Omega \Phi$ is that of F_Ω , the choice of the basis Φ is actually irrelevant. Exact reconstruction occurs (with overwhelming probability) when the signal is sparse in any fixed basis; of course, the recovery algorithm requires knowledge of this basis.

1.2 Power laws

In general, signals of practical interest may not be supported in space or in a transform domain on a set of relatively small size. Instead, the coefficients of elements taken from a signal class decay rapidly, typically like a power law [20, 22]. We now give two examples leaving mathematical rigor aside in the hope of being more concise.

- *Smooth signals.* It is well-known that if a continuous-time object has s bounded derivatives, then the n th largest entry of the wavelet or Fourier coefficient sequence is of size about $1/n^{s+1/2}$ in one dimension and more generally, $1/n^{s/d+1/2}$ in d dimensions [22]. Hence, the decay of Fourier or wavelet coefficients of smooth signals exhibits a power law.
- *Signals with bounded variations.* A popular model for signal/analysis is the space of objects with bounded variations. At the level of the continuum, the total-variation norm of an object is approximately the ℓ_1 norm of its gradient. In addition, there

are obvious discrete analogs for finite signals where the gradient is replaced by finite differences. Now a norm which is almost equivalent to the total-variation norm is the weak- ℓ_1 norm in the wavelet domain; that is, the reordered wavelet coefficients of a compactly supported object f approximately decay like $1/n$ [18]. At the discrete level, $\|f\|_{BV}$ essentially behaves like the ℓ_1 -norm of the Haar wavelet coefficients up to a multiplicative factor of at most $\log N$. Moreover, it is interesting to note that studies show that the empirical wavelet coefficients of photographs of natural scenes actually exhibit the $1/n$ -decay [41].

In fact, finding representations with rapidly decaying coefficients is a very active area of research known as Computational Harmonic Analysis and there are of course many other such examples. For instance, certain classes of oscillatory signals have rapidly decaying Gabor coefficients [29], certain types of images with discontinuities along edges have rapidly decaying curvelet coefficients [10] and so on.

Whereas [11] considered signals f of small support, we now consider objects whose coefficients in some basis decay like a power-law. We fix an orthonormal basis $\Phi = (\phi_t)_{1 \leq t \leq N}$ (which we call the *reference basis*), and rearrange the entries $\theta_t(f) := \langle f, \phi_t \rangle$ of the coefficient vector $\theta(f)$ in decreasing order of magnitude $|\theta|_{(1)} \geq |\theta|_{(2)} \geq \dots \geq |\theta|_{(N)}$. We say that $\theta(f)$ belongs to the weak- ℓ_p ball or radius R (and we will sometimes write $f \in w\ell_p(R)$) for some $0 < p < \infty$ and $C > 0$ if for each $1 \leq n \leq N$,

$$|\theta|_{(n)} \leq R \cdot n^{-1/p}. \quad (1.8)$$

In other words, p controls the speed of the decay: the smaller p , the faster the decay. The condition (1.8) is also equivalent to the estimate

$$|\{t \in \mathbb{Z}_N : |\theta_t(f)| \geq \lambda\}| \leq \frac{R^p}{\lambda^p}$$

holding for all $\lambda > 0$. We shall focus primarily on the case $0 < p < 1$.

It is well-known that the decay rate of the coefficients of f is linked to the ‘compressibility’ of f , compare the widespread use of transform coders in the area of lossy signal or image compression. Suppose for instance that all the coefficients $(\theta_t(f))_{1 \leq t \leq N}$ are known and consider the partial reconstruction $\theta_K(f)$ (where $1 \leq K \leq N$ is fixed) obtained by keeping the K largest entries of the vector $\theta(f)$ (and setting the others to zero). Then it immediately follows from (1.8) that the approximation error obeys

$$\|\theta(f) - \theta_K(f)\|_{\ell_2} \leq C_p \cdot R \cdot K^{-r}, \quad r := 1/p - 1/2,$$

for some constant C_p which only depends on p . And thus, it follows from Parseval that the approximate signal f_K obtained by keeping the largest coefficients in the expansion of f in the reference basis Φ obeys the same estimate, namely,

$$\|f - f_K\|_{\ell_2} \leq C_p \cdot R \cdot K^{-r}, \quad (1.9)$$

where C_p only depends on p .

1.3 Recovery of objects with power-law decay

We now return to the setup we discussed earlier, where we select K orthonormal vectors ψ_1, \dots, ψ_K in \mathbb{R}^N uniformly at random. Since applying a fixed orthonormal transform does

not change the problem, we may just as well assume that Φ is the identity and solve

$$(P_1) \quad \min_{g \in \mathbb{R}^N} \|g\|_{\ell_1} \quad \text{subject to} \quad F_\Omega g = F_\Omega f, \quad (1.10)$$

where as usual, $F_\Omega f = (\langle f, \psi_k \rangle)_{k \in \Omega}$. In the setting where f does not have small support, we do not expect the recovery procedure (1.10) to recover f exactly, but our first main theorem asserts that it will recover f *approximately*.

Note. From now on and for ease of exposition, we will take (P_1) as our abstract recovery procedure where it is understood that f is the sparse object of interest to be recovered; that is, f could be a signal in \mathbb{R}^N or its coefficients in some fixed basis Φ .

Theorem 1.1 (Optimal recovery of $w\ell_p$ from random measurements) *Suppose that $f \in \mathbb{R}^N$ obeys (1.8) for some fixed $0 < p < 1$ or $\|f\|_{\ell_1} \leq R$ for $p = 1$, and let $\alpha > 0$ be a sufficiently small number (less than an absolute constant). Assume that we are given K random measurements $F_\Omega f$ as described above. Then with probability 1, the minimizer f^\sharp to (1.10) is unique. Furthermore, with probability at least $1 - O(N^{-\rho/\alpha})$, we have the approximation*

$$\|f - f^\sharp\|_{\ell_2} \leq C_{p,\alpha} \cdot R \cdot (K/\log N)^{-r}, \quad r = 1/p - 1/2. \quad (1.11)$$

Here, $C_{p,\alpha}$ is a fixed constant depending on p and α but not on anything else. The implicit constant in $O(N^{-\rho/\alpha})$ is allowed to depend on α .

The result of this theorem may seem surprising. Indeed, (1.11) says that if one makes $O(K \log N)$ random measurements of a signal f , and then reconstructs an approximate signal from this limited set of measurements in a manner which requires no prior knowledge or assumptions on the signal (other than it perhaps obeys some sort of power law decay with unknown parameters) *one still obtains a reconstruction error which is equally as good as that one would obtain by knowing everything about f and selecting the K largest entries of the coefficient vector $\theta(f)$* ; thus the amount of “oversampling” incurred by this random measurement procedure compared to the optimal sampling for this level of error is only a multiplicative factor of $O(\log N)$. To avoid any ambiguity, when we say that no prior knowledge or information is required about the signal, we mean that the reconstruction algorithm does not depend upon unknown quantities such as p or R .

Below, we will argue that we cannot, in general, design a set of K measurements that would allow essentially better reconstruction errors by any method, no matter how intractable. As we will see later, Theorem 1.1 is a special case of Theorem 1.4 below (but for the uniqueness claim which is proved in Section 3).

1.4 Precedents

A natural question is whether the number of random samples we identified in Theorem 1.4 is, in some sense, optimal. Or would it be possible to obtain similar accuracies with far fewer observations? To make things concrete, suppose we are interested in the recovery of objects with bounded ℓ_1 -norm, e.g. the ℓ_1 -ball

$$\mathcal{B}_1 := \{f : \|f\|_{\ell_1} \leq 1\}.$$

Suppose we can make K linear measurements about $f \in \mathcal{B}_1$ of the form $y = F_\Omega f$. Then what is the best measurement/reconstruction pair so that the error

$$E_K(\mathcal{B}_1) = \sup_{f \in \mathcal{B}_1} \|f - D(y)\|_2, \quad y = F_\Omega f, \quad (1.12)$$

is minimum? In (1.12), D is the reconstruction algorithm.

To develop an insight about the intrinsic difficulty of our problem, consider the following geometric picture. Suppose we take K measurements $F_\Omega f$; this says that f belongs to an affine space $f_0 + S$ where S is a linear subspace of co-dimension less or equal to K . Now *the data available for the problem cannot distinguish any object belonging to that plane*. Assume f is known to belong to the ℓ_1 -ball \mathcal{B}_1 , say, then the data cannot distinguish between any two points in the intersection $\mathcal{B}_1 \cap f_0 + S$. Therefore, any reconstruction procedure $f^*(y)$ based upon $y = F_\Omega f$ would obey

$$\sup_{f \in \mathcal{F}} \|f - f^*\| \geq \frac{\text{diam}(\mathcal{B}_1 \cap S)}{2}. \quad (1.13)$$

(When we take the supremum over all f , we may just assume that f be orthogonal to the measurements ($y = 0$) since the diameter will of course be maximal in that case.) The goal is then to find S such that the above diameter be minimal. This connects with the agenda of approximation theory where this problem is known as finding the Gelfand n -width of the class \mathcal{B}_1 [45], as we explain below.

The *Gelfand* numbers of a set \mathcal{F} are defined as

$$d_K(\mathcal{F}) = \inf_S \{ \sup_{f \in \mathcal{F}} \|P_S f\| : \text{codim}(S) < K \}, \quad (1.14)$$

where P_S is, of course, the orthonormal projection on the subspace S . Then it turns out that $d_K(\mathcal{F}) \leq E_K(\mathcal{F}) \leq C_0 d_K(\mathcal{F})$. Now a seminal result of Kashin [37] and improved by Garnaev and Gluskin [31, 3] shows that for the ℓ_1 ball, the Gelfand numbers obey

$$C \sqrt{\frac{\log(N/K) + 1}{K}} \leq d_k(\mathcal{B}_1) \leq C' \sqrt{\frac{\log(N/K) + 1}{K}}, \quad (1.15)$$

where C, C' are universal constants. Gelfand numbers are also approximately known for weak- ℓ_p balls as well.

Viewed differently, Kashin, Garnaev and Gluskin assert that with K measurements, the minimal reconstruction error (1.12) one can hope for is bounded below by a constant times $(K/\log(N/K))^{-1/2}$. In this sense, Theorem 1.4 is optimal (within a multiplicative constant) at least for $K \asymp N^\beta$, with $\beta < 1^2$. Kashin also shows that if we take a random projection, $\text{diam}(\mathcal{B}_1 \cap S)$ is bounded above by the right-hand side of (1.15). We would also like to emphasize that similar types of recovery have also been known to be possible in the literature of theoretical computer science, at least in principle, for certain types of random measurements [1]. On the one hand, finding the Chebyshev center of $\text{diam}(\mathcal{B}_1 \cap S)$ is a convex problem, which would yield a near-optimal reconstruction algorithm. On the other hand, this problem is computationally intractable when $p < 1$. Further, one would need to know p and the radius of the weak- ℓ_p ball which is not realistic in practical applications.

²Note added in proof: since submission of this paper, we proved in [12] that Theorem 1.4 holds with $\log(N/K)$ instead of $\log N$ in (1.11).

The novelty here is that the information about f can be retrieved from those random coefficients by minimizing a simple linear program (1.10), and that the decoding algorithm adapts automatically to the weak- ℓ_p signal class, without knowledge thereof. Minimizing the ℓ_1 -norm gives nearly the best possible reconstruction error simultaneously over a wide range of sparse classes of signals; no information about p and the radius R are required. In addition and as we will see next, another novelty is the general nature of the measurement ensemble.

It should also be mentioned that when the measurement ensemble consists of Fourier coefficients on a random arithmetic progression, a very fast recovery algorithm that gives near-optimal results for arbitrary ℓ_2 data has recently been given in [33]. Since the preparation of this manuscript, we have learnt that results closely related to those in this paper have appeared in [24]. We compare our results with both these works in Section 9.2.

1.5 Other Measurement Ensembles

Underlying our results is a powerful machinery essentially relying on properties of random matrices which gives us very precise tools allowing to quantify how much of a signal one can reconstruct from random measurements. In fact, Theorem 1.1 holds for other *measurement ensembles*. For simplicity, we shall consider three types of measured data:

- *The Gaussian ensemble:* Here, we suppose that $1 \leq K \leq N$ and $\Omega := \{1, \dots, K\}$ are fixed, and the entries of F_Ω are identically and independently sampled from a standard normal distribution

$$F_\Omega(k, t) := \frac{1}{\sqrt{N}} X_{kt}, \quad X_{kt} \text{ i.i.d. } N(0, 1).$$

The Gaussian ensemble is invariant by rotation since for any fixed orthonormal matrix Φ , the distribution of F_Ω is that of $F_\Omega \Phi$.

- *The binary ensemble:* Again we take $1 \leq K \leq N$ and $\Omega := \{1, \dots, K\}$ to be fixed. But now we suppose that the entries of F_Ω are identically and independently sampled from a symmetric Bernoulli distribution

$$F_\Omega(k, t) := \frac{1}{\sqrt{N}} X_{kt}, \quad X_{kt} \text{ i.i.d. } P(X_{kt} = \pm 1) = 1/2.$$

- *The Fourier ensemble:* This ensemble was discussed earlier, and is obtained by randomly sampling rows from the orthonormal N by N Fourier matrix $\mathcal{F}(k, t) = \exp(-i2\pi kt/N)/\sqrt{N}$. Formally, we let $0 < \tau < 1$ be a fixed parameter, and then let Ω be the random set defined by

$$\Omega = \{k : I_k = 1\},$$

where the I_k 's are i.i.d. Bernoulli variables with $\mathbf{P}(I_k = 1) = \tau$. We then let $R_\Omega : \ell_2(\mathbb{Z}_N) \rightarrow \ell_2(\Omega)$ be the restriction map $(R_\Omega g)(k) = g(k)$ for all $k \in \Omega$ (so that the adjoint $R_\Omega^* : \ell_2(\Omega) \rightarrow \ell_2(\mathbb{Z}_N)$ is the embedding obtained by extending by zero outside of Ω), and set

$$F_\Omega := R_\Omega \mathcal{F}.$$

In this case, the role of K is played by the quantity $K := \mathbf{E}(|\Omega|) = \tau N$. (In fact $|\Omega|$ is usually very close to K ; see Lemma 6.6).

Just as Theorem 1.1 suggests, this paper will show that it is possible to derive recovery rates for all three measurement ensembles. The ability to recover a signal f from partial random measurements depends on key properties of those measurement ensembles that we now discuss.

1.6 Axiomatization

We shall now unify the treatment of all these ensembles by considering an abstract *measurement matrix* F_Ω , which is a random $|\Omega| \times N$ matrix following some probability distribution (e.g. the Gaussian, Bernoulli, or Fourier ensembles). We also allow the number of measurements $|\Omega|$ to be a random variable taking values between 1 and N , and set $K := \mathbf{E}(|\Omega|)$ —the expected number of measurements. For ease of exposition we shall restrict our attention to real-valued matrices F_Ω ; the modifications required to cover complex matrices such as those given by the Fourier ensemble are simple. We remark that we do not assume that the rows of the matrix F_Ω form an orthogonal family.

This section introduces two key properties on F_Ω which—if satisfied—will guarantee that the solution to the problem (1.10) will be a good approximation to the unknown signal f in the sense of Theorem 1.1.

First, as in [11], our arguments rely, in part, on the quantitative behavior of the singular values of the matrices $F_{\Omega T} := F_\Omega R_T^* : \ell_2(T) \rightarrow \ell_2(\Omega)$ which are the $|\Omega|$ by $|T|$ matrices obtained by extracting $|T|$ columns from F_Ω (corresponding to indices in a set T). More precisely, we shall need to assume the following hypothesis concerning the minimum and maximum eigenvalues of the square matrix $F_{\Omega T}^* F_{\Omega T} : \ell_2(T) \rightarrow \ell_2(T)$.

Definition 1.2 (UUP: Uniform Uncertainty Principle) *We say that a measurement matrix F_Ω obeys the uniform uncertainty principle with oversampling factor λ if for every sufficiently small $\alpha > 0$, the following statement is true with probability at least³ $1 - O(N^{-\rho/\alpha})$ for some fixed positive constant $\rho > 0$: for all subsets T such that*

$$|T| \leq \alpha \cdot K/\lambda, \tag{1.16}$$

the matrix $F_{\Omega T}$ obeys the bounds

$$\frac{1}{2} \cdot \frac{K}{N} \leq \lambda_{\min}(F_{\Omega T}^* F_{\Omega T}) \leq \lambda_{\max}(F_{\Omega T}^* F_{\Omega T}) \leq \frac{3}{2} \cdot \frac{K}{N}. \tag{1.17}$$

Note that (1.17) is equivalent to the inequality

$$\frac{1}{2} \frac{K}{N} \|f\|_{\ell_2}^2 \leq \|F_\Omega f\|_{\ell_2}^2 \leq \frac{3}{2} \frac{K}{N} \|f\|_{\ell_2}^2 \tag{1.18}$$

holding for all signals f with support size less or equal to $\alpha K/\lambda$.

There is nothing special about the constants $1/2$ and $3/2$ in (1.17), which we merely selected to make the **UUP** as concrete as possible. Apart from the size of certain numerical constants (in particular, implied constants in the $O()$ notation), nothing in our arguments depends on this special choice, and we could replace the pair $(1/2, 3/2)$ with a pair (a, b) where a

³Throughout this paper, we allow implicit constants in the $O()$ notation to depend on α .

and b are bounded away from zero and infinity. This remark is important to keep in mind when we will discuss the **UUP** for binary matrices.

To understand the content of (1.17), suppose that F_Ω is the partial Fourier transform and suppose we have a signal f supported on a set T obeying $|T| \leq \alpha K/\lambda$. Then (1.17) says that $\|\hat{f}\|_{\ell_2(\Omega)}$ is at most $\sqrt{3K/2N}\|f\|_{\ell_2}$ with overwhelming probability. Comparing this with Plancherel's identity $\|\hat{f}\|_{\ell_2(\mathbb{Z}_N)} = \|f\|_{\ell_2}$, we see that (with overwhelming probability) a sparse signal f cannot be concentrated in frequency on Ω *regardless of the exact support of f* , unless K is comparable to N . This justifies the terminology ‘‘Uncertainty Principle’’. A subtle but crucial point here is that, with overwhelming probability, we obtain the estimate (1.17) for *all* sets T obeying (1.16); this is stronger than merely asserting that each set T obeying (1.16) obeys (1.17) separately with overwhelming probability, since in the latter case the number of sets T obeying (1.16) is quite large and thus the union of all the exceptional probability events could thus also be quite large. This justifies the terminology ‘‘Uniform’’. As we will see in Section 3, the uniform uncertainty principle hypothesis is crucial to obtain estimates about the ℓ_2 distance between the reconstructed signal $f^\#$ and the unknown signal f .

The **UUP** is similar in spirit to several standard principles and results regarding random projection, such as the famous Johnson-Lindenstrauss lemma [35] regarding the preservation of distances between a finite number of points when randomly projected to a medium-dimensional space. There are however a number of notable features of the **UUP** that distinguish it from more standard properties of random projections. Firstly, there is a wide latitude in how to select the measurement ensemble F_Ω ; for instance, the entries do not have to be independent or Gaussian, and it is even conceivable that interesting classes of completely deterministic matrices obeying the **UUP** could be constructed. Secondly, the estimate (1.17) has to hold for *all* subsets T of a certain size; for various reasons in our applications, it would not be enough to have (1.17) merely on an overwhelming proportion of such sets T . This makes it somewhat trickier for us to verify the **UUP**; in the Fourier case we shall be forced to use some entropy counting methods of Bourgain.

We now introduce a second hypothesis (which appears implicitly in [11], [13]) whose significance is explained below.

Definition 1.3 (ERP: Exact Reconstruction Principle) *We say that a measurement matrix F_Ω obeys the exact reconstruction principle with oversampling factor λ if for all sufficiently small $\alpha > 0$, each fixed subset T obeying (1.16) and each ‘sign’ vector σ defined on T , $|\sigma(t)| = 1$, there exists with overwhelmingly large probability a vector $P \in \mathbb{R}^N$ with the following properties:*

- (i) $P(t) = \sigma(t)$, for all $t \in T$;
- (ii) P is a linear combination of the rows of F_Ω (i.e. $P = F_\Omega^* V$ for some vector V of length $|\Omega|$);
- (iii) and $|P(t)| \leq \frac{1}{2}$ for all $t \in T^c := \{0, \dots, N-1\} \setminus T$.

By ‘overwhelmingly large’, we mean that the probability be at least $1 - O(N^{-\rho/\alpha})$ for some fixed positive constant $\rho > 0$ (recall that the implied constant is allowed to depend on α).

Section 2 will make clear that **ERP** is crucial to check that the reconstruction f^\sharp is close, in the ℓ_1 -norm, to the vector obtained by truncating f , keeping only its largest entries. Note that, in contrast to the **UUP**, in **ERP** we allow a separate exceptional event of small probability for *each* set T , rather than having a uniform event of high probability that covers all T at once. There is nothing special about the factor $1/2$ in (iii); any quantity β strictly between 0 and 1 would suffice here.

To understand how **ERP** relates to our problem, suppose that f is a signal supported on a set T . Then using duality theory, it was shown in [11] that the solution to (1.10) is exact if and only if there exist a P with the above properties for $\sigma(t) = \text{sgn}(f)(t)$ —hence the name.

The hypotheses **UUP** and **ERP** are closely related. For instance, one can use **UUP** to prove a statement very similar to **ERP**, but in the ℓ_2 norm rather than the ℓ_∞ norm; see Corollary 3.1 below. One also has an implication of the form **UUP** \implies **ERP** for *generic* signals f assuming an additional weaker hypothesis **WERP**, see Section 5. In [11] and in [13], the property **UUP** was used (together with some additional arguments) to deduce⁴ **ERP**.

We now are in position to state the main result of this paper.

Theorem 1.4 *Let F_Ω be a measurement process such that **UUP** and **ERP** hold with oversampling factors λ_1 and λ_2 respectively. Put $\lambda = \max(\lambda_1, \lambda_2)$ and assume $K \geq \lambda$. Suppose that f is a signal in \mathbb{R}^N obeying (1.8) for some fixed $0 < p < 1$ or $\|f\|_{\ell_1} \leq R$ for $p = 1$, and let $r := 1/p - 1/2$. Then for any sufficiently small α , any minimizer f^\sharp to (1.10) will obey*

$$\|f - f^\sharp\|_{\ell_2} \leq C_{p,\alpha} \cdot R \cdot (K/\lambda)^{-r} \quad (1.19)$$

with probability at least $1 - O(N^{-\rho/\alpha})$. The implied constant may depend on p and α but not on anything else.

In this paper, we will show that the Gaussian and binary ensembles mentioned earlier obey **UUP** and **ERP** with $\lambda = \log N$, while the Fourier ensemble obeys **UUP** with $\lambda = (\log N)^6$ and **ERP** with $\lambda = \log N$. Hence given an object $f \in w\ell_p(R)$, we prove that if we collect $K \geq \log N$ Gaussian or binary measurements, then

$$\|f - f^\sharp\|_{\ell_2} \leq O(1) \cdot R \cdot (K/\log N)^{-r} \quad (1.20)$$

except for a set of probability at most $O(N^{-\rho/\alpha})$. For randomly sampled frequency data (with at least $(\log N)^6$ frequencies being sampled), the quality of the reconstruction now reads as

$$\|f - f^\sharp\|_{\ell_2} \leq O(1) \cdot R \cdot (K/(\log N)^6)^{-r}. \quad (1.21)$$

We prove this theorem in Section 3.2. Observe that our earlier Theorem 1.1 follows from (1.20) and is thus a special case of Theorem 1.4. Indeed, for a fixed F_Ω , (1.10) is equivalent to

$$\min_g \|\theta(g)\|_{\ell_1} \quad \text{subject to} \quad P_\Omega g = P_\Omega f.$$

⁴Note added in proof: In a sequel [12] to this paper, we show that a slight strengthening of the **UUP** (in which the constants $\frac{1}{2}$ and $\frac{3}{2}$ are replaced by other numerical constants closer to 1) in fact implies **ERP** unconditionally.

where P_Ω is the orthogonal projection onto the span of the rows of F_Ω . Now suppose as in the Gaussian ensemble that F_Ω is a matrix with i.i.d. $N(0, 1/N)$ entries, then P_Ω is simply the projection onto a random plane of dimension K (with probability 1) which, of course, is the setup of Theorem 1.1.

1.7 About the ℓ_1 norm

We would like to emphasize that the simple nonlinear reconstruction strategy which minimizes the ℓ_1 -norm subject to consistency with the measured observations is well-known in the literature of signal processing. For example in the mid-eighties, Santosa and Symes proposed this rule to reconstruct spike trains from incomplete data [49], see also [21]. We would also like to point out connections with total-variation approaches in the literature of image processing [48, 11] which are methods based on the minimization of the ℓ_1 -norm of the discrete gradient. Note that minimizing the ℓ_1 -norm is very different than standard least squares (i.e. ℓ_2) minimization procedures. With incomplete data, the least square approach would simply set to zero the ‘unobserved’ coefficients. Consider the Fourier case, for instance. The least-squares solution would set to zero all the unobserved frequencies so that the minimizer would have much smaller energy than the original signal. As is well known, the minimizer would also contain a lot of artifacts.

More recently, ℓ_1 -minimization perhaps best known under the name of *Basis Pursuit*, has been proposed as a convex alternative to the combinatorial norm ℓ_0 , which simply counts the number of nonzero entries in a vector, for synthesizing signals as sparse superpositions of waveforms [17]. Interestingly, these methods provided great practical success [17, 16] and were shown to enjoy remarkable theoretical properties and to be closely related to various kinds of uncertainty principles [26, 27, 34, 14].

On the practical side, an ℓ_1 -norm minimization problem (for real-valued signals) can be recast as a linear program (LP) [4]. For example, (1.10) is equivalent to minimizing $\sum_t u(t)$ subject to $F_\Omega g = F_\Omega f$ and $-u(t) \leq g(t) \leq u(t)$ for all t . This is interesting since there is a wide array of ever more effective computational strategies for solving LPs.

1.8 Applications

In many applications of practical interest, we often wish to reconstruct an object (a discrete signal, a discrete image and so on) from incomplete samples and it is natural to ask how much one can hope to recover. Actually, this work was motivated by the problem of reconstructing biomedical images from vastly undersampled Fourier data. Of special interest are problems in magnetic resonance (MR) angiography but it is expected that our methodology and algorithms will be suitable for other MR imagery, and to other acquisition techniques, such as tomography. In MR angiography, however, we observe few Fourier samples, and therefore if the images of interest are compressible in some transform domain such as in the wavelet domain for example, then ℓ_1 -based reconstructions might be especially well-suited.

Another application of these ideas might be to view the measurement/reconstruction procedure as a kind of lossy encoder/decoder pair where the measurement process would play the role of an encoder and the linear program (P_1) that of a decoder. We postpone this discussion to Section 8.

1.9 Organization of the Paper

This paper is roughly divided into three parts and is organized as follows. The first part (Sections 2 and 3), shows how **UUP** together with **ERP** give our main result, namely, Theorem 1.4. In Section 2, we establish that the solution to (1.10) is in some sense stable in the ℓ_1 -norm, while Section 3 introduces some ℓ_2 -theory and proves our main result. In the second part (Sections 4, 5, 6 and 7), we show that all three measurement ensembles obey **UUP** and **ERP**. Section 4 studies singular values of random matrices and shows that the **UUP** holds for the Gaussian and binary ensembles. Section 5 presents a weaker **ERP** which, in practice, is far easier to check. In Section 6, we prove that all three ensembles obey the **ERP**. In the case of the Fourier ensemble, the strategy for proving the **UUP** is very different than for Gaussian and binary measurements, and is presented in a separate Section 7. Finally, we will argue in the third part of the paper that one can think of the random measurement process as some kind of universal encoder (Section 8) and briefly discuss some of its very special properties. We conclude with a discussion section (Section 9) whose main purpose is to outline further work and point out connections with the work of others. The Appendix provides proofs of technical lemmas.

2 Stability in the ℓ_1 -norm

In this section, we establish ℓ_1 -properties of any minimizer to the problem (P_1) , when the initial signal is mostly concentrated (in an ℓ_1 sense) on a small set.

Lemma 2.1 *Assume that the measurement matrix F_Ω obeys **ERP**. We let f be a fixed signal of the form $f = f_0 + h$ where f_0 is a signal supported on a set T whose size obeys (1.16). Then with probability at least $1 - O(N^{-\rho/\alpha})$, any ℓ_1 -minimizer (1.10) obeys*

$$\|f^\sharp \cdot 1_{T^c}\|_{\ell_1} \leq 4 \|h\|_{\ell_1}. \quad (2.1)$$

Proof Observe that since f is of course feasible for (P_1) , we immediately have

$$\|f^\sharp\|_{\ell_1} \leq \|f\|_{\ell_1} \leq \|f_0\|_{\ell_1} + \|h\|_{\ell_1}. \quad (2.2)$$

Now because **ERP** holds, one can construct—with the required probability—a function $P = F_\Omega^* V$ for some $V \in \ell_2(K)$ such that $P = \text{sgn}(f_0)$ on T and $|P(t)| \leq 1/2$ away from T . Observe the identity

$$\langle f^\sharp, P \rangle = \langle f^\sharp, F_\Omega^* V \rangle = \langle F_\Omega f^\sharp, V \rangle = \langle F_\Omega(f_0 + h), V \rangle = \langle f_0 + h, F_\Omega^* V \rangle = \langle f_0 + h, P \rangle.$$

Then on the one hand

$$\langle f^\sharp, P \rangle = \langle f_0, P \rangle + \langle h, P \rangle \geq \|f_0\|_{\ell_1} - \|h\|_{\ell_1},$$

while on the other hand, the bounds on P give

$$\begin{aligned} |\langle f^\sharp, P \rangle| &\leq \sum_T |f^\sharp(t)P(t)| + \sum_{T^c} |f^\sharp(t)P(t)| \\ &\leq \sum_T |f^\sharp(t)| + \frac{1}{2} \sum_{T^c} |f^\sharp(t)| = \sum_{\mathbb{Z}_N} |f^\sharp(t)| - \frac{1}{2} \sum_{T^c} |f^\sharp(t)|. \end{aligned}$$

To conclude, we established that

$$\|f_0\|_{\ell_1} - \|h\|_{\ell_1} \leq \|f^\sharp\|_{\ell_1} - \frac{1}{2}\|f^\sharp \mathbf{1}_{T^c}\|_{\ell_1},$$

and together with (2.2) proved that

$$\|f^\sharp \mathbf{1}_{T^c}\|_{\ell_1} \leq 4\|h\|_{\ell_1},$$

as claimed. ■

This lemma says that any minimizer is approximately concentrated on the same set as the signal f . Indeed, suppose that f obeys (1.8) and consider T to be the set of largest values of $|f|$. Set $f_0 = f \cdot \mathbf{1}_T$. Then the property (1.8) gives

$$\|h\|_{\ell_1} = \|f \cdot \mathbf{1}_{T^c}\|_{\ell_1} \leq C_p \cdot |T|^{1-1/p}$$

for some constant C_p only depending on p , and therefore (2.1) gives

$$\|f^\sharp \cdot \mathbf{1}_{T^c}\|_{\ell_1} \leq 4C_p \cdot |T|^{1-1/p}. \quad (2.3)$$

Thus, f^\sharp puts ‘little mass’ outside of the set T .

Corollary 2.2 *Let f^\sharp be any ℓ_1 -minimizer to the problem (P_1) and rearrange the entries of f^\sharp in decreasing order of magnitude $|f^\sharp|_{(1)} \geq |f^\sharp|_{(2)} \geq \dots \geq |f^\sharp|_{(N)}$. Under the hypotheses of Lemma 2.1, the m th largest entry of f^\sharp obeys*

$$|f^\sharp|_{(m)} \leq C_p \cdot |T|^{-1/p}, \quad \forall m > 2|T|. \quad (2.4)$$

Proof Suppose T is the set of $|T|$ largest entries of f as above so that f^\sharp obeys (2.3). Denote by E_m the set of the m -largest values of the function f^\sharp . Obviously $|E_m \cap T^c| \geq m - |T|$ and, therefore,

$$\|f^\sharp\|_{\ell_1(E_m \cap T^c)} \geq (m - |T|) \cdot |f^\sharp|_{(m)} \geq |T| \cdot |f^\sharp|_{(m)}.$$

The claim then follows from

$$\|f^\sharp\|_{\ell_1(E_m \cap T^c)} \leq \|f^\sharp \mathbf{1}_{T^c}\|_{\ell_1} \leq C \cdot |T|^{1-1/p}. \quad \blacksquare$$

3 Stability in the ℓ_2 -norm

3.1 Extension lemma

As essentially observed in [11], a matrix obeying (1.17)— think of it as a partial Fourier transform— allows to extend a function from a small set to all of \mathbb{Z}_N while constraining its Fourier transform to a fixed random set:

Corollary 3.1 (Extension theorem) *Assume that F_Ω is a matrix obeying the uniform uncertainty principle UUP. Then with probability at least $1 - O(N^{-\rho/\alpha})$ the following statement holds: for all sets $T \subset \mathbb{Z}_N$ obeying the bound (1.17) and all functions $f \in \ell_2(T)$, there exists $f^{\text{ext}} \in \ell_2(\mathbb{Z}_N)$ which*

- agrees with f on T ($R_T f^{\text{ext}} = f$),
- belongs to the column space of F_Ω^* (i.e. $f^{\text{ext}} = F_\Omega^* V$ for some $V \in \ell_2(\Omega)$),
- and furthermore, we have the ℓ_2 estimates

$$\|f^{\text{ext}}\|_{\ell_2(E)} \leq C \left(1 + \frac{|E|}{\alpha K/\lambda}\right)^{1/2} \|f\|_{\ell_2(T)} \quad (3.1)$$

valid for all $E \subseteq \mathbb{Z}_N$.

Proof We may assume that we are on an event such that the conclusions of **UUP** hold. In particular, from (1.17), the operator $(F_{\Omega T}^* F_{\Omega T})$ is invertible and the inverse obeys

$$\|(F_{\Omega T}^* F_{\Omega T})^{-1}\| \leq 2N/K, \quad (3.2)$$

where $\|\cdot\|$ is the operator norm induced by the ℓ_2 norm. In the remainder of this paper and unless specified otherwise $\|A\|$ will always be the operator norm of $\|A\| := \sup_{\|x\|_{\ell_2}=1} \|Ax\|_{\ell_2}$. We now set f^{ext} as

$$f^{\text{ext}} := F_\Omega^* F_{\Omega T} (F_{\Omega T}^* F_{\Omega T})^{-1} f.$$

By construction, f^{ext} agrees with f on T , and is in the column space of F_Ω^* . Now we prove (3.1). It suffices to do so when $|E| \leq \alpha K/\lambda$, since the general claim then follows by decomposing larger E 's into smaller sets and then square-summing. But from (1.17), we see that $F_{\Omega T}$ and $F_{\Omega E}^*$ have operator norms of size at most $\sqrt{3K/2N}$, and the claim follows by composing these facts with (3.2). \blacksquare

3.2 Proof of Theorem 1.4

Let T_0 (resp. T_1) be the set of the S -largest values of $|f|$ (resp. $|f^\sharp|$) and put $T = T_0 \cup T_1$. By construction, $S \leq |T| \leq 2S$ and we assume that $|T|$ obeys the condition (1.16). Now observe that by construction of T , a consequence of Lemma 2.1 is that

$$\|f - f^\sharp\|_{\ell_1(T^c)} \leq \|f\|_{\ell_1(T_0^c)} + \|f^\sharp\|_{\ell_1(T_1^c)} \leq C_p \cdot |T|^{1-1/p}. \quad (3.3)$$

Furthermore, it follows from our assumption about f and (2.4) that

$$\|f - f^\sharp\|_{\ell_\infty(T^c)} \leq \|f\|_{\ell_\infty(T_0^c)} + \|f^\sharp\|_{\ell_\infty(T_1^c)} \leq C \cdot |T|^{-1/p}. \quad (3.4)$$

By interpolation, these last two inequalities give

$$\|f - f^\sharp\|_{\ell_2(T^c)} \leq C \cdot |T|^{1/2-1/p}, \quad (3.5)$$

and it remains to prove that the same bound holds over the set T .

In order to prove this fact, Corollary 3.1 assures us that one can find a function of the form $g = F_\Omega^* V$ which matches h on T and with the following property:

$$\sum_{t \in E} |g(t)|^2 \leq C \sum_{t \in T} |f(t) - f^\sharp(t)|^2, \quad (3.6)$$

for all sets E of cardinality $O(K/\lambda)$ that are disjoint from T . Here and in the rest of the proof, the constants C are allowed to depend on α . From the representation $g = F_\Omega^*V$ and the constraint $F_\Omega f = F_\Omega f^\sharp$ (from (1.10)), we have

$$\langle f - f^\sharp, g \rangle = \langle f - f^\sharp, F_\Omega^*V \rangle = \langle F_\Omega f - F_\Omega f^\sharp, V \rangle = 0,$$

and hence

$$\sum_{t \in \mathbb{Z}_N} (f - f^\sharp)(t) \overline{g(t)} = 0.$$

Splitting into T and T^c , we obtain

$$\sum_{t \in T} |f - f^\sharp|^2(t) = - \sum_{t \in T^c} (f - f^\sharp)(t) \overline{g(t)}. \quad (3.7)$$

We will use (3.7) to show that the left-hand side must be small since (3.5) and (3.6) assert that the right-hand side is not very large. Enumerate T^c as $n_1, n_2, \dots, n_{N-|T|}$ in decreasing order of magnitude of $|f - f^\sharp|$. We then group these into adjacent blocks B_J of size $|T|$ (except perhaps for the last one) $B_J := \{n_j, J|T| < j \leq (J+1)|T|\}$, $J = 0, 1, \dots$. From (3.6) and Cauchy-Schwarz, we have

$$\left| \sum_{j \in B_J} (f - f^\sharp)(n_j) \overline{g(n_j)} \right| \leq C \cdot \|f - f^\sharp\|_{\ell_2(T)} \cdot I_J, \quad (3.8)$$

where

$$I_J := \sqrt{\sum_{j=J|T|+1}^{(J+1)|T|} |(f - f^\sharp)(n_j)|^2}.$$

Because we are enumerating the values of n_j in decreasing order, we have $I_0 \leq |T|^{1/2} \cdot |f - f^\sharp|(n_1) \leq C \cdot |T|^{1/2-1/p}$ while for $J \geq 1$ we have

$$I_J \leq |T|^{1/2} \cdot |f - f^\sharp|(n_{J|T|+1}) \leq |T|^{1/2} \cdot |T|^{-1} \cdot \|f - f^\sharp\|_{\ell_1(B_{J-1})},$$

In other words,

$$\sum_J I_J \leq I_0 + \sum_{J \geq 1} I_J \leq C \cdot |T|^{-r} + |T|^{-1/2} \cdot \|f - f^\sharp\|_{\ell_1(T^c)}$$

and, therefore, it follows from (3.3) that the summation of the inequality (3.8) over the blocks B_J gives

$$\left| \sum_{t \in T^c} (f - f^\sharp)(t) \overline{g(t)} \right| \leq C \cdot |T|^{-r} \cdot \|f - f^\sharp\|_{\ell_2(T)}.$$

Inserting this back into (3.7), we established

$$\|f - f^\sharp\|_{\ell_2(T)} \leq C \cdot |T|^{-r}.$$

This concludes the proof of Theorem 1.4. \blacksquare

Note that by Cauchy-Schwarz, it follows from the proof of our Theorem that

$$\|f - f^\sharp\|_{\ell_1(T)} \leq C \cdot |T|^{1-1/p},$$

and, therefore, owing to (3.3), we also proved an ℓ_1 stability estimate

$$\|f - f^\sharp\|_{\ell_1} \leq C \cdot |T|^{1-1/p}. \quad (3.9)$$

Had we assumed that f belonged to the weak- ℓ_1 ball when $p = 1$, the right-hand side of (3.3) would read $C_1 \log(N/|T|)$ instead of just C_1 . This is the reason why we required ℓ_1 in the hypothesis of Theorem 1.4 and showed that we also have a near-optimal signal recovery result for the unit ball of ℓ_1 with no additional losses (logarithmic or otherwise).

3.3 Uniqueness of the minimizer for the Gaussian ensemble

The claim that the minimizer f^\sharp is unique with probability 1, for Gaussian measurements, can be easily established as follows. The claim is trivial for $f \equiv 0$ so we may assume f is not identically zero. Then $F_\Omega f$ is almost surely non-zero. Furthermore, if one considers each of the (finitely many) facets of the unit ball of $\ell_1(\mathbb{Z}_N)$, we see that with probability 1 the random Gaussian matrix F_Ω has maximal rank on each of these facets (i.e. the image of each facet under F_Ω has dimension equal to either K or the dimension of the facet, whichever is smaller). From this we see that every point on the boundary of the image of the unit ℓ_1 -ball under F_Ω arises from a unique point on that ball. Similarly for non-zero dilates of this ball. Thus the solution to the problem (1.10) is unique as claimed.

We remark that the question of establishing uniqueness with high probability for discretely randomized ensembles such as the binary and Fourier ensembles discussed below is an interesting one, but one which we will not pursue here.

4 Eigenvalues of random matrices

In this section, we show that all three ensembles obey the uniform uncertainty principle **UUP**.

4.1 The Gaussian ensemble

Let X be an n by p matrix with $p \leq n$ and with i.i.d. entries sampled from the normal distribution with mean zero and variance $1/n$. We are interested in the singular values of X or the eigenvalues of X^*X . A famous result due to Marchenko and Pastur [42] states that the eigenvalues of X^*X have a deterministic limit distribution supported by the interval $[(1 - \sqrt{c})^2, (1 + \sqrt{c})^2]$ as $n, p \rightarrow \infty$, with $p/n \rightarrow c < 1$. In fact, results from [51] show that the smallest (resp. largest) eigenvalue converges a.s. to $(1 - \sqrt{c})^2$ (resp. $(1 + \sqrt{c})^2$). In other words, the smallest singular value of X/\sqrt{n} converges a.s. to $1 - \sqrt{c}$ and the largest to $1 + \sqrt{c}$. In addition, there are remarkably fine statements concerning the speed of the convergence of the largest singular value [36].

To derive the **UUP**, we need a result about the concentration of the extreme singular values of a Gaussian matrix, and we borrow a most elegant estimate due to Davidson and Szarek [52]. We let $\lambda_1(X) \leq \dots \leq \lambda_p(X)$ be the ordered list of the singular values of X . Then in

[52], the authors prove that

$$\mathbf{P}\left(\lambda_p(X) > 1 + \sqrt{p/n} + r\right) \leq e^{-nr^2/2}, \quad (4.1)$$

$$\mathbf{P}\left(\lambda_1(X) < 1 - \sqrt{p/n} - r\right) \leq e^{-nr^2/2}. \quad (4.2)$$

Such inequalities about the concentration of the largest and smallest singular values of Gaussian matrices have been known for at least a decade or so. Estimates similar to (4.1)-(4.2) may be found in the work of Szarek [19], see also Ledoux [39].

Lemma 4.1 *The Gaussian ensemble obeys the uniform uncertainty principle (UUP) with oversampling factor $\lambda = \log N$.*

Proof Fix $K \geq \log N$ and let $\Omega := \{1, \dots, K\}$. Let T be a fixed index set and define the event E_T as

$$E_T := \{\lambda_{\min}(F_{\Omega T}^* F_{\Omega T}) < K/2N\} \cup \{\lambda_{\max}(F_{\Omega T}^* F_{\Omega T}) > 3K/2N\}.$$

Since the entries of $F_{\Omega T}$ are i.i.d. $N(0, 1/N)$, it follows from (4.1)-(4.2) by a simple renormalization that for each $|T| \leq K/16$,

$$\mathbf{P}(E_T) \leq 2e^{-cK}$$

where one can choose $c = 1/32$ by selecting $r = 1/4$ in (4.1)-(4.2). We now examine the tightness of the spectrum over all sets $T \in \mathcal{T}_m := \{T : |T| \leq m\}$ where we assume that m is less than $N/2$. We have

$$\mathbf{P}(\cup_{\mathcal{T}_m} E_T) \leq 2e^{-cK} \cdot |\mathcal{T}_m| = 2e^{-cK} \cdot \sum_{k=1}^m \binom{N}{k} \leq 2e^{-cK} \cdot m \binom{N}{m}.$$

We now use the well-known bound on the binomial coefficient

$$\binom{N}{m} \leq e^{NH(m/N)},$$

where for $0 < q < 1$, H is the binary entropy function

$$H(q) := -q \log q - (1 - q) \log(1 - q).$$

The inequality $-(1 - q) \log(1 - q) \leq q$ shows that $-(1 - m/N) \log(1 - m/N) \leq m/N$ and thus

$$m \cdot \binom{N}{m} = e^{m \log(N/m) + m + \log m}.$$

Whence,

$$\log \mathbf{P}(\cup_{\mathcal{T}_m} E_T) \leq \log 2 - cK + m(\log(N/m) + 1 + m^{-1} \log m) \leq \log 2 - \rho K$$

provided that $m(\log(N/m) + 1 + m^{-1} \log m) \leq (c - \rho)K$, which is what we needed to establish. (We need to assume that $K \geq (c - \rho)^{-1}(1 + \log N)$ for the claim not to be vacuous.) Note that we proved more than what we claimed since the **UUP** holds for an oversampling factor proportional to $\log N/K$. \blacksquare

4.2 The binary ensemble

The analysis is more complicated in the case where the matrix X is an n by p array with i.i.d. symmetric Bernoulli entries taking on values in $\{-1/\sqrt{n}, 1/\sqrt{n}\}$. To study the concentration of the largest singular values of X , we follow an approach proposed by Ledoux [39] which makes a simple use of the concentration property, see also [28].

As before, we let $\lambda_p(X)$ be the mapping that associates to a matrix X its largest singular values. Equip \mathbb{R}^{np} with the Frobenius norm

$$\|X\|_F^2 := \text{Tr}(X^*X) = \sum_{i,j=1}^p |X_{ij}|^2$$

(the Euclidean norm over \mathbb{R}^{np}). Then the mapping λ_p is convex and 1-Lipschitz in the sense that

$$|\lambda_p(X) - \lambda_p(X')| \leq \|X - X'\|_F$$

for all pairs (X, X') of n by p matrices. A classical application of the concentration inequality for binary measures [39] then gives

$$\mathbf{P}(\lambda_p(X) - m(\lambda_p(X)) \geq r) \leq e^{-nr^2/16}, \quad (4.3)$$

$m(\lambda_p(X))$ is either the mean or the median of $\lambda_p(X)$. Now the singular values still exhibit the same behavior; that is $\lambda_{\min}(X/\sqrt{n})$ and $\lambda_{\max}(X/\sqrt{n})$ converge a.s. to $1 - \sqrt{c}$ and $1 + \sqrt{c}$ respectively, as $n, p \rightarrow \infty$ with $p/n \rightarrow c$ [2]. As a consequence, for each ϵ_0 and n sufficiently large, one can show that the medians belong to the fixed interval $[1 - \sqrt{p/n} - \epsilon_0, 1 + \sqrt{p/n} + \epsilon_0]$ which gives

$$\mathbf{P}(\lambda_p(X) > 1 + \sqrt{p/n} + \epsilon_0 + r) \leq e^{-nr^2/16}. \quad (4.4)$$

This is a fairly well-established result [28].

The problem is that this method does not apply to the minimum singular value which is 1-Lipshitz but not convex. Fortunately, Litvak, Pajor, Rudelson and Tomczak-Jaegermann [40][Theorem 3.1] have recently announced a result which gives exponential concentration for the lowest singular value. They proved that whenever $n \geq (1 + \delta)p$ where δ is greater than a small constant,

$$\mathbf{P}(\lambda_1(X) < c_1) \leq e^{-c_2 n}, \quad (4.5)$$

where c_1 and c_2 are universal positive constants.

Just as (4.1)-(4.2) implied the uniform uncertainty principle **UUP** for Gaussian matrices, (4.4)-(4.5) gives the same conclusion for the binary ensemble with the proviso that the condition about the lowest singular value reads $\lambda_{\min}(F_{\Omega T}^* F_{\Omega T}) > c_1 K/N$; i.e., c_1 substitutes $1/2$ (recall the remark following the definition of the **UUP**).

Lemma 4.2 *The binary ensemble obeys the uniform uncertainty principle (**UUP**) with oversampling factor $\lambda = \log N$.*

The proof is of course identical to that of Lemma 4.1. If we define E_T as

$$E_T := \{\lambda_{\min}(F_{\Omega T}^* F_{\Omega T}) < c_1 K/N\} \cup \{\lambda_{\max}(F_{\Omega T}^* F_{\Omega T}) > 3K/2N\},$$

we have $\mathbf{P}(E_T) \leq 2e^{-cK}$ for some constant $c > 0$. The rest of the proof is as before.

4.3 The Fourier ensemble

The analysis for the Fourier ensemble is much more delicate than that for the Gaussian and binary cases, in particular requiring entropy arguments as used for instance by Bourgain [6], [7]. We prove the following lemma in the separate Section 7.

Lemma 4.3 *The Fourier ensemble obeys the uniform uncertainty principle **UUP** with oversampling factor $\lambda = (\log N)^6$.*

The exponent of 6 can almost certainly be lowered⁵, but we will not attempt to seek the optimal exponent of $\log N$ here.

5 Generic signals and the weak ERP

In some cases, it might be difficult to prove that the exact reconstruction principle **ERP** holds, and it is interesting to observe that **UUP** actually implies **ERP** for ‘generic’ sign functions $\sigma = \pm 1$ supported on a small set T . More precisely, if we fix T and define σ to be supported on T with the i.i.d. Bernoulli distribution (independently of F_Ω), thus

$$\mathbf{P}(\sigma(t) = \pm 1) = 1/2, \quad \text{for all } t \in T.$$

then we shall construct a P obeying the conditions (i)-(iii) in the definition of **ERP**. Indeed, we shall construct P explicitly as

$$P := F_\Omega^* F_{\Omega T} (F_{\Omega T}^* F_{\Omega T})^{-1} R_T \sigma; \tag{5.1}$$

one can view this choice of $P = F_\Omega^* V$ as the unique solution to (i) and (ii) which minimizes the ℓ_2 norm of V , and can thus be viewed as a kind of least-squares extension of σ using the rows of F_Ω .

It is immediate to check that P obeys (i) and (ii) above. Indeed, the restriction of P to T is given by

$$R_T P = R_T F_\Omega^* F_{\Omega T} (F_{\Omega T}^* F_{\Omega T})^{-1} R_T \sigma = F_{\Omega T}^* F_{\Omega T} (F_{\Omega T}^* F_{\Omega T})^{-1} R_T \sigma = R_T \sigma$$

and, therefore, (i) is verified. Further, it follows from the definition that P is a linear combination of the columns of F_Ω^* and thus, (ii) holds. Therefore, we only need to check that for all $t \in T^c$, $|P(t)| \leq \frac{1}{2}$ with sufficiently high probability. In order to do this, we rewrite $P(t)$ as

$$P(t) = \langle W_t, R_T \sigma \rangle,$$

where for each $t \in T^c$, W_t is the $|T|$ dimensional vector

$$W_t := (F_{\Omega T}^* F_{\Omega T})^{-1} F_{\Omega T}^* F_t$$

and F_t is the t -th column of F_Ω . We now introduce another condition which is far easier to check than **ERP**.

⁵Note added in proof: since the submission of this paper, Rudelson and Vershynin, in a very recent piece of work [47], have improved the oversampling factor to $(\log N)^4$.

WERP (Weak ERP). We say that the measurement process obeys the weak **ERP**, if for each fixed T obeying (1.16) and any $0 < \gamma \leq 1$, F_Ω obeys

$$\|F_{\Omega T}^* F_t\|_{\ell_2} \leq \gamma \cdot \sqrt{K |T|/N} \quad \text{for all } t \in T^c \quad (5.2)$$

with probability at least $1 - O(N^{-\rho/\gamma})$ for some fixed positive constant $\rho > 0$.

For example, it is an easy exercise in large deviation theory to show that **WERP** holds for Gaussian and binary measurements. One can also check that **WERP** holds for random frequency samples. We omit the proof of these facts, however, since we will show the stronger version, namely, **ERP** in all three cases. Instead, we would like to emphasize that **UUP** together with **WERP** actually imply **ERP** for most sign patterns σ .

We begin by recalling the classical Hoeffding inequality: let $X_1, \dots, X_N = \pm 1$ be independent symmetric Bernoulli random variables and consider the sum $S = \sum_{j=1}^N a_j X_j$. Then

$$P(|S| > \lambda) \leq 2 \exp\left(-\frac{\lambda^2}{2\|a\|_{\ell_2}^2}\right). \quad (5.3)$$

Suppose now that the $\sigma(t)$'s are independent Bernoulli, and independent from F_Ω . Then (5.3) gives

$$\mathbf{P}(|P(t)| > \lambda \mid \|W_t\|_{\ell_2} = \rho) \leq 2 \exp\left(-\frac{\lambda^2}{2\rho^2}\right).$$

If we now assume that both **UUP** and **WERP** hold, then for any $0 < \gamma \leq 1$ we have

$$\|W_t\|_{\ell_2} \leq \|(F_{\Omega T}^* F_{\Omega T})^{-1}\| \cdot \|F_{\Omega T}^* F_t\|_{\ell_2} \leq 2\gamma \cdot \sqrt{\frac{|T|}{K}}.$$

with probability at least $1 - O(N^{-\rho/\gamma})$. This shows that

$$\begin{aligned} \mathbf{P}(|P(t)| > \lambda) &\leq \mathbf{P}\left(|P(t)| > \lambda \mid \|W_t\|_{\ell_2} \leq 2\gamma \cdot \sqrt{|T|/K}\right) + \mathbf{P}\left(\|W_t\|_{\ell_2} > 2\gamma \cdot \sqrt{|T|/K}\right) \\ &\leq 2 \exp\left(-\frac{\lambda^2}{8\gamma^2} \frac{K}{|T|}\right) + O(N^{-\rho/\gamma}) = 2e^{-\rho_0 K/|T|} + O(N^{-\rho/\gamma}). \end{aligned}$$

Hence, if $|T| \leq \alpha K / \log N$, then

$$\mathbf{P}\left(\sup_{t \in T^c} |P(t)| > 1/2\right) \leq O(N \cdot N^{-\rho_0/\alpha}) + O(N^{-\rho/\alpha}).$$

Therefore, if α is chosen small enough, then for some small $\rho' > 0$

$$\mathbf{P}\left(\sup_{t \in T^c} |P(t)| > 1/2\right) = O(N^{-\rho'/\alpha}).$$

In other words, **ERP** holds for most sign patterns σ . That is, if one is only interested in providing good reconstruction to nearly all signals (but not all) in the sense discussed above, then it is actually sufficient to check that both conditions **UUP** and **WERP** are valid.

6 About the exact reconstruction principle

In this section, we show that all the three ensembles obey the exact reconstruction principle **ERP**.

6.1 The Gaussian ensemble

To show that there is function P obeying the conditions (i)-(iii) in the definition of **ERP**, we take an approach that resembles that of Section 5, and establish that P defined as in (5.1),

$$P := F_{\Omega}^* F_{\Omega T} (F_{\Omega T}^* F_{\Omega T})^{-1} R_T \sigma$$

obeys the three conditions (i)-(iii).

We already argued that P obeys (i) and (ii). Put $P^c = R_{T^c} P$ to be the restriction of P to T^c . We need to show that

$$\sup_{T^c} |P^c(t)| \leq 1/2$$

with high probability. Begin by factorizing P^c as

$$P^c = F_{\Omega T^c}^* V, \text{ where } V := F_{\Omega T} (F_{\Omega T}^* F_{\Omega T})^{-1} R_T \sigma.$$

The crucial observation is that the random matrix $F_{\Omega T^c}^*$ and the random variable V are independent since they are functions of disjoint sets of independent variables.

Proposition 6.1 *Conditional on V , the components of P^c are i.i.d. Gaussian with*

$$P^c(t) \sim N(0, \|V\|^2/N).$$

Proof Suppose V is fixed. By definition,

$$P^c(t) = \frac{1}{\sqrt{N}} \sum_{k \in \Omega} X_{k,t} V_k$$

and, therefore, it follows from the independence between the $X_{k,t}$'s and V for each $t \in T^c$ that the conditional distribution of $P^c(t)$ is normal with mean 0 and variance $\|V\|_{\ell_2}^2/N$. The independence between the components of P^c is a simple consequence of the independence between the columns of F . \blacksquare

Lemma 6.2 *Let $\alpha > 0$ be sufficiently small, and suppose that $|T|$ is chosen as in (1.16) so that **UUP** holds. The components of $P^c(t)$ obey*

$$\mathbf{P}(|P^c(t)| > \lambda) \leq \mathbf{P}\left(|Z| > \lambda \cdot \sqrt{K/6|T|}\right) + O(N^{-\rho/\alpha}).$$

where $Z \sim N(0, 1)$ is a standard normal random variable.

Proof Observe that

$$\|V\|_{\ell_2} \leq \|F_{\Omega T}\| \cdot \|(F_{\Omega T}^* F_{\Omega T})^{-1}\| \cdot \|R_T \sigma\|_{\ell_2}.$$

On the event such that the conclusions of **UUP** hold, $\|F_{\Omega T}\| \leq \sqrt{3K/2N}$ and also $\|(F_{\Omega T}^* F_{\Omega T})^{-1}\| \leq 2N/K$. Since $\|R_T \sigma\|_{\ell_2} = \sqrt{|T|}$, this gives

$$\|V\|_{\ell_2} \leq \sqrt{\frac{6N \cdot |T|}{K}}.$$

Therefore,

$$\mathbf{P}(|P^c(t)| > \lambda) \leq \mathbf{P}\left(|P^c(t)| > \lambda \mid \|V\|_{\ell_2} \leq \sqrt{6N|T|/K}\right) + \mathbf{P}\left(\|V\|_{\ell_2} > \sqrt{6N|T|/K}\right).$$

The first term is bounded by Proposition 6.1 while the second is bounded via the uniform uncertainty principle **UUP**. This establishes the lemma. \blacksquare

The previous lemma showed that for $|T| \leq \alpha \cdot K/\log N$,

$$\mathbf{P}\left(\sup_{t \in T^c} |P^c(t)| > 1/2\right) \leq N \cdot \mathbf{P}\left(|Z| > \frac{1}{2} \sqrt{\frac{\log N}{6\alpha}}\right) + O(N^{-\rho/\alpha}) \leq 2N^{1-1/(48\alpha)} + O(N^{-\rho/\alpha}).$$

Therefore, if α is chosen small enough, then for some small $\rho' > 0$

$$\mathbf{P}\left(\sup_{t \in T^c} |P^c(t)| > 1/2\right) = O(N^{-\rho'/\alpha}).$$

In short, we proved:

Lemma 6.3 *The Gaussian ensemble obeys the exact reconstruction principle **ERP** with oversampling factor $\lambda = \log N$.*

6.2 The binary ensemble

The strategy in the case where the entries of F are independent Bernoulli variables is nearly identical and we only discuss the main differences. Define P and V as above; obviously, $F_{\Omega T^c}^*$ and V are still independent.

Proposition 6.4 *Conditional on V , the components of P^c are independent and obey*

$$\mathbf{P}(|P^c(t)| > \lambda \mid V) \leq 2 \exp\left(-\frac{\lambda^2 N}{2\|V\|_{\ell_2}^2}\right).$$

Proof The conditional independence of the components is as before. As far as the tail-bound is concerned, we observe that $P^c(t)$ is a weighted sum of independent Bernoulli variables and the claim follows from the Hoeffding inequality (5.3). \blacksquare

The rest of the argument is as before. If $|T|$ is selected as in (1.16) such that **UUP** holds, one has

$$\mathbf{P}(|P^c(t)| > 1/2) \leq 2N^{-1/48\alpha} + O(N^{-\rho/\alpha}).$$

And, of course, identical calculations now give

Lemma 6.5 *The binary ensemble obeys the exact reconstruction principle **ERP** with oversampling factor $\lambda = \log N$.*

6.3 The Fourier ensemble

It turns out that the exact reconstruction principle also holds for the Fourier ensemble although the argument is considerably more involved [11]. We do not reproduce the proof here but merely indicate the strategy for proving that P (defined as before) also obeys the desired bound on the complement of T with sufficiently high probability. We first remark that $|\Omega|$ is concentrated around K . To see this, recall the Bernstein's inequality [5] which states that if X_1, \dots, X_m are independent random variables with mean-zero and obeying $|X_i| \leq c$, then

$$\mathbf{P} \left(\left| \sum_{i=1}^N X_i \right| > \lambda \right) \leq 2 \exp \left(-\frac{\lambda^2}{2\sigma^2 + 2c\lambda/3} \right), \quad (6.1)$$

where $\sigma^2 = \sum_{i=1}^m \text{Var}(X_i)$. Specializing this inequality gives the following lemma which we shall need later in this paper.

Lemma 6.6 *Fix $\tau \in (0, 1)$ and let $I_k \in \{0, 1\}$ be an i.i.d. sequence of random variables obeying $\mathbf{P}(I_k = 1) = \tau$. Let $a \in \ell_2(\mathbf{Z}_N)$ be arbitrary, and set $\sigma^2 := \tau(1 - \tau)\|a\|_{\ell_2}^2$. Then*

$$\mathbf{P} \left(\left| \sum_{k \in \mathbf{Z}_N} (I_k - \tau)a(k) \right| > \lambda \right) \leq 4 \exp \left(-\frac{\lambda^2}{4(\sigma^2 + \lambda\|a\|_{\infty}/3\sqrt{2})} \right).$$

Proof Letting S be the sum $\sum_{k \in \mathbf{Z}_N} (I_k - \tau)a(k)$, the proof follows from (6.1) by simply stating that $P(|S| > \lambda)$ is bounded above by the sum $P(|S_1| > \lambda/\sqrt{2}) + P(|S_2| > \lambda/\sqrt{2})$, where S_1 and S_2 are the real and imaginary parts of S respectively. ■

Thus the bound on the quantity $|\sum_{k \in \mathbf{Z}_N} (I_k - \tau)a(k)|$ exhibits a Gaussian-type behavior at small thresholds λ , and an exponential-type behavior at large thresholds.

Recall that $K = \mathbf{E}(|\Omega|) = N\tau$. Applying Lemma 6.6 with $a \equiv 1$ (so $\sigma^2 = N\tau(1 - \tau)$), we have that $P(K/2 \leq |\Omega| \leq 2K)$ with probability $O(N^{-\rho/\alpha})$ provided that $K \geq C\alpha^{-1} \log N$, which we will assume as the claim is vacuous otherwise. In the sequel, we assume that we are on an event $\{K/2 \leq |\Omega| \leq 2K\}$.

Decompose $F_{\Omega T}^* F_{\Omega T}$ as

$$F_{\Omega T}^* F_{\Omega T} = \frac{|\Omega|}{N} (I - H),$$

where H is the matrix defined by $H(t, t') = |\Omega|^{-1} \sum_{k \in \Omega} e^{i2\pi k(t-t')}$ if $t \neq t'$ and 0 otherwise. We then expand the inverse as a truncated Neumann series

$$(F_{\Omega T}^* F_{\Omega T})^{-1} = \frac{N}{|\Omega|} (I + H + \dots + H^n + E),$$

where E is small remainder term. This allows to express P^c as

$$P^c = \frac{N}{|\Omega|} \cdot F_{\Omega T^c}^* F_{\Omega T} \cdot (I + H + \dots + H^n + E),$$

and one can derive bounds on each individual terms so that the sum obeys the desired property. By pursuing this strategy, the following claim was proved in [11].

Lemma 6.7 *The Fourier ensemble obeys the exact reconstruction principle **ERP** with oversampling factor $\lambda = \log N$.*

7 Uniform Uncertainty Principles for the Fourier Ensemble

In this section, we prove Lemma 4.3. The ideas here are inspired by an entropy argument sketched in [7], as well as by related arguments in [6], [43]. These methods have since become standard in the high-dimensional geometry literature, but we shall give a mostly self-contained presentation here.

We remark that the arguments in this section (and those in the Appendix) do not use any algebraic properties of the Fourier transform other than the Plancherel identity and the fact that the maximum entry of the Fourier matrix is bounded by $1/\sqrt{N}$. Indeed a simple modification of the arguments we give below also gives the UUP for randomly sampled rows of orthonormal matrices, see also [6] and [13] for further discussion of this issue. Suppose that $\sup_{i,j} |U_{ij}| \leq \mu$ and let U_Ω be the matrix obtained by randomly selected rows. Then the UUP holds for

$$|T| \leq C \cdot \frac{\mathbf{E}|\Omega|}{\mu^2 \log^6 N}.$$

In the case where one observes a few coefficients in the basis Φ when the signal is sparse in another basis Ψ , $\mu = \sqrt{N} \sup_{i,j} |\langle \phi_i, \psi_j \rangle|$ is interpreted as the mutual coherence between Φ and Ψ [27].

For sake of concreteness, we now return to the Fourier ensemble. Let us first set up what we are trying to prove. Fix $\alpha > 0$, which we shall assume to be sufficiently small. We may take N to be large depending on α , as the claim is vacuous when N is bounded depending on α . If T is empty then the claim is trivial, so from (1.16) we may assume that

$$K = \tau N \geq C \log^6 N \tag{7.1}$$

for some (possibly) large constant C .

We need to prove (1.17). By self-adjointness, it would suffice to show that with probability at least $1 - O(N^{-\rho/\alpha})$

$$\left| \langle F_{\Omega T}^* F_{\Omega T} f, f \rangle_{\ell_2(T)} - \tau \|f\|_{\ell_2(T)}^2 \right| \leq \frac{1}{4} \tau \cdot \|f\|_{\ell_2(T)}^2$$

for all $f \in \ell_2(T)$ and all T obeying (1.16), thus $|T| \leq m$, where

$$m := \alpha \tau N / \log^6 N. \tag{7.2}$$

For any fixed T and f , the above type of estimate can easily be established with high probability by standard tools such as Lemma 6.6. The main difficulty is that there are an exponentially large number of possible T to consider, and for each fixed T there is a $|T|$ -dimensional family of f to consider. The strategy is to cover the set of all f of interest by various finite nets at several scales, obtain good bounds on the size of such nets, obtain large deviation estimates for the contribution caused by passing from one net to the net at the next scale, and sum using the union bound.

We turn to the details. We can rewrite our goal as

$$\left| \sum_{k \in \Omega} |\hat{f}(k)|^2 - \tau \|f\|_{\ell_2(T)}^2 \right| \leq \frac{1}{4} \tau \cdot \|f\|_{\ell_2(T)}^2$$

whenever $|T| \leq m$. From Parseval's identity, this is the same as asking that

$$\left| \sum_{k \in \mathbf{Z}_N} (I_k - \tau) |\hat{f}(k)|^2 \right| \leq \frac{1}{4} \tau \cdot \|f\|_{\ell_2(T)}^2$$

whenever $|T| \leq m$. Now let $U_m \subseteq \ell_2(\mathbf{Z}_N)$ denote the set

$$U_m := \bigcup \{B_{\ell_2(E)} : E \subseteq \mathbf{Z}_N, |E| = m\} = \{f \in \ell_2(\mathbf{Z}_N) : \|f\|_{\ell_2(\mathbf{Z}_N)} \leq 1, |\text{supp}(f)| \leq m\}.$$

Then the previous goal is equivalent to showing that

$$\sup_{f \in U_m} \left| \sum_{k \in \mathbf{Z}_N} (I_k - \tau) |\hat{f}(k)|^2 \right| \leq \frac{1}{4} \tau$$

with probability at least $1 - O(N^{-\rho/\alpha})$ for some $\rho > 0$. In fact we shall obtain the stronger estimate

$$\mathbf{P} \left(\sup_{f \in U_m} \left| \sum_{k \in \mathbf{Z}_N} (I_k - \tau) |\hat{f}(k)|^2 \right| > \frac{1}{4} \tau \right) = O \left(\exp \left(-\frac{1}{\beta} \log^2 N \right) \right) \quad (7.3)$$

for some constant $\beta > 0$.

It remains to prove (7.3). The left-hand side of (7.3) is the large deviation probability of a supremum of random sums over U_m . This type of expression can be handled by entropy estimates on U_m , as was done in [6], [7], [43]. To follow their approach, we need some notation. For any $f \in \ell_2(\mathbf{Z}_N)$, we let \hat{f} be its discrete Fourier transform (1.3) and define the X norm of f by

$$\|f\|_X := \sqrt{N} \cdot \|\hat{f}\|_{\ell_\infty}.$$

Intuitively, if f is a “generic” function bounded in $\ell_2(\mathbf{Z}_N)$ we expect the X norm of f to be also bounded (by standard large deviation estimates). We shall need this type of control in order to apply Lemma 6.6 effectively. To formalize this intuition we shall need entropy estimates on U_m in the X norm. Let B_X be the unit ball of X in $\ell_2(\mathbf{Z}_N)$. Thus for instance U_m is contained inside the ball $\sqrt{m} \cdot B_X$, thanks to Cauchy-Schwarz. However we have much better entropy estimates available on U_m in the X norm, which we now state.

Definition 7.1 (Kolmogorov entropy) *Let X be a (finite-dimensional) normed vector space with norm $\|\cdot\|_X$, and let $B_X := \{x \in X : \|x\|_X < 1\}$ be the unit ball of X . If $U \subset X$ is any bounded non-empty subset of X and $r > 0$, we define the covering number $N(U, B_X, r) \in \mathbf{Z}^+$ to be the least integer N such that there exist elements $x_1, \dots, x_N \in X$ such that the balls $x_j + rB_X = \{x \in X : \|x - x_j\|_X < r\}$, $1 \leq j \leq N$ cover U , and the Kolmogorov entropy as*

$$\mathcal{E}(U, B_X, r) := \log_2(N(U, B_X, r)).$$

Proposition 7.2 *We have*

$$\mathcal{E}(U_m, B_X, r) \leq C \cdot m \log N \cdot \min(r^{-2} \log N, 1) \quad (7.4)$$

for all $r > N^{-2}$.

This proposition is essentially contained in [6], [7], [43]; for sake of completeness we give a proof of the proposition in the Appendix. Let us assume this proposition for the moment and conclude the proof of Lemma 4.3. Set

$$\{J_0, \dots, J_1\} = \{j \in \mathbf{Z} : N^2 \leq 2^j \leq \sqrt{m}\}.$$

and fix $r = 2^j$ in Lemma 7.2. By (7.4) one can find a finite subset A_j of U_m of cardinality

$$|A_j| \leq \exp\left(\frac{C}{1+2^{2j}} \cdot m \log^2 N\right), \quad (7.5)$$

such that for all $f \in U_m$, there exists $f_j \in A_j$ such that $\|f - f_j\|_X \leq 2^j$. Let us fix such sets A_j . Then for any $f \in U_m$, we have the telescoping decomposition

$$f = f_{-\infty} + \sum_{J_0 \leq j \leq J_1} f_{j+1} - f_j.$$

where $f_j \in A_j$ and $f_{-\infty} = f - f_{J_0}$; here we have the convention that $f_j = 0$ and $A_j = \{0\}$ if $j > J_1$. By construction, $\|f_j - f_{j+1}\|_X \leq 2^{j+2}$, and $\|f_{-\infty}\|_X \leq 2N^{-2}$. We write $g_j := f_{j+1} - f_j$, thus $\|g_j\|_X \leq 2^{j+2}$. Fix k and observe the crude estimates

$$|\hat{f}(k)| \leq \|\hat{f}\|_{\ell_2} = \|f\|_{\ell_2} = 1 \quad \text{and} \quad |\hat{f}_{-\infty}(k)| \leq \|f_{-\infty}\|_X / \sqrt{N} \leq 2N^{-5/2}.$$

It then follows from $|a + b|^2 \leq |a|^2 + |b|^2 + 2|a||b|$ that

$$|\hat{f}(k)|^2 = \left| \sum_{J_0 \leq j \leq J_1} \hat{g}_j(k) \right|^2 + O(N^{-5/2}).$$

Multiplying this by $I_k - \tau$ and summing, we obtain

$$\begin{aligned} \left| \sum_{k \in \mathbf{Z}_N} (I_k - \tau) |\hat{f}(k)|^2 \right| &= \left| \sum_{k \in \mathbf{Z}_N} (I_k - \tau) \left| \sum_{J_0 \leq j \leq J_1} \hat{g}_j(k) \right|^2 \right| + O(N^{-3/2}) \\ &\leq 2 \sum_{J_0 \leq j \leq j' \leq J_1} Q(g_j, g_{j'}) + O(N^{-3/2}). \end{aligned}$$

where $Q(g_j, g_{j'})$ is the nonnegative random variable

$$Q(g_j, g_{j'}) := \left| \sum_{k \in \mathbf{Z}_N} (I_k - \tau) \operatorname{Re}(\hat{g}_j(k) \overline{\hat{g}_{j'}(k)}) \right|.$$

By (7.1), the error term $O(N^{-3/2})$ is less than $\tau/20$ (say) if N is sufficiently large. Thus to prove (7.3) it suffices to show that

$$\mathbf{P} \left(\sum_{J_0 \leq j \leq j' \leq J_1} \sup_{\substack{g_j \in A_j - A_{j+1} \\ \|g_j\|_X \leq 2^{j+2}}} \sup_{\substack{g_{j'} \in A_{j'} - A_{j'+1} \\ \|g_{j'}\|_X \leq 2^{j'+2}}} Q(g_j, g_{j'}) > \frac{\tau}{10} \right) = O \left(\exp \left(-\frac{1}{\beta} \log^2 N \right) \right). \quad (7.6)$$

The main difficulty is of course the presence of the suprema. On the other hand, the fact that the functions $g_j, g_{j'}$ are well controlled both in entropy and in X norm will allow us to

handle these suprema by relatively crude tools such as the union bound. By the pigeonhole principle, we can bound the left-hand side of (7.6) by

$$\sum_{J_0 \leq j \leq j' \leq J_1} \mathbf{P} \left(\sup_{\substack{g_j \in A_j - A_{j+1} \\ \|g_j\|_X \leq 2^{j+2}}} \sup_{\substack{g_{j'} \in A_{j'} - A_{j'+1} \\ \|g_{j'}\|_X \leq 2^{j'+2}}} Q(g_j, g_{j'}) > \frac{c_0}{\log^2 N} \cdot \tau \right),$$

for some small absolute constant c_0 . Since the number of pairs (j, j') is $O(\log^2 N)$, which is much smaller than $\exp(\frac{1}{\beta} \log^2 N)$, it now suffices to show (after adjusting the value of β slightly) that

$$\mathbf{P} \left(\sup_{\substack{g_j \in A_j - A_{j+1} \\ \|g_j\|_X \leq 2^{j+2}}} \sup_{\substack{g_{j'} \in A_{j'} - A_{j'+1} \\ \|g_{j'}\|_X \leq 2^{j'+2}}} Q(g_j, g_{j'}) > \frac{c_0}{\log^2 N} \cdot \tau \right) = O \left(\exp \left(-\frac{1}{\beta} \log^2 N \right) \right).$$

whenever $J_0 \leq j \leq j' \leq J_1$.

Fix j, j' as above. From (7.5) the number of possible values of $g_{j'}$ is at most $\exp(\frac{C}{1+2^{2j'}} \cdot m \log N)$. Thus by the union bound it suffices to show that

$$\mathbf{P} \left(\sup_{\substack{g_j \in A_j - A_{j+1} \\ \|g_j\|_X \leq 2^{j+2}}} Q(g_j, g_{j'}) > \frac{c_0}{\log^2 N} \cdot \tau \right) = O \left(\exp \left(-\frac{C}{1+2^{2j'}} \cdot m \log^2 N \right) \right);$$

for each $g_{j'} \in A_{j'} - A_{j'+1}$ with $\|g_{j'}\|_X \leq 2^{j'+2}$; note that we can absorb the $\exp(-\frac{1}{\beta} \log^2 N)$ factor since $2^{2j'} \leq m$.

Fix $g_{j'}$. We could also apply the union bound to eliminate the supremum over g_j , but it turns out that this will lead to inferior bounds at this stage, because of the poor ℓ_∞ control on $g_{j'}$. We have to first split the frequency domain \mathbf{Z}_N into two sets $\mathbf{Z}_N = E_1 \cup E_2$, where

$$E_1 := \{k \in \mathbf{Z}_N : |\hat{g}_{j'}(k)| \geq \frac{C_0 2^j \log^2 N}{\sqrt{N}}\} \text{ and } E_2 := \{k \in \mathbf{Z}_N : |\hat{g}_{j'}(k)| < \frac{C_0 2^j \log^2 N}{\sqrt{N}}\}$$

for some large absolute constant C_0 . Note that these sets depend on $g_{j'}$ but not on g_j . It thus suffices to show that

$$\mathbf{P} \left(\sup_{g_j \in A_j - A_{j+1} : \|g_j\|_X \leq 2^{j+2}} Q_i(g_j, g_{j'}) > \frac{c_0}{\log^2 N} \cdot \tau \right) = O \left(\exp \left(-\frac{C}{1+2^{2j'}} \cdot m \log^2 N \right) \right) \quad (7.7)$$

for $i = 1, 2$, where we have substituted (7.2), and $Q_i(g_j, g_{j'})$ is the random variable

$$Q_i(g_j, g_{j'}) := \left| \sum_{k \in E_i} (I_k - \tau) \operatorname{Re}(\hat{g}_j(k) \overline{\hat{g}_{j'}(k)}) \right|.$$

We treat the cases $i = 1, 2$ separately.

Proof of (7.7) when $i = 1$. For the contribution of the large frequencies E_1 we will take absolute values everywhere, which is fairly crude but conveys the major advantage that we will be able to easily eliminate the supremum in g_j . Note that since

$$|\hat{g}_{j'}(k)| \leq \|g_{j'}\|_X / \sqrt{N} \leq 2^{j'+2} / \sqrt{N},$$

we see that this case is vacuous unless

$$\frac{2^{j'+2}}{\sqrt{N}} \geq \frac{C_0 2^j \log^2 N}{\sqrt{N}},$$

or in other words,

$$2^{j'-j} \geq \frac{C_0 \log^2 N}{4}. \quad (7.8)$$

We then use the crude bound

$$|\hat{g}_j(k)| \leq \|g_j\|_X / \sqrt{N} \leq 2^{j+2} / \sqrt{N}$$

and the triangle inequality to conclude

$$\sup_{g_j \in A_j - A_{j+1}} Q_1(g_j, g_{j'}) \leq \frac{2^{j+2}}{\sqrt{N}} \sum_{k \in E_1} (I_k + \tau) |\hat{g}_{j'}(k)|.$$

By definition of E_1 , we have

$$\begin{aligned} \sum_{k \in E_1} 2\tau |\hat{g}_{j'}(k)| &\leq \frac{2\tau\sqrt{N}}{C_0 2^j \log^2 N} \sum_{k \in \mathbf{Z}_N} |\hat{g}_{j'}(k)|^2 \\ &= \frac{2\tau\sqrt{N}}{C_0 2^j \log^2 N} \|g_{j'}\|_{\ell_2}^2 \\ &\leq \frac{C\tau\sqrt{N}}{C_0 2^j \log^2 N}. \end{aligned}$$

Writing $I_k + \tau = (I_k - \tau) + 2\tau$, we conclude that

$$\sup_{g_j \in A_j - A_{j+1}} Q_1(g_j, g_{j'}) \leq \frac{2^{j+2}}{\sqrt{N}} \sum_{k \in E_1} (I_k - \tau) |\hat{g}_{j'}(k)| + \frac{C\tau}{C_0 \log^2 N}$$

and hence to prove (7.7) when $i = 1$, it would suffice (if C_0 is chosen sufficiently large) to show that

$$\mathbf{P} \left(\frac{2^{j+2}}{\sqrt{N}} \sum_{k \in E_1} (I_k - \tau) |\hat{g}_{j'}(k)| > \frac{c_0}{\log^2 N} \cdot \tau \right) = O \left(\exp \left(-\frac{C}{1 + 2^{2j'}} \cdot m \log^2 N \right) \right).$$

It thus suffices to show

$$\mathbf{P} \left(\left| \sum_{k \in E_1} (I_k - \tau) a(k) \right| \geq \gamma \right) = O \left(\exp \left(-\frac{C}{1 + 2^{2j'}} \cdot m \log^2 N \right) \right)$$

where $a(k) := |\hat{g}_{j'}(k)|$ and $\gamma := \frac{c_0 \tau \sqrt{N}}{C 2^j \log^2 N}$. Recall that

$$\|a(k)\|_{\ell_\infty(E_1)} = O(\|g_{j'}\|_X / \sqrt{N}) = O(2^{j'} / \sqrt{N})$$

and

$$\|a(k)\|_{\ell_2(E_1)} \leq \|\hat{g}_{j'}\|_{\ell_2} = \|g_{j'}\|_{\ell_2} = O(1).$$

We apply Lemma 6.6 and obtain

$$\mathbf{P} \left(\left| \sum_{k \in E_1} (I_k - \tau)a(k) \right| \geq \gamma \right) = O \left(\exp \left(-\frac{C\gamma^2}{4\tau + \gamma 2^{j'}/\sqrt{N}} \right) \right).$$

Using (7.8), we see that $\gamma 2^{j'}/\sqrt{N} \geq c \cdot \tau$ for some absolute constant $c > 0$, and conclude that

$$\mathbf{P} \left(\left| \sum_{k \in E_1} (I_k - \tau)a(k) \right| \geq \gamma \right) = O \left(\exp \left(-\frac{C \cdot c_0 \cdot \tau N}{2^{j+j'} \log^2 N} \right) \right).$$

Taking logarithms, we deduce that this contribution will be acceptable if

$$\frac{1}{1 + 2^{2j'}} \cdot m \log^2 N \leq C \cdot \frac{c_0 \cdot \tau N}{2^{j+j'} \log^2 N}$$

which holds (with some room to spare) thanks to (7.2).

Proof of (7.7) when $i = 2$. For the contribution of the small frequencies E_2 we use (7.5) and the union bound, and reduce to showing that

$$\mathbf{P} \left(Q_2(g_j, g_{j'}) > \frac{c_0}{\log^2 N} \cdot \tau \right) = O \left(\exp \left(-\frac{C}{1 + 2^{2j}} m \log^2 N \right) \right) \quad (7.9)$$

for any $g_j \in A_j - A_{j+1}$.

Fix g_j , and set $a(k) := \operatorname{Re}(\hat{g}_j(k) \overline{\hat{g}_{j'}(k)})$; thus

$$Q_2(g_j, g_{j'}) = \sum_{k \in E_2} (I_k - \tau)a(k).$$

By definition of E_2 , we have

$$\|a(k)\|_{\ell_\infty(E_2)} \leq O \left(\frac{2^j \log^2 N}{\sqrt{N}} \frac{\|g_j\|_X}{\sqrt{N}} \right) = O \left(\frac{2^{2j} \log^2 N}{N} \right)$$

while from Hölder and Plancherel

$$\|a(k)\|_{\ell_2(E_2)} \leq \|\hat{g}_{j'}\|_{\ell_2} \|\hat{g}_j\|_{\ell_\infty} = \frac{\|g_{j'}\|_{\ell_2} \|g_j\|_X}{\sqrt{N}} = O \left(\frac{2^j}{\sqrt{N}} \right).$$

We can apply Lemma 6.6 to conclude

$$\begin{aligned} \mathbf{P} \left(Q_2(g_j, g_{j'}) > \frac{c_0}{\log^2 N} \cdot \tau \right) &= O \left(\exp \left(-\frac{C \cdot \frac{c_0^2}{\log^4 N} \cdot \tau^2}{\tau[2^{2j}/N] + \tau[c_0 \log^2 N 2^{2j}/\log^2 NN]} \right) \right) \\ &= O \left(\exp \left(-\frac{C \cdot \log^{-4} N \cdot \tau N}{2^{2j}} \right) \right). \end{aligned}$$

Taking logarithms, we thus see that this contribution will be acceptable if

$$\frac{1}{1 + 2^{2j}} \cdot m \log^2 N \leq C \cdot \frac{\tau N}{2^{2j} \log^4 N}$$

which holds thanks to (7.2). This concludes the proof of Lemma 4.3 (assuming Proposition 7.2). \blacksquare

8 ‘Universal’ Encoding

Our results interact with the agenda of coding theory. In fact, one can think of the process of taking random measurements as a kind of universal coding strategy that we explain below. In a nutshell, consider an encoder/decoder pair which would operate roughly as follows:

- The Encoder and the Decoder share a collection of random vectors (X_k) where the X_k 's are independent Gaussian vectors with standard normal entries. In practice, we can imagine that the encoder would send the seed of a random generator so that the decoder would be able to reconstruct those ‘pseudo-random’ vectors.
- *Encoder.* To encode a discrete signal f , the encoder simply calculates the coefficients $y_k = \langle f, X_k \rangle$ and quantizes the vector y .
- *Decoder.* The decoder then receives the quantized values and reconstructs a signal by solving the linear program (1.10).

This encoding/decoding scheme is of course very different from those commonly discussed in the literature of information theory. In this scheme, the encoder would not try to know anything about the signal, nor would exploit any special structure of the signal; it would blindly correlate the signal with noise and quantize the output—effectively doing very little work. In other words, the encoder would treat each signal in exactly the same way, hence the name “universal encoding.” There are several aspects of such a strategy which seem worth exploring:

- *Robustness.* A fundamental problem with most existing coding strategies is their fragility vis a vis bit-loss. Take JPEG 2000, the current digital still-picture compression standard, for example. All the bits in JPEG 2000 do not have the same value and if important bits are missing (e.g. because of packet loss), then there is simply no way the information can be retrieved accurately.

The situation is very different when one is using the scheme suggested above. Suppose for example that with a little more than K coefficients one achieves the distortion obeying the power-law

$$\|f - f^\# \|^2 \lesssim 1/K. \quad (8.1)$$

(This would correspond to the situation where our objects are bounded in ℓ_1 .) Thus receiving a little more than K random coefficients essentially allows to reconstruct a signal as precisely as if one knew the K largest coefficients.

Now suppose that in each packet of information, we have both encoded the (quantized) value of the coefficients y_k but also the label of the corresponding coefficients k . Consider now a situation in which half of the information is lost in the sense that only half of the coefficients are actually received. What is the accuracy of the decoded message $f_{50\%}^\#$? This essentially corresponds to reducing the number of randomly sampled coefficients by a factor of two, and so by (8.1) we see that the distortion would obey

$$\|f - f_{50\%}^\# \|^2 \lesssim 2/K \quad (8.2)$$

and, therefore, losses would have minimal effect.

- *Security.* Suppose that someone would intercept the message. Then he/she would not be able to decode the message because he/she would not know in which random basis the coefficients are expressed. (In practice, in the case where one would exchange the seed of a random generator, one could imagine protecting it with standard technologies such as RSA. Thus this scheme can be viewed as a variant of the standard stream cipher, based on applying a XOR operation between the plain text and a pseudorandom keystream, but with the advantage of robustness.)
- *Cost Efficiency.* Nearly all coding scenarios work roughly as follows: we acquire a large number of measurements about an object of interest, which we then encode. This encoding process effectively discards most of the measured data so that only a fraction of the measurement is being transmitted. For concreteness, consider JPEG 2000, a prototype of a transform coder. We acquire a large number N of sample values of a digital image f . The encoder then computes all the N wavelet coefficients of f , and quantizes only the $B \ll N$ largest, say. Hence only a very small fraction of the wavelet coefficients of f are actually transmitted.

In stark contrast, our encoder makes measurements that are immediately used. Suppose we could design sensors which could actually measure the correlations $\langle f, X_k \rangle$. Then not only the decoded object would be nearly as good (in the ℓ_2 -distance) as that obtained by knowing all the wavelet coefficients and selecting the largest (it is expected that the ℓ_1 -reconstruction is well-behaved vis a vis quantization), but we would effectively encode all the measured coefficients and thus, we would not discard any data available about f (except for the quantization).

Even if one could make all of this practical, a fundamental question remains: is this an efficient strategy? That is, for a class of interesting signals, e.g. a class of digital images with bounded variations, would it be possible to adapt the ideas presented in this paper to show that this scheme does not use many more bits than what is considered necessary? In other words, it appears interesting to subject this compression scheme to a rigorous information theoretic analysis. This analysis would need to address 1) how one would want to efficiently quantize the values of the coefficients $\langle f, X_k \rangle$ and 2) how the quantization quantitatively affects the precision of the reconstructed signal.

9 Discussion

9.1 Robustness

To be widely applicable, we need noise-aware variants of the ideas presented in this paper which are robust against the effects of quantization, measurement noise and modeling error, as no real-world sensor can make perfectly accurate measurements. We view these issues as important research topics. For example, suppose that the measurements $y_k = \langle f, \psi_k \rangle$ are rounded up to the nearest multiple of q , say, so that the available information is of the form y_k^q with $-q/2 \leq y_k^q - y_k \leq q/2$. Then we would like to know whether the solution $f^\#$ to (1.10) or better, of the variant

$$\min_g \|g\|_{\ell_1}, \quad \text{subject to} \quad \|F_\Omega g - y^q\|_{\ell_\infty} \leq q/2$$

still obeys error estimates such as those introduced in Theorem 1.4. Our analysis seems to be amenable to this situation and work in progress shows that the quality of the reconstruction degrades gracefully as q increases. Precise quantitative answers would help establishing the information theoretic properties of the scheme introduced in Section 8.

9.2 Connections with other works

Our results are connected with very recent work of A. Gilbert, S. Muthukrishnan, and M. Strauss [32], [33]. In this work, one considers a discrete signal of length N which one would like to represent as a sparse superposition of sinusoids. In [32], the authors develop a randomized algorithm that essentially samples the signal f in the time domain $O(B^2 \text{poly}(\log N))$ times ($\text{poly}(\log N)$ denotes a polynomial term in $\log N$) and returns a vector of approximate Fourier coefficients. They show that under certain conditions, this vector gives, with positive probability, an approximation to the discrete Fourier transform of \hat{f} which is almost as good as that obtained by keeping the B -largest entries of the discrete Fourier transform of \hat{f} . In [33], the algorithm was refined so that (1) only $O(B \text{poly}(\log N))$ samples are needed and (2) so that the algorithm runs in $O(B \text{poly}(\log N))$ time which truly is a remarkable feat. To achieve this gain, however, one has to sample the signal on highly structured random grids.

Our approach is different in several aspects. First and foremost, we are given a fixed set of nonadaptive measurements. In other words, the way in which we stated the problem does not give us the ‘luxury’ of adaptively sampling the signals as in [33]. In this context, it is unclear how the methodology presented in [32, 33] would allow reconstructing the signal f from $O(B \text{poly}(\log N))$ arbitrary sampled values. In contrast, our results guarantee that an accurate reconstruction is possible for nearly all possible measurements sets taken from ensembles obeying **UUP** and **ERP**. Second, the methodology there essentially concerns the recovery of spiky signals from frequency samples and do not address other setups. Yet, there certainly is a similar flavor in the statements of their results. Of special interest is whether some of the ideas developed by this group of researchers might be fruitful to attack problems such as those discussed in this article.

While finishing the write-up of this paper, we became aware of very recent and independent work by David Donoho on a similar project [24]. In that paper which appeared one month before ours, Donoho essentially proves Theorem 1.1 for Gaussian ensembles. He also shows that if a measurement matrix obeys 3 conditions (CS1-CS3), then one can obtain the estimate (1.11). There is some overlap in methods, in particular the estimates of Szarek [52] on the condition numbers of random matrices (CS1) also play a key role in those papers, but there is also a greater reliance in those papers on further facts from high-dimensional geometry, in particular in understanding the shape of random sections of the ℓ_1 ball (CS2-CS3). Our proofs are completely different in style and approach, and most of our claims are different. While [24] only derives results for the Gaussian ensemble, this paper establishes that other types of ensembles such as the binary and the Fourier ensembles and even arbitrary measurement/synthesis pairs will work as well. This is important because this shows that concrete sensing mechanisms may be used in concrete applications.

In a companion [12] to this paper we actually improve on the results presented here and show that Theorem 1.4 holds for general measurement ensembles obeying the **UUP**. The implication for the Gaussian ensemble is that the recovery holds with an error in (1.11) of

size at most a constant times $(K/\log(N/K))^{-r}$.

10 Appendix: Proof of entropy estimate

In this section we prove Proposition 7.2. The material here is to a large extent borrowed from that in [6], [7], [43].

The entropy of the unit ball of a Hilbert space can be estimated using the *dual Sudakov inequality* of Pajor and Tomczak-Jaegerman [44] (See [8], [43] for a short “volume packing” proof, and [43] for further discussion):

Lemma 10.1 [44] *Let H be a n -dimensional Hilbert space with norm $\|\cdot\|_H$, and let B_H be the associated unit ball. Let e_1, \dots, e_n be an orthonormal basis of the Hilbert space H , and let $Z_1, \dots, Z_n \sim N(0, 1)$ be i.i.d. standard Gaussian random variables. Let $\|\cdot\|_Y$ be any other norm on \mathbf{C}^n . Then we have*

$$\mathcal{E}(B_H, B_Y, r) \leq Cr^{-2} \cdot \mathbf{E}(\|\sum_{j=1}^n Z_j e_j\|_Y)^2$$

where C is an absolute constant (independent of n).

To apply this Lemma, we need to estimate the X norm of certain randomized signs. Fortunately, this is easily accomplished:

Lemma 10.2 *Let $f \in \ell_2(\mathbb{Z}_N)$ and $Z(t)$, $t \in \mathbb{Z}_N$, be i.i.d. standard Gaussian random variables. Then*

$$\mathbf{E}(\|Zf\|_X) \leq C \cdot \sqrt{\log N} \cdot \|f\|_{\ell_2}.$$

The same statement holds if the Z 's are i.i.d. Bernoulli symmetric random variables ($Z(t) = \pm 1$ with equal probability).

Proof Let us normalize $\|f\|_{\ell_2} = 1$. For any $\lambda > 0$, we have

$$\begin{aligned} \mathbf{P}(\|Zf\|_X > \lambda) &= \mathbf{P}\left(\left|\sum_{t \in \mathbb{Z}_N} Z(t)f(t)e^{-2\pi itk/N}\right| > \lambda \text{ for some } k \in \mathbb{Z}_N\right) \\ &\leq N \sup_{k \in \mathbb{Z}_N} \mathbf{P}\left(\left|\sum_{t \in \mathbb{Z}_N} Z(t)f(t)e^{-2\pi itk/N}\right| > \lambda\right). \end{aligned}$$

If the $Z(t)$ are i.i.d. normalized Gaussians, then for each fixed k , $\sum_{t \in \mathbb{Z}_N} Z(t)f(t)e^{-2\pi itk/N}$ is a Gaussian with mean zero and standard deviation $\|f\|_{\ell_2} = 1$. Hence

$$\mathbf{P}(\|Zf\|_X > \lambda) \leq C \cdot N \cdot e^{-\lambda^2/2}.$$

Combining this with the trivial bound $\mathbf{P}(\|Zf\|_X > \lambda) \leq 1$ and then integrating in λ gives the result. The claim for i.i.d. Bernoulli variables is similar but uses Hoeffding's inequality; we omit the standard details. \blacksquare

Combining this lemma with Lemma 10.1, we immediately obtain

Corollary 10.3 *Let E be a non-empty subset of \mathbb{Z}_N ; note that $\ell_2(E)$ is both a Hilbert space (with the usual Hilbert space structure), as well as a normed vector space with the X norm. For all $r > 0$, we have*

$$\mathcal{E}(B_{\ell_2(E)}, B_X, r) \leq Cr^{-2} \cdot |E| \cdot \log N.$$

Now we turn to the set U_m introduced in the preceding section. Since the number of sets E of cardinality m is $\binom{N}{m} \leq N^m$, we have the crude bound

$$N(U_m, B_X, r) \leq N^m \sup_{E \subseteq \mathbb{Z}_N, |E|=m} N(\mathcal{E}, B_X, r)$$

and hence by Corollary 10.3

$$\mathcal{E}(U_m, B_X, r) \leq C(1 + r^{-2}) m \log N. \quad (10.1)$$

This already establishes (7.4) in the range $C^{-1} \leq r \leq C\sqrt{\log N}$. However, this bound is quite poor when r is large. For instance, when $r \geq m^{1/2}$ we have

$$\mathcal{E}(U_m, B_X, r) = 0 \quad (10.2)$$

since we have $\|f\|_X < m^{1/2}$ whenever $\|f\|_{\ell_2} \leq 1$ and $|\text{supp}(f)| \leq m$. In the regime $1 \ll r \leq m^{1/2}$ we can use the following support reduction trick of Bourgain to obtain a better bound:

Lemma 10.4 [6] *If $r \geq C\sqrt{\log N}$ and $m \geq C$, then*

$$N(U_m, B_X, r) \leq N(U_{m/2+C\sqrt{m}}, B_X, \frac{r}{\sqrt{2} + C/\sqrt{m}} - C\sqrt{\log N}).$$

Proof Let $f \in U_m$ and $E := \text{supp}(f)$, thus $\|f\|_{\ell_2} \leq 1$ and $|E| \leq m$. Let $\sigma(t) = \pm 1$ be i.i.d. Bernoulli symmetric variables. We write $f = \sigma f + (1 - \sigma)f$. From Lemma 10.2 for Bernoulli variables we have

$$\mathbf{E}(\|\sigma f\|_X) \leq C\sqrt{\log N}$$

and hence by Markov's inequality

$$\mathbf{P}(\|\sigma f\|_X \geq C\sqrt{\log N}) \leq \frac{1}{10}$$

for a suitable absolute constant C . Also observe that

$$\|(1 - \sigma)f\|_{\ell_2}^2 = \sum_{t \in E; \sigma(t) = -1} 4|f(t)|^2 = 2\|f\|_{\ell_2}^2 - 2 \sum_{t \in E} \sigma(t)|f(t)|^2,$$

and hence by Hoeffding's or Khintchine's inequalities and the normalization $\|f\|_{\ell_2} \leq 1$

$$\mathbf{P}(\|(1 - \sigma)f\|_{\ell_2} \geq \sqrt{2} + C/\sqrt{m}) \leq \frac{1}{10}$$

for a suitable absolute constant C . In a similar spirit, we have

$$\text{supp}((1 - \sigma)f) = |\{t \in \text{supp}(f) : \sigma(t) = -1\}| = \frac{1}{2} \text{supp}(f) - \frac{1}{2} \sum_{t \in E} \sigma(t),$$

and hence

$$\mathbf{P}(\text{supp}((1 - \sigma)f) \geq \frac{m}{2} + C\sqrt{m}) \leq \frac{1}{10}$$

for a suitable absolute constant C . Combining all these estimates together, we see that there exists a deterministic choice of signs $\sigma(t) = \pm 1$ (depending on f and E) such that

$$\|\sigma f\|_X \leq C\sqrt{\log N}; \|(1 - \sigma)f\|_{\ell_2} \leq \sqrt{2} + C/\sqrt{m}; \text{supp}((1 - \sigma)f) \leq \frac{m}{2} + C\sqrt{m}.$$

In particular, f is within $C\sqrt{\log N}$ (in X norm) from $(\sqrt{2} + C/\sqrt{m}) \cdot U_{m/2+C\sqrt{m}}$. We thus have

$$N(U_m, B_X, r) \leq N((\sqrt{2} + C/\sqrt{m})U_{m/2+C\sqrt{m}}, B_X, r - C\sqrt{\log N})$$

and the claim follows. ■

Iterating this lemma roughly $\log_{\sqrt{2}} \frac{r}{\sqrt{\log N}}$ times to reduce m and r , and then applying (10.1) once r becomes comparable with $\sqrt{\log N}$, we obtain

$$\mathcal{E}(U_m, B_X, r) \leq Cr^{-2} m (\log N)^2 \text{ whenever } C\sqrt{\log N} \leq r \leq m^{1/2},$$

which (together with (10.2)) yields (7.4) for all $r \geq C\sqrt{\log N}$.

It remains to address the case of small r , say $N^{-2} < r < 1/2$. A simple covering argument (see [43, Lemma 2.7]; the basic point is that $B_{\ell_2(E)}$ can be covered by $O(r^{-C|E|})$ translates of $r \cdot B_{\ell_2}(E)$) gives the general inequality

$$\mathcal{E}(B_{\ell_2(E)}, B_X, r) \leq C|E| \log \frac{1}{r} + \mathcal{E}(B_{\ell_2(E)}, B_X, 1)$$

for $0 < r < 1/2$, and hence by Corollary 10.3

$$\mathcal{E}(B_{\ell_2(E)}, B_X, r) \leq C|E| \log \frac{1}{r} + C|E| \log N.$$

Arguing as in the proof of (10.1) we thus have

$$\mathcal{E}(U_m, B_X, r) \leq Cm(\log N + \log \frac{1}{r}),$$

which gives (7.4) in the range $N^{-2} < r < 1/2$. This completes the proof of Proposition 7.2. ■

References

- [1] A. Noga, Y. Matias, B. Szegedy, The space complexity of approximating the frequency moments, *Proc. of the 28th ACM STOC*, ACM Press (1996), 20-29. Also; *J. Comp. Sys. Sci.* **58** (1999), 137-147.
- [2] Z. D. Bai and Y. Q. Yin, Limit of the smallest eigenvalue of a large-dimensional sample covariance matrix. *Ann. Probab.* **21** (1993), 1275-1294.
- [3] K. M. Ball. An elementary introduction to modern convex geometry. In S. Levy (Ed.), *Flavors of Geometry*, Number 31 in MSRI Publications, 1-58. New York: Cambridge U. Press, 1997.

- [4] P. Bloomfield and W. Steiger, *Least Absolute Deviations: Theory, Applications, and Algorithms*. Birkhäuser, Boston, 1983.
- [5] S. Boucheron, G. Lugosi, and P. Massart, A sharp concentration inequality with applications, *Random Structures Algorithms* **16** (2000), 277–292.
- [6] J. Bourgain, Bounded orthogonal systems and the $\Lambda(p)$ -set problem, *Acta Math.* **162** (1989), 227–245.
- [7] J. Bourgain, Remarks on Halasz-Montgomery type inequalities, Geometric aspects of functional analysis (Israel, 1992–1994), 25–39, *Oper. Theory Adv. Appl.*, **77**, Birkhäuser, Basel, 1995.
- [8] J. Bourgain, J. Lindenstrauss, V. Milman, Approximation of zonoids by zonotopes, *Acta Math.* **162** (1989), 73–141.
- [9] E. J. Candès, and P. S. Loh, Image reconstruction with ridgelets, SURF Technical report, California Institute of Technology, 2002.
- [10] E. J. Candès, D. L. Donoho, New Tight Frames of Curvelets and Optimal Representations of Objects with Piecewise C^2 Singularities, *Comm. on Pure and Appl. Math.*, **57** (2004), 219-266.
- [11] E. J. Candès, J. Romberg, and T. Tao, Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information, *to appear, IEEE Inf. Theory*.
- [12] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory* **51** 4203-4215.
- [13] E. J. Candès, and J. Romberg, The Role of Sparsity and Incoherence for Exactly Reconstructing a Signal from Limited Measurements, Technical Report, California Institute of Technology.
- [14] E. J. Candès, and J. Romberg, Quantitative Robust Uncertainty Principles and Optimally Sparse Decompositions, Technical Report, California Institute of Technology.
- [15] B. Carl, Entropy numbers, s -numbers, and eigenvalue problems. *J. Funct. Anal.* **41** (1981), 290–306.
- [16] S. S. Chen, Basis Pursuit, Stanford Ph. D. Thesis, 1995.
- [17] S. S. Chen, D. L. Donoho, and M. A. Saunders, Atomic decomposition by basis pursuit, *SIAM J. Scientific Computing* **20** (1999), 33–61.
- [18] A. Cohen, R. DeVore, P. Petrushev and H. Xu, Nonlinear approximation and the space $BV(\mathbb{R}^2)$. *Amer. J. Math.* **121** (1999), 587628.
- [19] K. R. Davidson and S. J. Szarek, Local operator theory, random matrices and Banach spaces, In: *Handbook in Banach Spaces* Vol I, ed. W. B. Johnson, J. Lindenstrauss, Elsevier (2001), 317–366.
- [20] R. A. DeVore, Nonlinear approximation, *Acta Numerica, 1998*, 51–150, *Acta Numer.* **7**, Cambridge Univ. Press, Cambridge, 1998.

- [21] D. C. Dobson, and F. Santosa, Recovery of blocky images from noisy and blurred data, *SIAM J. Appl. Math.* **56** (1996), 1181–1198.
- [22] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies. Data compression and harmonic analysis, *IEEE Trans. Inform. Theory* **44** (1998), 2435–2476.
- [23] D. L. Donoho, For Most Large Underdetermined Systems of Linear Equations the Minimal ℓ_1 -norm Solution is also the Sparsest Solution, Manuscript, September 2004.
- [24] D. L. Donoho, Compressed Sensing, Manuscript, September 2004.
- [25] D.L. Donoho, P.B. Stark, Uncertainty principles and signal recovery, *SIAM J. Appl. Math.* **49** (1989), 906–931.
- [26] D.L. Donoho and X. Huo, Uncertainty principles and ideal atomic decomposition, *IEEE Transactions on Information Theory*, **47** (2001), 2845–2862.
- [27] D. L. Donoho and M. Elad, Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization, *Proc. Natl. Acad. Sci. USA* **100** (2003), 2197–2202.
- [28] El Karoui, New Results about Random Covariance Matrices and Statistical Applications, Stanford Ph. .D. Thesis, August 2004.
- [29] H. G. Feichtinger, Atomic characterizations of modulation spaces through Gabor-type representations. *Rocky Mountain J. Math.* **19** (1989), 113–125.
- [30] J. J. Fuchs, On sparse representations in arbitrary redundant bases, *IEEE Transactions on Information Theory*, **50** (2004), 1341–1344.
- [31] A. Garnaev and E. Gluskin. The widths of a Euclidean ball. *Dokl. A. N. USSR* **277** (1984), 1048–1052. In Russian.
- [32] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, M. Strauss, Near-optimal sparse Fourier representations via sampling, *34th ACM Symposium on Theory of Computing*, Montréal, May 2002.
- [33] A. C. Gilbert, S. Muthukrishnan, and M. Strauss, Beating the B^2 bottleneck in estimating B -term Fourier representations, unpublished manuscript, May 2004.
- [34] R. Gribonval and M. Nielsen, Sparse representations in unions of bases. *IEEE Trans. Inform. Theory* **49** (2003), 3320–3325.
- [35] W.B. Johnson, J. Lindenstrauss, Extensions of Lipschitz mappings into Hilbert space, *Contemp. Math.* **26** (1984), 189–206.
- [36] I. M. Johnstone, On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* **29** (2001), 295–327.
- [37] B. Kashin. The widths of certain finite dimensional sets and classes of smooth functions, *Izvestia* **41** (1977), 334351.
- [38] T. Kühn. A lower estimate for entropy numbers. *J. Approx. Theory* **110** (2001), 120–124.

- [39] Ledoux, M, *The concentration of measure phenomenon*. Mathematical Surveys and Monographs **89**, . American Mathematical Society, Providence, RI, 2001.
- [40] A. E. Litvak, A. Pajor M. Rudelson, and N. Tomczak-Jaegermann, Smallest singular value of random matrices and geometry of random polytopes, manuscript, August 2004.
- [41] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1998.
- [42] V. A. Marchenko and L. A. Pastur, Distribution of eigenvalues in certain sets of random matrices, *Mat. Sb. (N.S.)* **72** (1967), 407–535 (in Russian).
- [43] G. Mockenhaupt and W. Schlag, *On the Hardy-Littlewood problem for random sets*, preprint.
- [44] A. Pajor, N. Tomczak-Jaegermann, Subspaces of small codimension of finite-dimensional Banach spaces, *Proc. Amer. Math. Soc.* **97** (1986), 637–642.
- [45] A. Pinkus, *N-Widths in Approximation Theory*. Springer-Verlag, Berlin, 1985.
- [46] G. Pisier. *The Volume of Convex Bodies and Banach Space Geometry*. Cambridge Tracts in Mathematics **94**, Cambridge University Press, Cambridge, 1989.
- [47] M. Rudelson, R. Vershynin, Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements, preprint.
- [48] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithm, *Physica D* **60**, 259–268.
- [49] F. Santosa, and W. W. Symes, Linear inversion of band-limited reflection seismograms, *SIAM J. Sci. Statist. Comput.* **7** (1986), 1307–1330.
- [50] C. Schutt. Entropy numbers of diagonal operators between symmetric Banach spaces. *J. Approx. Theory* **40** (1984), 121–128.
- [51] J. W. Silverstein, The smallest eigenvalue of a large dimensional Wishart matrix, *Ann. Probab.* **13** (1985), 1364–1368.
- [52] S. J. Szarek, Condition numbers of random matrices, *J. Complexity* **7** (1991), 131–149.