J Fourier Anal Appl (2018) 24:935–956
https://doi.org/10.1007/s00041-017-9550-x

Discrete Uncertainty Principles and Sparse Signal Processing

Afonso S. Bandeira[1] · **Megan E. Lewis**[2] · **Dustin G. Mixon**[3]

Received: 7 December 2015 / Revised: 22 May 2017 / Published online: 19 June 2017

Abstract We develop new discrete uncertainty principles in terms of numerical sparsity, which is a continuous proxy for the 0-norm. Unlike traditional sparsity, the continuity of numerical sparsity naturally accommodates functions which are nearly sparse. After studying these principles and the functions that achieve exact or near equality in them, we identify certain consequences in a number of sparse signal processing applications.

Keywords Uncertainty principle · Sparsity · Compressed sensing

1 Introduction

Uncertainty principles have maintained a significant role in both science and engineering for most of the past century. In 1927, the concept was introduced by Werner Heisenberg in the context of quantum mechanics [24], in which a particle's position and momentum are represented by wavefunctions $f, g \in L^2(\mathbb{R})$, and g happens to be the Fourier transform of f. Measuring the position or momentum of a particle amounts to drawing a random variable whose probability density function is a normal-

Communicated by Roman Vershynin.

✉ Dustin G. Mixon
dustin.mixon@gmail.com

1 Department of Mathematics, Courant Institute of Mathematical Sciences, New York University, New York, NY, USA

2 Detachment 5, Air Force Operational Test and Evaluation Center, Edwards AFB, CA, USA

3 Department of Mathematics and Statistics, Air Force Institute of Technology, Wright-Patterson AFB, OH, USA

ized version of $|f|^2$ or $|g|^2$, respectively. Heisenberg's uncertainty principle postulates a fundamental limit on the precision with which one can measure both position and momentum; in particular, the variance of the position measurement is small only if the momentum measurement exhibits large variance. From a mathematical perspective, this physical principle can be viewed as an instance of a much broader meta-theorem in harmonic analysis:

A nonzero function and its Fourier transform cannot be simultaneously localized.

Heisenberg's uncertainty principle provides a lower bound on the product of the variances of the probability density functions corresponding to f and $\hat{f}$. In the time since, various methods have emerged for quantifying localization. For example, instead of variance, one might consider entropy [6], the size of the density's support [2], or how rapidly it decays [23]. Furthermore, the tradeoff in localization need not be represented by a product—as we will see, it is sometimes more telling to consider a sum.

Beyond physics, the impossibility of simultaneous localization has had significant consequences in signal processing. For example, when working with the short-time Fourier transform, one is forced to choose between temporal and frequency resolution. More recently, the emergence of digital signal processing has prompted the investigation of uncertainty principles underlying the discrete Fourier transform, notably by Donoho and Stark [17], Tao [42], and Tropp [45]. Associated with this line of work is the uniform uncertainty principle of Candès and Tao [12], which played a key role in the development of compressed sensing. The present paper continues this investigation of discrete uncertainty principles with an eye on applications in sparse signal processing.

1.1 Background and Overview

For any finite abelian group G, let $\ell(G)$ denote the set of functions $x\colon G \to \mathbb{C}$, and $\widehat{G} \subseteq \ell(G)$ the group of characters over G. Then taking inner products with these characters and normalizing leads to the (unitary) Fourier transform $F\colon \ell(G) \to \ell(\widehat{G})$, namely

$$(Fx)[\chi] := \frac{1}{\sqrt{|G|}} \sum_{g\in G} x[g]\overline{\chi[g]} \qquad \forall \chi \in \widehat{G}.$$

The reader who is unfamiliar with Fourier analysis over finite abelian groups is invited to learn more in [43]. In the case where $G = \mathbb{Z}/n\mathbb{Z}$ (which we denote by $\mathbb{Z}_n$ in the sequel), the above definition coincides with the familiar discrete Fourier transform after one identifies characters with their frequencies. The following theorem provides two uncertainty principles in terms of the so-called 0-norm $\|\cdot\|_0$, defined to be number of nonzero entries in the argument.

Theorem 1 ([17, Theorem 1], [42, Theorem 1.1]) *Let G be a finite abelian group, and let $F\colon \ell(G) \to \ell(\widehat{G})$ denote the corresponding Fourier transform. Then*

$$\|x\|_0\|Fx\|_0 \geq |G| \qquad \forall x \in \ell(G)\backslash\{0\}. \tag{1}$$

Furthermore, if $|G|$ is prime, then

$$\|x\|_0 + \|Fx\|_0 \geq |G| + 1 \qquad \forall x \in \ell(G)\backslash\{0\}. \tag{2}$$

Proof Sketch For (1), apply the fact that the ℓ_1/ℓ_∞-induced norm of F is given by $\|F\|_{1\to\infty} = 1/\sqrt{|G|}$, along with Cauchy–Schwarz and Parseval's identity:

$$\begin{aligned}\|Fx\|_\infty &\leq \frac{1}{\sqrt{|G|}}\|x\|_1 \leq \sqrt{\frac{\|x\|_0}{|G|}}\|x\|_2 = \sqrt{\frac{\|x\|_0}{|G|}}\|Fx\|_2 \\ &\leq \sqrt{\frac{\|x\|_0\|Fx\|_0}{|G|}}\|Fx\|_\infty,\end{aligned}$$

where the last step bounds a sum in terms of its largest summand. Rearranging gives the result.

For (2), suppose otherwise that there exists $x \neq 0$ which violates the claimed inequality. Denote $\mathcal{J} = \operatorname{supp}(x)$ and pick some $\mathcal{I} \subseteq \widehat{G}\backslash\operatorname{supp}(Fx)$ with $|\mathcal{I}| = |\mathcal{J}|$. Then $0 = (Fx)_\mathcal{I} = F_{\mathcal{I}\mathcal{J}}x_\mathcal{J}$. Since the submatrix $F_{\mathcal{I}\mathcal{J}}$ is necessarily invertible by a theorem of Chebotarëv [39], we conclude that $x_\mathcal{J} = 0$, a contradiction. □

We note that the additive uncertainty principle above is much stronger than its multiplicative counterpart. Indeed, with the help of the arithmetic mean–geometric mean inequality, (1) immediately implies

$$\|x\|_0 + \|Fx\|_0 \geq 2\sqrt{\|x\|_0\|Fx\|_0} \geq 2\sqrt{|G|} \qquad \forall x \in \ell(G), \tag{3}$$

which is sharp when $G = \mathbb{Z}_n$ and n is a perfect square (simply take x to be a Dirac comb, specifically, the indicator function 1_K of the subgroup K of size $\sqrt{n}$). More generally, if n is not prime, then $n = ab$ with integers $a, b \in [2, n/2]$, and so $a + b \leq n/2 + 2 < n + 1$; as such, taking x to be an indicator function of the subgroup of size a (whose Fourier transform necessarily has 0-norm b) will violate (2). Overall, the hypothesis that $|G|$ is prime cannot be weakened. Still, something can be said if one slightly strengthens the hypothesis on x. For example, Theorem A in [44] gives that for every $S \subseteq G$,

$$\|x\|_0 + \|Fx\|_0 > \sqrt{|G|\|x\|_0}$$

for almost every $x \in \ell(G)$ supported on S. This suggests that extreme functions like the Dirac comb are atypical, i.e., (3) is "barely sharp".

One could analogously argue that, in some sense, (2) is "barely true" when $|G|$ is prime. For an illustration, Fig. 1 depicts a discrete version of the Gaussian function, which is constructed by first periodizing the function $f(t) = e^{-n\pi t^2}$ over the real line in order to have unit period, and then sampling this periodized function at multiples of $1/n$. As we verify in Sect. 3.2, the resulting function $x \in \ell(\mathbb{Z}_n)$ satisfies $Fx = x$, analogous to the fact that a Gaussian function in $L^2(\mathbb{R})$ with the proper width is fixed by the Fourier transform. Given its resemblance to the fast-decaying Gaussian function over $\mathbb{R}$, it comes as no surprise that many entries of this function are nearly zero. In the depicted case where $n = 211$ (which is prime), only 99 entries of this function manage to be larger than machine precision, and so from a numerical perspective, this function appears to contradict Theorem 1: $99 + 99 = 198 < 212 = 211 + 1$.

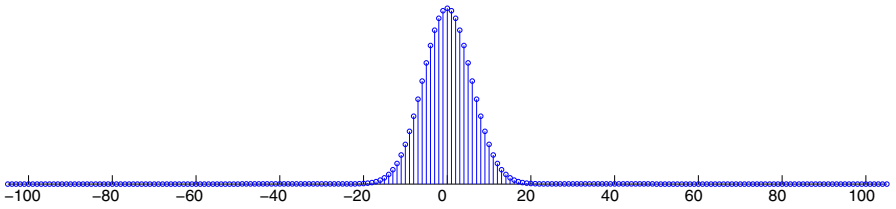

Fig. 1 Discrete Gaussian function, obtained by periodizing the function $f(t) = e^{-n\pi t^2}$ with period 1 before sampling at multiples of $1/n$. The resulting function in $\mathbb{Z}_n$ is fixed by the $n \times n$ discrete Fourier transform. In this figure, we take $n = 211$, and only 99 entries are larger than machine precision (i.e., 2.22×10^{-16}). As such, an unsuspecting signal processor might think $\|x\|_0$ and $\|Fx\|_0$ are both 99 instead of 211. Since 211 is prime and $99 + 99 = 198 < 212 = 211 + 1$, this illustrates a lack of numerical robustness in the additive uncertainty principle of Theorem 1. By contrast, our main result (Theorem 2) provides a robust alternative in terms of numerical sparsity, though the result is not valid for the discrete Fourier transform, but rather a random unitary matrix.

To help resolve this discrepancy, we consider a numerical version of traditional sparsity which is aptly named *numerical sparsity*:

$$\mathrm{ns}(x) := \frac{\|x\|_1^2}{\|x\|_2^2} \qquad \forall x \in \mathbb{C}^n \backslash \{0\}.$$

See Fig. 2 for an illustration. This ratio appeared as early as 1978 in the context of geophysics [20]. More recently, it has been used as a proxy for sparsity in various signal processing applications [14,25,31,36,40]. The numerical rank of a matrix is analogously defined as the square ratio of the nuclear and Frobenius norms, and has been used, for example, in Alon's work on extremal combinatorics [1]. We note that numerical sparsity is invariant under nonzero scaling, much like traditional sparsity. In addition, one bounds the other:

$$\mathrm{ns}(x) \leq \|x\|_0. \tag{4}$$

To see this, apply Cauchy–Schwarz to get

$$\|x\|_1 = \langle |x|, 1_{\mathrm{supp}(x)} \rangle \leq \|x\|_2 \|1_{\mathrm{supp}(x)}\|_2 = \|x\|_2 \sqrt{\|x\|_0},$$

where $|x|$ denotes the entrywise absolute value of x. Rearranging then gives (4). For this paper, the most useful feature of numerical sparsity is its continuity, as this will prevent near-counterexamples like the one depicted in Fig. 1. What follows is our main result, which leverages numerical sparsity to provide uncertainty principles that are analogous to those in Theorem 1:

Theorem 2 (Main result[1]) *Let U be an $n \times n$ unitary matrix. Then*

$$\mathrm{ns}(x)\,\mathrm{ns}(Ux) \geq \frac{1}{\|U\|_{1\to\infty}^2} \qquad \forall x \in \mathbb{C}^n \backslash \{0\}, \tag{5}$$

[1] Recall that $f(n) = O(g(n))$ if there exists $C, n_0 > 0$ such that $f(n) \leq Cg(n)$ for all $n > n_0$. We write $f(n) = O_\delta(g(n))$ if the constant C is a function of δ. Also, $f(n) = \Omega(g(n))$ if $g(n) = O(f(n))$, and $f(n) = o(g(n))$ if $f(n)/g(n) \to 0$ as $n \to \infty$.

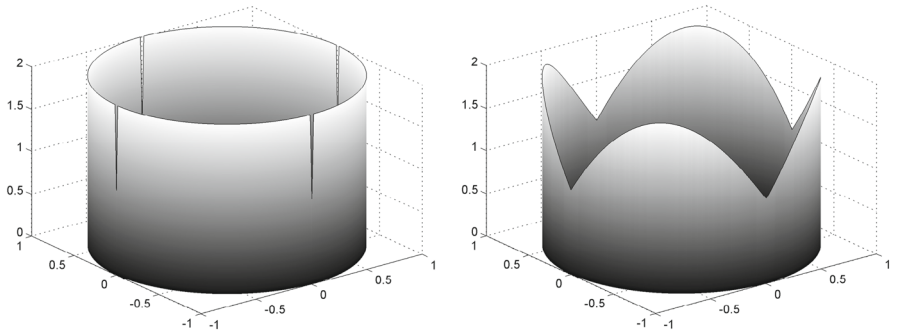

Fig. 2 Traditional sparsity $\|x\|_0$ (*left*) and numerical sparsity $\mathrm{ns}(x)$ (*right*) for all x in the unit circle in $\mathbb{R}^2$. This illustrates how numerical sparsity is a continuous analog of traditional sparsity; we leverage this feature to provide robust alternatives to the uncertainty principles of Theorem 1. In this case, one may verify that $\mathrm{ns}(x) \leq \|x\|_0$ by visual inspection.

where $\|\cdot\|_{1\to\infty}$ *denotes the induced matrix norm. Furthermore, there exists a universal constant* $c > 0$ *such that if* U *is drawn uniformly from the unitary group* $\mathrm{U}(n)$*, then with probability* $1 - e^{-\Omega(n)}$*,*

$$\mathrm{ns}(x) + \mathrm{ns}(Ux) \geq (c - o(1))n \qquad \forall x \in \mathbb{C}^n \backslash \{0\}. \tag{6}$$

Perhaps the most glaring difference between Theorems 1 and 2 is our replacement of the Fourier transform with an arbitrary unitary matrix. Such generalizations have appeared in the quantum physics literature (for example, see [29]), as well as in the sparse signal processing literature [15,16,21,40,44,45]. Our multiplicative uncertainty principle still applies when $U = F$, in which case $\|U\|_{1\to\infty} = 1/\sqrt{n}$. Considering (4), the uncertainty principle in this case immediately implies the analogous principle in Theorem 1. Furthermore, the proof is rather straightforward: Apply Hölder's inequality to get

$$\mathrm{ns}(x)\,\mathrm{ns}(Ux) = \frac{\|x\|_1^2}{\|x\|_2^2} \cdot \frac{\|Ux\|_1^2}{\|Ux\|_2^2} \geq \frac{\|x\|_1^2}{\|x\|_2^2} \cdot \frac{\|Ux\|_2^2}{\|Ux\|_\infty^2} = \frac{\|x\|_1^2}{\|Ux\|_\infty^2} \geq \frac{1}{\|U\|_{1\to\infty}^2}. \tag{7}$$

By contrast, the proof of our additive uncertainty principle is not straightforward, and it does not hold if we replace U with F. Indeed, as we show in Sect. 3.2, the discrete Gaussian function depicted in Fig. 1 has numerical sparsity $O(\sqrt{n})$, thereby violating (6); recall that the same function is a near-counterexample of the analogous principle in Theorem 1. Interestingly, our uncertainty principle establishes that the Fourier transform is rare in that the vast majority of unitary matrices offer much more uncertainty in the worst case. This naturally leads to the following question:

Problem 3 For each n, what is the largest $c = c(n)$ for which there exists a unitary matrix U that satisfies $\mathrm{ns}(x) + \mathrm{ns}(Ux) \geq cn$ for every $x \in \mathbb{C}^n \backslash \{0\}$?

Letting $x = e_1$ gives $\mathrm{ns}(x)+\mathrm{ns}(Ux) \leq 1+\|Ux\|_0 \leq n+1$, and so $c(n) \leq 1+o(1)$; a bit more work produces a strict inequality $c(n) < 1+1/n$ for $n \geq 4$. Also, our proof of the uncertainty principle implies $\liminf_{n\to\infty} c(n) \geq 1/540{,}000$.

1.2 Outline

The primary focus of this paper is Theorem 2. Having already proved the multiplicative uncertainty principle in (7), it remains to prove the additive counterpart, which we do in the following section. Next, Sect. 3 considers functions which achieve either exact or near equality in (5) when U is the discrete Fourier transform. Surprisingly, exact equality occurs in (5) precisely when it occurs in (1). We also show that the discrete Gaussian depicted in Fig. 1 achieves near equality in (5). We conclude in Sect. 4 by studying a few applications, specifically, sparse signal demixing, compressed sensing with partial Fourier operators, and the fast detection of sparse signals.

2 Proof of Additive Uncertainty Principle

In this section, we prove the additive uncertainty principle in Theorem 2. The following provides a more explicit statement of the principle we prove:

Theorem 4 *Draw U uniformly from the unitary group $\mathrm{U}(n)$. Then with probability $\geq 1-8e^{-n/4096}$,*

$$\mathrm{ns}(x)+\mathrm{ns}(Ux) \geq \frac{1}{9}\left\lfloor \frac{n}{60{,}000} \right\rfloor \qquad \forall x \in \mathbb{C}^n\backslash\{0\}.$$

For the record, we did not attempt to optimize the constants. Our proof of this theorem makes use of several ideas from the compressed sensing literature:

Definition 5 Take any $m \times n$ matrix $\Phi = [\varphi_1 \cdots \varphi_n]$.

(a) We say Φ exhibits (k,θ)-**restricted orthogonality** if

$$|\langle \Phi x, \Phi y\rangle| \leq \theta \|x\|_2 \|y\|_2$$

for every $x, y \in \mathbb{C}^n$ with $\|x\|_0, \|y\|_0 \leq k$ and disjoint support.

(b) We say Φ satisfies the (k,δ)-**restricted isometry property** if

$$(1-\delta)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1+\delta)\|x\|_2^2$$

for every $x \in \mathbb{C}^n$ with $\|x\|_0 \leq k$.

(c) We say Φ satisfies the (k,c)-**width property** if

$$\|x\|_2 \leq \frac{c}{\sqrt{k}}\|x\|_1$$

for every x in the nullspace of Φ.

The restricted isometry property is a now-standard sufficient condition for uniformly stable and robust reconstruction from compressed sensing measurements (for example, see [11]). As the following statement reveals, restricted orthogonality implies the restricted isometry property:

Lemma 6 [4, Lemma 11] *If a matrix satisfies (k, θ)-restricted orthogonality and its columns have unit norm, then it also satisfies the (k, δ)-restricted isometry property with $\delta = 2\theta$.*

To prove Theorem 4, we will actually make use of the width property, which was introduced by Kashin and Temlyakov [27] to characterize uniformly stable ℓ_1 reconstruction for compressed sensing. Luckily, the restricted isometry property implies the width property:

Lemma 7 ([9, Theorem 11], cf. [27]) *If a matrix satisfies the (k, δ)-restricted isometry property for some positive integer k and $\delta < 1/3$, then it also satisfies the $(k, 3)$-width property.*

What follows is a stepping-stone result that we will use to prove Theorem 4, but it is also of independent interest:

Theorem 8 *Draw U uniformly from the unitary group $\mathrm{U}(n)$. Then $[I\ U]$ satisfies the (k, δ)-restricted isometry property with probability $\geq 1 - 8e^{-\delta^2 n/256}$ provided $\delta < 1$ and*

$$n \geq \frac{256}{\delta^2} k \log\left(\frac{en}{k}\right). \tag{8}$$

This is perhaps not surprising, considering various choices of structured random matrices are known to form restricted isometries with high probability [3,8,12,28,34,35,37]. To prove Theorem 8, we show that the structured matrix enjoys restricted orthogonality with high probability, and then appeal to Lemma 6. Before proving this result, we first motivate it by proving the desired uncertainty principle:

Proof of Theorem 4 Take $k = \lfloor n/60{,}000 \rfloor$ and $\delta = 1/4$. We will show $\mathrm{ns}(x) + \mathrm{ns}(Ux) \geq k/9$ for every nonzero $x \in \mathbb{C}^n$. If $k = 0$, the result is immediate, and so $n \geq 60{,}000$ without loss of generality. In this regime, we have $k \in [n/120{,}000, n/60{,}000]$, and so

$$\frac{256}{\delta^2} k \log\left(\frac{en}{k}\right) \leq 4096 \log(120{,}000e) \cdot k \leq 60{,}000k \leq n.$$

Theorem 8 and Lemma 7 then give that $[I\ U]$ satisfies the $(k, 3)$-width property with probability $\geq 1 - 8e^{-n/4096}$. Observe that $z = [Ux; -x]$ resides in the nullspace of $[I\ U]$ regardless of $x \in \mathbb{C}^n$. In the case where x (and therefore z) is nonzero, the width property and the arithmetic mean–geometric mean inequality together give

$$\frac{k}{9} \le \frac{\|z\|_1^2}{\|z\|_2^2} = \frac{(\|x\|_1 + \|Ux\|_1)^2}{\|x\|_2^2 + \|Ux\|_2^2} = \frac{\|x\|_1^2 + 2\|x\|_1\|Ux\|_1 + \|Ux\|_1^2}{2\|x\|_2^2}$$
$$\le \mathrm{ns}(x) + \mathrm{ns}(Ux).$$

□

Proof of Theorem 8 Take $[I\ U] = [\varphi_1 \cdots \varphi_{2n}]$, and let k be the largest integer satisfying (8). We will demonstrate that $[I\ U]$ satisfies the (k, δ)-restricted isometry property, which will then imply the (k', δ)-restricted isometry property for all $k' < k + 1$, and therefore all k satisfying (8). To this end, define the random quantities

$$\theta^\star(U) := \max_{\substack{x,y\in\mathbb{C}^{2n} \\ \|x\|_0,\|y\|_0\le k \\ \mathrm{supp}(x)\cap\mathrm{supp}(y)=\emptyset}} \frac{|\langle \Phi x, \Phi y\rangle|}{\|x\|_2\|y\|_2}, \qquad \theta(U) := \max_{\substack{x,y\in\mathbb{C}^{2n} \\ \|x\|_0,\|y\|_0\le k \\ \mathrm{supp}(x)\subseteq[n] \\ \mathrm{supp}(y)\subseteq[n]^c}} \frac{|\langle \Phi x, \Phi y\rangle|}{\|x\|_2\|y\|_2}.$$

We first claim that $\theta^\star(U) \le \theta(U)$. To see this, for any x, y satisfying the constraints in $\theta^\star(U)$, decompose $x = x_1 + x_2$ so that x_1 and x_2 are supported in $[n]$ and $[n]^c$, respectively, and similarly $y = y_1 + y_2$. For notational convenience, let S denote the set of all 4-tuples (a, b, c, d) of k-sparse vectors in $\mathbb{C}^{2n}$ such that a and b are disjointly supported in $[n]$, while c and d are disjointly supported in $[n]^c$. Then $(x_1, y_1, x_2, y_2) \in S$. Since $\mathrm{supp}(x)$ and $\mathrm{supp}(y)$ are disjoint, and since I and U each have orthogonal columns, we have

$$\langle \Phi x, \Phi y\rangle = \langle \Phi x_1, \Phi y_2\rangle + \langle \Phi x_2, \Phi y_1\rangle.$$

As such, the triangle inequality gives

$$\begin{aligned}
\theta^\star(U) &= \max_{\substack{x,y\in\mathbb{C}^{2n} \\ \|x\|_0,\|y\|_0\le k \\ \mathrm{supp}(x)\cap\mathrm{supp}(y)=\emptyset}} \frac{|\langle \Phi x_1, \Phi y_2\rangle + \langle \Phi x_2, \Phi y_1\rangle|}{\|x\|_2\|y\|_2} \\
&\le \max_{(x_1,y_1,x_2,y_2)\in S} \frac{|\langle \Phi x_1, \Phi y_2\rangle| + |\langle \Phi x_2, \Phi y_1\rangle|}{\sqrt{\|x_1\|_2^2 + \|x_2\|_2^2}\sqrt{\|y_1\|_2^2 + \|y_2\|_2^2}} \\
&\le \left(\max_{(x_1,y_1,x_2,y_2)\in S} \frac{\|x_1\|_2\|y_2\|_2 + \|x_2\|_2\|y_1\|_2}{\sqrt{\|x_1\|_2^2 + \|x_2\|_2^2}\sqrt{\|y_1\|_2^2 + \|y_2\|_2^2}} \right) \theta(U) \\
&\le \theta(U),
\end{aligned}$$

where the last step follows from squaring and applying the arithmetic mean–geometric mean inequality:

$$\left(\frac{\sqrt{ad} + \sqrt{bc}}{\sqrt{(a+b)(c+d)}} \right)^2 = \frac{ad + bc + 2\sqrt{acbd}}{(a+b)(c+d)} \le \frac{ad + bc + (ac + bd)}{(a+b)(c+d)} = 1.$$

At this point, we seek to bound the probability that $\theta(U)$ is large. First, we observe an equivalent expression:

$$\theta(U) = \max_{\substack{x,y\in\mathbb{C}^n \\ \|x\|_2=\|y\|_2=1 \\ \|x\|_0,\|y\|_0\le k}} |\langle x, Uy\rangle|.$$

To estimate the desired probability, we will pass to an ϵ-net $\mathcal{N}_\epsilon$ of k-sparse vectors with unit 2-norm. A standard volume-comparison argument gives that the unit sphere in $\mathbb{R}^m$ enjoys an ϵ-net of size $\le (1+2/\epsilon)^m$ (see [47, Lemma 5.2]). As such, for each choice of k coordinates, we can cover the corresponding copy of the unit sphere in $\mathbb{C}^k = \mathbb{R}^{2k}$ with $\le (1+2/\epsilon)^{2k}$ points, and unioning these produces an ϵ-net of size

$$|\mathcal{N}_\epsilon| \le \binom{n}{k}\left(1+\frac{2}{\epsilon}\right)^{2k}.$$

To apply this ϵ-net, we note that $\|x-x'\|_2, \|y-y'\|_2 \le \epsilon$ and $\|x'\|_2 = \|y'\|_2 = 1$ together imply

$$\begin{aligned}|\langle x, Uy\rangle| &= |\langle x'+x-x', U(y'+y-y')\rangle| \\ &\le |\langle x', Uy'\rangle| + \|x-x'\|_2 + \|y-y'\|_2 + \|x-x'\|_2\|y-y'\|_2 \\ &\le |\langle x', Uy'\rangle| + 3\epsilon,\end{aligned}$$

where the last step assumes $\epsilon \le 1$. As such, the union bound gives

$$\begin{aligned}&\Pr(\theta(U) > t) \\ &= \Pr\left(\exists x, y \in \mathbb{C}^n,\ \|x\|_2 = \|y\|_2 = 1,\ \|x\|_0, \|y\|_0 \le k \text{ s.t. } |\langle x, Uy\rangle| > t\right) \\ &\le \Pr\left(\exists x, y \in \mathcal{N}_\epsilon \text{ s.t. } |\langle x, Uy\rangle| > t - 3\epsilon\right) \\ &\le \sum_{x,y\in\mathcal{N}_\epsilon} \Pr\left(|\langle x, Uy\rangle| > t - 3\epsilon\right) \\ &= \binom{n}{k}^2\left(1+\frac{2}{\epsilon}\right)^{4k} \Pr\left(|\langle e_1, Ue_1\rangle| > t - 3\epsilon\right), \end{aligned} \tag{9}$$

where the last step uses the fact that the distribution of U is invariant under left- and right-multiplication by any deterministic unitary matrix (e.g., unitary matrices that send e_1 to x and y to e_1, respectively). It remains to prove tail bounds on $U_{11} := \langle e_1, Ue_1\rangle$. First, we apply the union bound to get

$$\begin{aligned}\Pr(|U_{11}| > u) &\le \Pr\left(|\operatorname{Re}(U_{11})| > \frac{u}{\sqrt{2}}\right) + \Pr\left(|\operatorname{Im}(U_{11})| > \frac{u}{\sqrt{2}}\right) \\ &= 4\Pr\left(\operatorname{Re}(U_{11}) > \frac{u}{\sqrt{2}}\right),\end{aligned} \tag{10}$$

where the last step uses the fact that $\operatorname{Re}(U_{11})$ has even distribution. Next, we observe that $\operatorname{Re}(U_{11})$ has the same distribution as $g/\sqrt{h}$, where g has standard normal distri-

bution and h has chi-squared distribution with $2n$ degrees of freedom. Indeed, this can be seen from one method of constructing the matrix U: Start with an $n \times n$ matrix G with iid $N(0,1) + iN(0,1)$ complex Gaussian entries and apply Gram–Schmidt to the columns; the first column of U is then the first column of G divided by its norm $\sqrt{h}$. Let $s > 0$ be arbitrary (to be selected later). Then $g/\sqrt{h} > u/\sqrt{2}$ implies that either $g > \sqrt{s}u/\sqrt{2}$ or $h < s$. As such, the union bound implies

$$\Pr\left(\operatorname{Re}(U_{11}) > \frac{u}{\sqrt{2}}\right) \le 2\max\left\{\Pr\left(g > \sqrt{s}\frac{u}{\sqrt{2}}\right), \Pr(h < s)\right\}. \tag{11}$$

For the first term, Proposition 7.5 in [19] gives

$$\Pr\left(g > \sqrt{s}\frac{u}{\sqrt{2}}\right) \le e^{-su^2/4}. \tag{12}$$

For the second term, Lemma 1 in [30] gives $\Pr(h < 2n - \sqrt{8nx}) \le e^{-x}$ for any $x > 0$. Picking $x = (2n-s)^2/(8n)$ then gives

$$\Pr(h < s) \le e^{-(2n-s)^2/(8n)}. \tag{13}$$

We use the estimate $\binom{n}{k} \le (en/k)^k$ when combining (9)–(13) to get

$$\log\Big(\Pr(\theta(U) > t)\Big) \le 2k\log\left(\frac{en}{k}\right) + 4k\log\left(1 + \frac{2}{\epsilon}\right) + \log 8 - \min\left\{\frac{s(t-3\epsilon)^2}{4}, \frac{(2n-s)^2}{8n}\right\}.$$

Notice $n/k \ge (256/\delta^2)\log(en/k) \ge 256$ implies that taking $\epsilon = \sqrt{(k/n)\log(en/k)}$ gives

$$\sqrt{\frac{en}{k}} - \frac{2}{\epsilon} = \left(1 - \frac{2}{\sqrt{e\log(n/k)}}\right)\sqrt{\frac{en}{k}} \ge \left(1 - \frac{2}{\sqrt{e\log(256)}}\right)\sqrt{256e} \ge 1,$$

which can be rearranged to get

$$\log\left(1 + \frac{2}{\epsilon}\right) \le \frac{1}{2}\log\left(\frac{en}{k}\right).$$

As such, we also pick $s = n$ and $t = \sqrt{(64k/n)\log(en/k)}$ to get

$$\log\Big(\Pr(\theta(U) > t)\Big) \le 4k\log\left(\frac{en}{k}\right) + \log 8 - \frac{25}{4}k\log\left(\frac{en}{k}\right) \le \log 8 - 2k\log\left(\frac{en}{k}\right).$$

Since we chose k to be the largest integer satisfying (8), we therefore have $\theta(U) \leq \sqrt{(64k/n)\log(n/k)}$ with probability $\geq 1 - 8e^{-\delta^2 n/256}$. Lemma 6 then gives the result. □

3 Low Uncertainty with the Discrete Fourier Transform

In this section, we study functions which achieve either exact or near equality in our multiplicative uncertainty principle (6) in the case where the unitary matrix U is the discrete Fourier transform.

3.1 Exact Equality in the Multiplicative Uncertainty Principle

We seek to understand when equality is achieved in (6) in the special case of the discrete Fourier transform. For reference, the analogous result for (1) is already known:

Theorem 9 [17, Theorem 13] *Suppose $x \in \ell(\mathbb{Z}_n)$ satisfies $\|x\|_0\|Fx\|_0 = n$. Then x has the form $x = cT^a M^b 1_K$, where $c \in \mathbb{C}$, K is a subgroup of $\mathbb{Z}_n$, and $T, M : \ell(\mathbb{Z}_n) \to \ell(\mathbb{Z}_n)$ are translation and modulation operators defined by*

$$(Tx)[j] := x[j-1], \qquad (Mx)[j] := e^{2\pi ij/n}x[j] \qquad \forall j \in \mathbb{Z}_n.$$

Here, i denotes the imaginary unit $\sqrt{-1}$.

In words, equality is achieved in (1) by indicator functions of subgroups, namely, the so-called Dirac combs (as well as their scalar multiples, translations, modulations). We seek an analogous characterization for our uncertainty principle (6). Surprisingly, the characterization is identical:

Theorem 10 *Suppose $x \in \ell(\mathbb{Z}_n)$. Then $\mathrm{ns}(x)\,\mathrm{ns}(Fx) = n$ if and only if $\|x\|_0\|Fx\|_0 = n$.*

Proof ($\Leftarrow$) This follows directly from (4), along with Theorems 1 and 2.

($\Rightarrow$) It suffices to show that $\mathrm{ns}(x) = \|x\|_0$ and $\mathrm{ns}(Fx) = \|Fx\|_0$. Note that both F and F^{-1} are unitary operators and $\|F\|_{1\to\infty}^2 = \|F^{-1}\|_{1\to\infty}^2 = 1/n$. By assumption, taking $y := Fx$ then gives

$$\mathrm{ns}(F^{-1}y)\,\mathrm{ns}(y) = \mathrm{ns}(x)\,\mathrm{ns}(Fx) = n.$$

We will use the fact that x and y each achieve equality in the first part of Theorem 2 with $U = F$ and $U = F^{-1}$, respectively. Notice from the proof (7) that equality occurs only if x and y satisfy equality in Hölder's inequality, that is,

$$\|x\|_1\|x\|_\infty = \|x\|_2^2, \qquad \|y\|_1\|y\|_\infty = \|y\|_2^2. \tag{14}$$

To achieve the first equality in (14),

$$\sum_{j\in\mathbb{Z}_n} |x[j]|^2 = \|x\|_2^2 = \|x\|_1\|x\|_\infty = \sum_{j\in\mathbb{Z}_n} |x[j]| \max_{k\in\mathbb{Z}_n} |x[k]|.$$

This implies that $|x[j]| = \max_k |x[k]|$ for every j with $x[j] \neq 0$. Similarly, in order for the second equality in (14) to hold, $|y[j]| = \max_k |y[k]|$ for every j with $y[j] \neq 0$. As such, $|x| = a1_A$ and $|y| = b1_B$ for some $a, b > 0$ and $A, B \subseteq \mathbb{Z}_n$. Then

$$\mathrm{ns}(x) = \frac{\|x\|_1^2}{\|x\|_2^2} = \frac{(a|A|)^2}{a^2|A|} = |A| = \|x\|_0,$$

and similarly, $\mathrm{ns}(y) = \|y\|_0$. □

3.2 Near Equality in the Multiplicative Uncertainty Principle

Having established that equality in the new multiplicative uncertainty principle (5) is equivalent to equality in the analogous principle (1), we wish to separate these principles by focusing on near equality. For example, in the case where n is prime, $\mathbb{Z}_n$ has no nontrivial proper subgroups, and so by Theorem 9, equality is only possible with identity basis elements and complex exponentials. On the other hand, we expect the new principle to accommodate nearly sparse vectors, and so we appeal to the discrete Gaussian depicted in Fig. 1:

Theorem 11 *Define $x \in \ell(\mathbb{Z}_n)$ by*

$$x[j] := \sum_{j' \in \mathbb{Z}} e^{-n\pi(\frac{j}{n}+j')^2} \qquad \forall j \in \mathbb{Z}_n. \tag{15}$$

Then $Fx = x$ and $\mathrm{ns}(x)\,\mathrm{ns}(Fx) \leq (2 + o(1))n$.

In words, the discrete Gaussian achieves near equality in the uncertainty principle (5). Moreover, numerical evidence suggests that $\mathrm{ns}(x)\,\mathrm{ns}(Fx) = (2 + o(1))n$, i.e., the 2 is optimal for the discrete Gaussian. Note that this does not depend on whether n is prime or a perfect square. Recall that a function $f \in C^\infty(\mathbb{R})$ is *Schwarz* if $\sup_{x\in\mathbb{R}} |x^\alpha f^{(\beta)}(x)| < \infty$ for every pair of nonnegative integers α and β. We use this to quickly prove a well-known lemma that will help us prove Theorem 11:

Lemma 12 *Suppose $f \in C^\infty(\mathbb{R})$ is Schwarz and construct a discrete function $x \in \ell(\mathbb{Z}_n)$ by periodizing and sampling f as follows:*

$$x[j] = \sum_{j' \in \mathbb{Z}} f\left(\frac{j}{n} + j'\right) \qquad \forall j \in \mathbb{Z}_n. \tag{16}$$

Then the discrete Fourier transform of x is determined by $\hat{f}(\xi) := \int_{-\infty}^{\infty} f(t)e^{-2\pi i\xi t}\,dt$:

$$(Fx)[k] = \sqrt{n} \sum_{k' \in \mathbb{Z}} \hat{f}(k + k'n) \qquad \forall k \in \mathbb{Z}_n.$$

Proof Since f is Schwarz, we may apply the Poisson summation formula:

$$x[j] = \sum_{j' \in \mathbb{Z}} f\left(\frac{j}{n} + j'\right) = \sum_{l \in \mathbb{Z}} \hat{f}(l) e^{2\pi i j l/n}.$$

Next, the geometric sum formula gives

$$\begin{aligned}(Fx)[k] &= \frac{1}{\sqrt{n}} \sum_{j \in \mathbb{Z}_n} \left(\sum_{l \in \mathbb{Z}} \hat{f}(l) e^{2\pi i j l/n} \right) e^{-2\pi i j k/n} \\ &= \frac{1}{\sqrt{n}} \sum_{l \in \mathbb{Z}} \hat{f}(l) \sum_{j \in \mathbb{Z}_n} \left(e^{2\pi i (l-k)/n} \right)^j = \sqrt{n} \sum_{\substack{l \in \mathbb{Z} \\ l \equiv k \bmod n}} \hat{f}(l).\end{aligned}$$

The result then follows from a change of variables. □

Proof of Theorem 11 It is straightforward to verify that the function $f(t) = e^{-n\pi t^2}$ is Schwarz. Note that defining x according to (16) then produces (15). Considering $\hat{f}(\xi) = n^{-1/2} e^{-\pi \xi^2/n}$, one may use Lemma 12 to quickly verify that $Fx = x$. To prove Theorem 11, it then suffices to show that $\mathrm{ns}(x) \leq (\sqrt{2} + o(1))\sqrt{n}$. We accomplish this by bounding $\|x\|_2$ and $\|x\|_1$ separately.

To bound $\|x\|_2$, we first expand a square to get

$$\|x\|_2^2 = \sum_{j \in \mathbb{Z}_n} \left(\sum_{j' \in \mathbb{Z}} e^{-n\pi (\frac{j}{n} + j')^2} \right)^2 = \sum_{j \in \mathbb{Z}_n} \sum_{j' \in \mathbb{Z}} \sum_{j'' \in \mathbb{Z}} e^{-n\pi [(\frac{j}{n} + j')^2 + (\frac{j}{n} + j'')^2]}.$$

Since all of the terms in the sum are nonnegative, we may infer a lower bound by discarding the terms for which $j'' \neq j'$. This yields the following:

$$\|x\|_2^2 \geq \sum_{j \in \mathbb{Z}_n} \sum_{j' \in \mathbb{Z}} e^{-2n\pi (\frac{j}{n} + j')^2} = \sum_{k \in \mathbb{Z}} e^{-2\pi k^2/n} \geq \int_{-\infty}^{\infty} e^{-2\pi x^2/n} dx - 1 = \sqrt{\frac{n}{2}} - 1,$$

where the last inequality follows from an integral comparison. Next, we bound $\|x\|_1$ using a similar integral comparison:

$$\|x\|_1 = \sum_{j \in \mathbb{Z}_n} \sum_{j' \in \mathbb{Z}} e^{-n\pi (\frac{j}{n} + j')^2} = \sum_{k \in \mathbb{Z}} e^{-\pi k^2/n} \leq \int_{-\infty}^{\infty} e^{-\pi x^2/n} dx + 1 = \sqrt{n} + 1.$$

Overall, we have

$$\mathrm{ns}(x) = \frac{\|x\|_1^2}{\|x\|_2^2} \leq \frac{(\sqrt{n} + 1)^2}{\sqrt{n/2} - 1} = (\sqrt{2} + o(1))\sqrt{n}.$$

□

4 Applications

Having studied the new uncertainty principles in Theorem 2, we now take some time to identify certain consequences in various sparse signal processing applications. In particular, we report consequences in sparse signal demixing, in compressed sensing with partial Fourier operators, and in the fast detection of sparse signals.

4.1 Sparse Signal Demixing

Suppose a signal x is sparse in the Fourier domain and corrupted by noise ϵ which is sparse in the time domain (such as speckle). The goal of demixing is to recover the original signal x given the corrupted signal $z = x + \epsilon$; see [32] for a survey of various related demixing problems. Provided Fx and ϵ are sufficiently sparse, it is known that this recovery can be accomplished by solving

$$v^\star \quad := \quad \operatorname{argmin} \quad \|v\|_1 \quad \text{subject to} \quad [I \ F]v = Fz, \tag{17}$$

where, if successful, the solution $v^\star$ is the column vector obtained by concatenating Fx and ϵ; see [38] for an early appearance of this sort of approach. To some extent, we know how sparse Fx and ϵ must be for this ℓ_1 recovery method to succeed. Coherence-based guarantees in [15,16,21] show that it suffices for $v^\star$ to be k-sparse with $k = O(\sqrt{n})$, while restricted isometry–based guarantees [5,11] allow for $k = O(n)$ if $[I \ F]$ is replaced with a random matrix. This disparity is known as the *square-root bottleneck* [46]. In particular, does $[I \ F]$ perform similarly to a random matrix, or is the coherence-based sufficient condition on k also necessary?

In the case where n is a perfect square, it is well known that the coherence-based sufficient condition is also necessary. Indeed, let K denote the subgroup of $\mathbb{Z}_n$ of size $\sqrt{n}$ and suppose $x = 1_K$ and $\epsilon = -1_K$. Then $[Fx; \epsilon]$ is $2\sqrt{n}$-sparse, and yet $z = 0$, thereby forcing $v^\star = 0$. On the other hand, if n is prime, then the additive uncertainty principle of Theorem 1 implies that every member of the nullspace of $[I \ F]$ has at least $n + 1$ nonzero entries, and so $v^\star \neq 0$ in this setting. Still, considering Fig. 1, one might expect a problem from a stability perspective. In this section, we use numerical sparsity to show that $\Phi = [I \ F]$ cannot break the square-root bottleneck, even if n is prime. To do this, we will make use of the following theorem:

Theorem 13 (see [9,27]) *Denote* $\Delta(y) := \operatorname{argmin} \|x\|_1$ *subject to* $\Phi x = y$. *Then*

$$\|\Delta(\Phi x) - x\|_2 \leq \frac{C}{\sqrt{k}} \|x - x_k\|_1 \qquad \forall x \in \mathbb{R}^n \tag{18}$$

if and only if Φ *satisfies the* (k, c)*-width property. Furthermore,* $C \asymp c$ *in both directions of the equivalence.*

Take x as defined in (15). Then $[x; -x]$ lies in the nullspace of $[I \ F]$ and

$$\operatorname{ns}\big([x; -x]\big) = \frac{(2\|x\|_1)^2}{2\|x\|_2^2} = 2\operatorname{ns}(x) \leq (2\sqrt{2} + o(1))\sqrt{n},$$

where the last step follows from the proof of Theorem 11. As such, $[I\ F]$ satisfies the (k, c)-width property for some c independent of n only if $k = O(\sqrt{n})$. Furthermore, Theorem 13 implies that stable demixing by ℓ_1 reconstruction requires $k = O(\sqrt{n})$, thereby proving the necessity of the square-root bottleneck in this case.

It is worth mentioning that the restricted isometry property is a sufficient condition for (18) (see [11], for example), and so by Theorem 8, one can break the square-root bottleneck by replacing the F in $[I\ F]$ with a random unitary matrix. This gives a uniform demixing guarantee which is similar to those provided by McCoy and Tropp [33], though the convex program they consider differs from (17).

4.2 Compressed Sensing with Partial Fourier Operators

Consider the random $m \times n$ matrix obtained by drawing rows uniformly with replacement from the $n \times n$ discrete Fourier transform matrix. If $m = \Omega_\delta(k \operatorname{polylog} n)$, then the resulting partial Fourier operator satisfies the restricted isometry property, and this fact has been dubbed the *uniform uncertainty principle* [12]. A fundamental problem in compressed sensing is determining the smallest number m of random rows necessary. To summarize the progress to date, Candès and Tao [12] first found that $m = \Omega_\delta(k \log^6 n)$ rows suffice, then Rudelson and Vershynin [37] proved $m = \Omega_\delta(k \log^4 n)$, and recently, Bourgain [8] achieved $m = \Omega_\delta(k \log^3 n)$; Nelson, Price and Wootters [34] also achieved $m = \Omega_\delta(k \log^3 n)$, but using a slightly different measurement matrix. In this subsection, we provide a lower bound: in particular, $m = \Omega_\delta(k \log n)$ is necessary whenever k divides n. Our proof combines ideas from the multiplicative uncertainty principle and the classical problem of coupon collecting.

The coupon collector's problem asks how long it takes to collect all k coupons in an urn if you repeatedly draw one coupon at a time randomly with replacement. It is a worthwhile exercise to prove that the expected number of trials scales like $k \log k$. We will require even more information about the distribution of the random number of trials:

Theorem 14 (see [13,18]) *Let T_k denote the random number of trials it takes to collect k different coupons, where in each trial, a coupon is drawn uniformly from the k coupons with replacement.*

(a) *For each $a \in \mathbb{R}$,*

$$\lim_{k\to\infty} \Pr\left(T_k \le k \log k + ak\right) = e^{-e^{-(a+\gamma)}},$$

where $\gamma \approx 0.5772$ denotes the Euler–Mascheroni constant.

(b) *There exists $c > 0$ such that for each k,*

$$\sup_{a\in\mathbb{R}} \left| \Pr\left(T_k \le k \log k + ak\right) - e^{-e^{-(a+\gamma)}} \right| \le \frac{c \log k}{k}.$$

Lemma 15 *Suppose k divides n, and draw m iid rows uniformly from the $n \times n$ discrete Fourier transform matrix to form a random $m \times n$ matrix Φ. If $m < k \log k$, then the*

nullspace of Φ *contains a* k*-sparse vector with probability* $\geq 0.4 - c(\log k)/k$*, where* c *is the constant from Theorem 14(b).*

Proof Let K denote the subgroup of $\mathbb{Z}_n$ of size k, and let 1_K denote its indicator function. We claim that some modulation of 1_K resides in the nullspace of Φ with the probability reported in the lemma statement. Let H denote the subgroup of $\mathbb{Z}_n$ of size n/k. Then the Fourier transform of each modulation of 1_K is supported on some coset of H. Letting M denote the random row indices that are drawn uniformly from $\mathbb{Z}_n$, a modulation of 1_K resides in the nullspace of Φ precisely when M fails to intersect the corresponding coset of H. As there are k cosets, each with probability $1/k$, this amounts to a coupon-collecting problem (explicitly, each "coupon" is a coset, and we "collect" the cosets that M intersects). The result then follows immediately from Theorem 14(b):

$$\Pr(T_k \leq m) \leq e^{-e^{-(m/k - \log k + \gamma)}} + \frac{c \log k}{k} \leq e^{-e^{-\gamma}} + \frac{c \log k}{k} \leq 0.6 + \frac{c \log k}{k}.$$

□

Presumably, one may remove the divisibility hypothesis in Lemma 15 at the price of weakening the conclusion. We suspect that the new conclusion would declare the existence of a vector x of numerical sparsity k such that $\|\Phi x\|_2 \ll \|x\|_2$. If so, then Φ fails to satisfy the so-called *robust width property*, which is necessary and sufficient for stable and robust reconstruction by ℓ_1 minimization [9]. For the sake of simplicity, we decided not to approach this, but we suspect that modulations of the discrete Gaussian would adequately fill the role of the current proof's modulated indicator functions.

What follows is the main result of this subsection:

Theorem 16 *Let* k *be sufficiently large, suppose* k *divides* n*, and draw* m *iid rows uniformly from the* $n \times n$ *discrete Fourier transform matrix to form a random* $m \times n$ *matrix* Φ*. Take* $\delta < 1/3$*. Then* Φ *satisfies the* (k, δ)*-restricted isometry property with probability* $\geq 2/3$ *only if*

$$m \geq C(\delta) k \log(en),$$

where $C(\delta)$ *is some constant depending only on* δ*.*

Proof In the event that Φ satisfies (k, δ)-RIP, we know that no k-sparse vector lies in the nullspace of Φ. Therefore, Lemma 15 implies

$$m \geq k \log k, \tag{19}$$

since otherwise Φ fails to be (k, δ)-RIP with probability $\geq 0.4 - c(\log k)/k > 1/3$, where the last step uses the fact that k is sufficiently large. Next, we leverage standard techniques from compressed sensing: (k, δ)-RIP implies (18) with $C = C_1(\delta)$ (see [10, Theorem 3.3]), which in turn implies

$$m \geq C_2(\delta) k \log\left(\frac{en}{k}\right) \tag{20}$$

by Theorem 11.7 in [19]. Since Φ is (k,δ)-RIP with positive probability, we know there exists an $m \times n$ matrix which is (k,δ)-RIP, and so m must satisfy (20). Combining with (19) then gives

$$m \geq \max\left\{k \log k, C_2(\delta) k \log\left(\frac{en}{k}\right)\right\}.$$

The result then follows from applying the bound $\max\{a,b\} \geq (a+b)/2$ and then taking $C(\delta) := (1/2)\min\{1, C_2(\delta)\}$. □

We note that the necessity of $k \log n$ random measurements contrasts with the proportional-growth asymptotic adopted in [7] to study the restricted isometry property of Gaussian matrices. Indeed, it is common in compressed sensing to consider phase transitions in which k, m and n are taken to infinity with fixed ratios k/m and m/n. However, since random partial Fourier operators fail to be restricted isometries unless $m = \Omega_\delta(k \log n)$, such a proportional-growth asymptotic fails to capture the so-called *strong phase transition* of these operators [7].

The proof of Theorem 16 relies on the fact that the measurements are drawn at random. By contrast, it is known that *every* $m \times n$ partial Hadamard operator fails to satisfy (k,δ)-RIP unless $m = \Omega_\delta(k \log n)$ [22,41]. We leave the corresponding deterministic result in the Fourier case for future work.

4.3 Fast Detection of Sparse Signals

The previous subsection established fundamental limits on the number of Fourier measurements necessary to perform compressed sensing with a uniform guarantee. However, for some applications, signal reconstruction is unnecessary. In this subsection, we consider one such application, namely sparse signal detection, in which the goal is to test the following hypotheses:

$$\begin{aligned} H_0 &: \ x = 0 \\ H_1 &: \ \|x\|_2^2 = \frac{n}{k}, \ \|x\|_0 \leq k. \end{aligned}$$

Here, we assume we know the 2-norm of the sparse vector we intend to detect, and we set it to be $\sqrt{n/k}$ without loss of generality (this choice of scaling will help us interpret our results later). We will assume the data is accessed according to the following query–response model:

Definition 17 (*Query–response model*) If the ith query is $j_i \in \mathbb{Z}_n$, then the ith response is $(Fx)[j_i] + \epsilon_i$, where the ϵ_i's are iid complex random variables with some distribution such that

$$\mathbb{E}|\epsilon_i| = \alpha, \qquad \mathbb{E}|\epsilon_i|^2 = \beta^2.$$

The coefficient of variation v of $|\epsilon_i|$ is defined as

$$v = \frac{\sqrt{\mathrm{Var}\,|\epsilon_i|}}{\mathbb{E}|\epsilon_i|} = \frac{\sqrt{\beta^2 - \alpha^2}}{\alpha}. \tag{21}$$

Note that for any scalar $c \neq 0$, the mean and variance of $|c\epsilon_i|$ are $|c|\alpha$ and $|c|^2 \operatorname{Var} |\epsilon_i|$, respectively. As such, v is scale invariant and is simply a quantification of the "shape" of the distribution of $|\epsilon_i|$. We will evaluate the responses to our queries with an ℓ_1 detector, defined below.

Definition 18 (ℓ_1 *detector*) Fix a threshold τ. Given responses $\{y_i\}_{i=1}^m$ from the query–response model, if

$$\sum_{i=1}^{m} |y_i| > \tau,$$

then reject H_0.

The following is the main result of this section:

Theorem 19 *Suppose $\alpha \leq 1/(8k)$. Randomly draw m indices uniformly from $\mathbb{Z}_n$ with replacement, input them into the query–response model and apply the ℓ_1 detector with threshold $\tau = 2m\alpha$ to the responses. Then*

$$\Pr\left(\textit{reject } H_0 \;\middle|\; H_0\right) \leq p \tag{22}$$

and

$$\Pr\left(\textit{fail to reject } H_0 \;\middle|\; H_1\right) \leq p \tag{23}$$

provided $m \geq (8k + 2v^2)/p$, where v is the coefficient of variation defined in (21).

In words, the probability that the ℓ_1 detector delivers a false positive is at most p, as is the probability that it delivers a false negative. These error probabilities can be estimated better given more information about the distribution of the random noise, and presumably, the threshold τ can be modified to decrease one error probability at the price of increasing the other. Notice that we only use $O(k)$ samples in the Fourier domain to detect a k-sparse signal. Since the sampled indices are random, it will take $O(\log n)$ bits to communicate each query, leading to a total computational burden of $O(k \log n)$ operations. This contrasts with the state-of-the-art sparse fast Fourier transform algorithms which require $\Omega(k \log(n/k))$ samples and take $O(k \operatorname{polylog} n)$ time (see [26] and references therein). We suspect k-sparse signals cannot be detected with substantially fewer samples (in the Fourier domain or any domain).

We also note that the acceptable noise magnitude $\alpha = O(1/k)$ is optimal in some sense. To see this, consider the case where k divides n and x is a properly scaled indicator function of the subgroup of size k. Then Fx is the indicator function of the subgroup of size n/k. (Thanks to our choice of scaling, each nonzero entry in the Fourier domain has unit magnitude.) Since a proportion of $1/k$ entries is nonzero in the Fourier domain, we can expect to require $O(k)$ random samples in order to observe a nonzero entry, and the ℓ_1 detector will not distinguish the entry from accumulated noise unless $\alpha = O(1/k)$.

Before proving Theorem 19, we first prove a couple of lemmas. We start by estimating the probability of a false positive:

Lemma 20 *Take $\epsilon_1, \ldots, \epsilon_m$ to be iid complex random variables with $\mathbb{E}|\epsilon_i| = \alpha$ and $\mathbb{E}|\epsilon_i|^2 = \beta^2$. Then*

$$\Pr\left(\sum_{i=1}^{m} |\epsilon_i| > 2m\alpha\right) \leq p$$

provided $m \geq v^2/p$, where v is the coefficient of variation of $|\epsilon_i|$ defined in (21).

Proof Denoting $X := \sum_{i=1}^{m} |\epsilon_i|$, we have $\mathbb{E}X = m\alpha$ and $\operatorname{Var} X = m(\beta^2 - \alpha^2)$. Chebyshev's inequality then gives

$$\Pr\left(\sum_{i=1}^{m} |\epsilon_i| - m\alpha > t\right) \leq \Pr(|X - \mathbb{E}X| > t) \leq \frac{\operatorname{Var} X}{t^2} = \frac{m(\beta^2 - \alpha^2)}{t^2}.$$

Finally, we take $t = m\alpha$ to get

$$\Pr\left(\sum_{i=1}^{m} |\epsilon_i| > 2m\alpha\right) \leq m\frac{(\beta^2 - \alpha^2)}{(m\alpha)^2} = \frac{\beta^2 - \alpha^2}{m\alpha^2} \leq \frac{\beta^2 - \alpha^2}{\alpha^2} \cdot \frac{p}{v^2} = p.$$

□

Next, we leverage the multiplicative uncertainty principle in Theorem 2 to estimate moments of noiseless responses:

Lemma 21 *Suppose $\|x\|_0 \leq k$ and $\|x\|_2^2 = n/k$. Draw j uniformly from $\mathbb{Z}_n$ and define $Y := |(Fx)[j]|$. Then*

$$\mathbb{E}Y \geq \frac{1}{k}, \qquad \mathbb{E}Y^2 = \frac{1}{k}.$$

Proof Recall that $\operatorname{ns}(x) \leq \|x\|_0 \leq k$. With this, Theorem 2 gives

$$n \leq \operatorname{ns}(x) \operatorname{ns}(Fx) \leq k \operatorname{ns}(Fx).$$

We rearrange and apply the definition of numerical sparsity to get

$$\frac{n}{k} \leq \operatorname{ns}(Fx) = \frac{\|Fx\|_1^2}{\|Fx\|_2^2} = \frac{\|Fx\|_1^2}{\|x\|_2^2} = \frac{\|Fx\|_1^2}{n/k},$$

where the second to last equality is due to Parseval's identity. Thus, $\|Fx\|_1 \geq n/k$. Finally,

$$\mathbb{E}Y = \frac{1}{n}\sum_{j \in \mathbb{Z}_n} |(Fx)[j]| = \frac{1}{n}\|Fx\|_1 \geq \frac{1}{k}$$

and

$$\mathbb{E}Y^2 = \frac{1}{n}\sum_{j\in\mathbb{Z}_n} |(Fx)[j]|^2 = \frac{1}{n}\|Fx\|_2^2 = \frac{1}{k}.$$

□

Proof of Theorem 19 Lemma 20 gives (22), and so it remains to prove (23). Denoting $Y_i := |(Fx)[j_i]|$, we know that $|y_i| \geq Y_i - |\epsilon_i|$, and so

$$\Pr\left(\sum_{i=1}^m |y_i| \leq 2ma\right) \leq \Pr\left(\sum_{i=1}^m Y_i - \sum_{i=1}^m |\epsilon_i| \leq 2ma\right). \tag{24}$$

For notational convenience, put $Z := \sum_{i=1}^m Y_i - \sum_{i=1}^m |\epsilon_i|$. We condition on the size of the noise and apply Lemma 20 with the fact that $m \geq v^2/(p/2)$ to bound (24):

$$\begin{aligned}\Pr(Z \leq 2m\alpha) &= \Pr\left(Z \leq 2m\alpha \;\middle|\; \sum_{i=1}^m |\epsilon_i| > 2m\alpha\right)\Pr\left(\sum_{i=1}^m |\epsilon_i| > 2m\alpha\right)\\ &\quad + \Pr\left(Z \leq 2m\alpha \;\middle|\; \sum_{i=1}^m |\epsilon_i| \leq 2m\alpha\right)\Pr\left(\sum_{i=1}^m |\epsilon_i| \leq 2m\alpha\right)\\ &\leq \frac{p}{2} + \Pr\left(\sum_{i=1}^m Y_i \leq 4m\alpha\right).\end{aligned} \tag{25}$$

Now we seek to bound the second term of (25). Taking $X = \sum_{i=1}^m Y_i$, Lemma 21 gives $\mathbb{E}X \geq m/k$ and $\operatorname{Var} X = m \operatorname{Var} Y_i \leq m\mathbb{E}Y_i^2 = m/k$. As such, applying Chebyshev's inequality gives

$$\Pr\left(\sum_{i=1}^m Y_i < \frac{m}{k} - t\right) \leq \Pr(X \leq \mathbb{E}X - t) \leq \Pr(|X - \mathbb{E}X| > t) \leq \frac{\operatorname{Var}(X)}{t^2} \leq \frac{m}{kt^2}.$$

Recalling that $\alpha \leq 1/(8k)$, we take $t = m/(2k)$ to get

$$\begin{aligned}\Pr\left(\sum_{i=1}^m Y_i \leq 4m\alpha\right) &\leq \Pr\left(\sum_{i=1}^m Y_i \leq \frac{m}{2k}\right)\\ &= \Pr\left(\sum_{i=1}^m Y_i \leq \frac{m}{k} - t\right) \leq \frac{m}{kt^2} = \frac{4k}{m} \leq \frac{p}{2},\end{aligned} \tag{26}$$

where the last step uses the fact that $m \geq 8k/p$. Combining (24), (25), and (26) gives the result. □

Acknowledgements The authors thank Laurent Duval, Joel Tropp, and the anonymous referees for multiple suggestions that significantly improved the presentation of our results and our discussion of the relevant literature. ASB was supported by AFOSR Grant No. FA9550-12-1-0317. DGM was supported by an AFOSR Young Investigator Research Program award, NSF Grant No. DMS-1321779, and AFOSR Grant No. F4FGA05076J002. The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

References

1. Alon, N.: Problems and results in extremal combinatorics. Discret. Math. **273**, 31–53 (2003)
2. Amrein, W.O., Berthier, A.M.: On support properties of L^p-functions and their Fourier transforms. J. Funct. Anal. **24**, 258–267 (1977)
3. Bandeira, A.S., Fickus, M., Mixon, D.G., Moreira, J.: Derandomizing restricted isometries via the Legendre symbol. arXiv:1406.4089
4. Bandeira, A.S., Fickus, M., Mixon, D.G., Wong, P.: The road to deterministic matrices with the restricted isometry property. J. Fourier Anal. Appl. **19**, 1123–1149 (2013)
5. Baraniuk, R., Davenport, M., DeVore, R., Wakin, M.: A simple proof of the restricted isometry property for random matrices. Constr. Approx. **28**, 253–263 (2008)
6. Beckner, W.: Inequalities in Fourier analysis. Ann. Math. **102**, 159–182 (1975)
7. Blanchard, J.D., Cartis, C., Tanner, J.: Compressed sensing: How sharp is the restricted isometry property? SIAM Rev. **53**, 105–125 (2011)
8. Bourgain, J.: An improved estimate in the restricted isometry problem. Lect. Notes Math. **2116**, 65–70 (2014)
9. Cahill, J., Mixon, D.G.: Robust width: A characterization of uniformly stable and robust compressed sensing. arXiv:1408.4409
10. Cai, T.T., Zhang, A.: Sharp RIP bound for sparse signal and low-rank matrix recovery. Appl. Comput. Harmon. Anal. **35**, 74–93 (2013)
11. Candès, E.J.: The restricted isometry property and its implications for compressed sensing. C. R. Acad. Sci. Paris Ser. I **346**, 589–592 (2008)
12. Candès, E.J., Tao, T.: Near-optimal signal recovery from random projections: Universal encoding strategies? IEEE Trans. Inf. Theory **52**, 5406–5425 (2006)
13. Csörgő, S.: A rate of convergence for coupon collectors. Acta Sci. Math. (Szeged) **57**, 337–351 (1993)
14. Demanet, L., Hand, P.: Scaling law for recovering the sparsest element in a subspace. arXiv:1310.1654
15. Donoho, D.L., Elad, M.: Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization. Proc. Natl. Acad. Sci. USA **100**, 2197–2202 (2003)
16. Donoho, D.L., Huo, X.: Uncertainty principles and ideal atomic decomposition. IEEE Trans. Inf. Theory **47**, 2845–2862 (2001)
17. Donoho, D.L., Stark, P.B.: Uncertainty principles and signal recovery. SIAM J. Appl. Math. **49**, 906–931 (1989)
18. Erdős, P., Rényi, A.: On a classical problem of probability theory. Magy. Tud Akad. Mat. Kutató Int. Közl. **6**, 215–220 (1961)
19. Foucart, A., Rauhut, H.: A Mathematical Introduction to Compressive Sensing. Applied and Numerical Harmonic Analysis. Birkhäuser, Basel (2013)
20. Gray, W.C.: Variable norm deconvolution, Tech. Rep. SEP-14, Stanford University (1978)
21. Gribonval, R., Nielsen, M.: Highly sparse representations from dictionaries are unique and independent of the sparseness measure. Appl. Comput. Harmon. Anal. **22**, 335–355 (2007)
22. Guédon, O., Mendelson, S., Pajor, A., Tomczak-Jaegermann, N.: Majorizing measures and proportional subsets of bounded orthonormal systems. Rev. Mat. Iberoam. **24**, 1075–1095 (2008)
23. Hardy, G.H.: A theorem concerning Fourier transforms. J. Lond. Math. Soc. **8**, 227–231 (1933)
24. Heisenberg, W.: Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. Z. Phys. (in German) **43**, 172–198 (1927)
25. Hurley, N., Rickard, S.: Comparing measures of sparsity. IEEE Trans. Inf. Theory **55**, 4723–4741 (2009)
26. Indyk P, Kapralov M.: Sample-optimal Fourier sampling in any constant dimension. In: foundations of computer science (FOCS), 2014 IEEE 55th Annual Symposium on 2014, pp. 514–523
27. Kashin, B.S., Temlyakov, V.N.: A remark on compressed sensing. Math. Notes **82**, 748–755 (2007)

28. Krahmer, F., Mendelson, S., Rauhut, H.: Suprema of chaos processes and the restricted isometry property. Commun. Pure Appl. Math. **67**, 1877–1904 (2014)
29. Kraus, K.: Complementary observables and uncertainty relations. Phys. Rev. D **35**, 3070–3075 (1987)
30. Laurent, B., Massart, P.: Adaptive estimation of a quadratic functional by model selection. Ann. Stat. **28**, 1302–1338 (2000)
31. Lopes, M.E.: Estimating unknown sparsity in compressed sensing. JMLR W&CP **28**, 217–225 (2013)
32. McCoy, M.B., Cevher, V., Dinh, Q.T., Asaei, A., Baldassarre, L.: Convexity in source separation: models, geometry, and algorithms. Signal Proc. Mag. **31**, 87–95 (2014)
33. McCoy, M.B., Tropp, J.A.: Sharp recovery bounds for convex demixing, with applications. Found. Comput. Math. **14**, 503–567 (2014)
34. Nelson, J., Price, E., Wootters, M.: New constructions of RIP matrices with fast multiplication and fewer rows. In: SODA, pp. 1515–1528 (2014)
35. Rauhut, H.: Compressive sensing and structured random matrices. Theor. Found. Numer. Methods Sparse Recover. **9**, 1–92 (2010)
36. Repetti, A., Pham, M.Q., Duval, L., Chouzenoux, É., Pesquet, J.-C.: Euclid in a taxicab: sparse blind deconvolution with smoothed ℓ_1/ℓ_2 regularization. IEEE Signal Process. Lett. **22**, 539–543 (2014)
37. Rudelson, M., Vershynin, R.: On sparse reconstruction from Fourier and Gaussian measurements. Comm. Pure Appl. Math. **61**, 1025–1045 (2008)
38. Santosa, F., Symes, W.W.: Linear inversion of band-limited reflection seismograms. SIAM J. Sci. Stat. Comput. **7**, 1307–1330 (1986)
39. Stevenhagen, P., Lenstra, H.W.: Chebotarëv and his density theorem. Math. Intell. **18**, 26–37 (1996)
40. Studer, C.: Recovery of signals with low density. arXiv:1507.02821
41. Talagrand, M.: Selecting a proportion of characters. Israel J. Math. **108**, 173–191 (1998)
42. Tao, T.: An uncertainty principle for cyclic groups of prime order. Math. Res. Lett. **12**, 121–127 (2005)
43. Terras, A.: Fourier Analysis on Finite Groups and Applications. Cambridge University Press, Cambridge (1999)
44. Tropp, J.A.: The sparsity gap: uncertainty principles proportional to dimension. In: CISS, pp. 1–6 (2010)
45. Tropp, J.A.: On the linear independence of spikes and sines. J. Fourier Anal. Appl. **14**, 838–858 (2008)
46. Tropp, J.A.: On the conditioning of random subdictionaries. Appl. Comput. Harmon. Anal. **25**, 1–24 (2008)
47. Vershynin, R.: Introduction to the non-asymptotic analysis of random matrices. Theory and Applications, Cambridge University Press, Compressed Sensing (2012)