

UC Merced

UC Merced Previously Published Works

Title

The Noise Collector for sparse recovery in high dimensions

Permalink

<https://escholarship.org/uc/item/0dk492jk>

Journal

Proceedings of the National Academy of Sciences of the United States of America,
117(21)

ISSN

0027-8424

Authors

Moscoso, Miguel
Novikov, Alexei
Papanicolaou, George
et al.

Publication Date

2020-05-26

DOI

10.1073/pnas.1913995117

Peer reviewed



The Noise Collector for sparse recovery in high dimensions

Miguel Moscoso^{a,1}, Alexei Novikov^b, George Papanicolaou^{c,1}, and Chryssoula Tsogka^d

^aDepartment of Mathematics, Universidad Carlos III de Madrid, Leganes, Madrid 28911, Spain; ^bDepartment of Mathematics, Pennsylvania State University, University Park, PA 16802; ^cDepartment of Mathematics, Stanford University, Stanford, CA 94305; and ^dDepartment of Applied Mathematics, University of California, Merced, CA 95343

Contributed by George Papanicolaou, March 25, 2020 (sent for review August 13, 2019; reviewed by Albert C. Fannjiang and Jack Xin)

The ability to detect sparse signals from noisy, high-dimensional data is a top priority in modern science and engineering. It is well known that a sparse solution of the linear system $\mathcal{A}\rho = b_0$ can be found efficiently with an ℓ_1 -norm minimization approach if the data are noiseless. However, detection of the signal from data corrupted by noise is still a challenging problem as the solution depends, in general, on a regularization parameter with optimal value that is not easy to choose. We propose an efficient approach that does not require any parameter estimation. We introduce a no-phantom weight τ and the Noise Collector matrix \mathcal{C} and solve an augmented system $\mathcal{A}\rho + \mathcal{C}\eta = b_0 + e$, where e is the noise. We show that the ℓ_1 -norm minimal solution of this system has zero false discovery rate for any level of noise, with probability that tends to one as the dimension of b_0 increases to infinity. We obtain exact support recovery if the noise is not too large and develop a fast Noise Collector algorithm, which makes the computational cost of solving the augmented system comparable with that of the original one. We demonstrate the effectiveness of the method in applications to passive array imaging.

high-dimensional probability | convex geometry | sparsity-promoting algorithms | noisy data

We want to find sparse solutions $\rho \in \mathbb{R}^K$ for

$$\mathcal{A}\rho = b \tag{1}$$

from highly incomplete measurement data $b = b_0 + e \in \mathbb{R}^N$ corrupted by noise e , where $1 \ll N < K$. In the noiseless case, ρ can be found exactly by solving the optimization problem (1)

$$\rho_* = \arg \min_{\rho} \|\rho\|_{\ell_1}, \text{ subject to } \mathcal{A}\rho = b, \tag{2}$$

provided that the measurement matrix $\mathcal{A} \in \mathbb{R}^{N \times K}$ satisfies additional conditions (e.g., decoherence or restricted isometry properties) (2, 3) and that the solution vector ρ has a small number M of nonzero components or degrees of freedom. When measurements are noisy, exact recovery is no longer possible. However, the exact support of ρ can still be determined if the noise is not too strong. The most commonly used approach is to solve the ℓ_2 -relaxed form of Eq. 2:

$$\rho_{\lambda} = \arg \min_{\rho} \left(\lambda \|\rho\|_{\ell_1} + \|\mathcal{A}\rho - b\|_{\ell_2}^2 \right), \tag{3}$$

which is known as Lasso in the statistics literature (4). There are sufficient conditions for the support of ρ_{λ} to be contained within the true support [e.g., the works of Fuchs (5), Tropp (6), Wainwright (7), and Maleki et al. (8)]. These conditions depend on the signal-to-noise ratio (SNR), which is not known and must be estimated, and on the regularization parameter λ , which must be carefully chosen and/or adaptively changed (9). Although such an adaptive procedure improves the outcome, the resulting solutions tend to include a large number of “false positives” in practice (10). Belloni et al. (11) proposed to solve the square root

Lasso minimization problem instead of Eq. 3, which makes the regularization parameter λ independent of the SNR. Our contribution is a computationally efficient method for exact support recovery with no false positives in noisy settings. It also does not require an estimate on SNR.

Main Results

Suppose that ρ is an M -sparse solution of system [1] with no noise, where the columns of \mathcal{A} have unit length. Our main result ensures that we can still recover the support of ρ when the data are noisy by looking at the support of ρ_{τ} found as

$$(\rho_{\tau}, \eta_{\tau}) = \arg \min_{\rho, \eta} \left(\tau \|\rho\|_{\ell_1} + \|\eta\|_{\ell_1} \right), \tag{4}$$

subject to $\mathcal{A}\rho + \mathcal{C}\eta = b_0 + e$,

with an $O(1)$ weight τ and an appropriately chosen Noise Collector matrix $\mathcal{C} \in \mathbb{R}^{N \times \Sigma}$, $\Sigma \gg K$. The minimization problem [4] can be understood as a relaxation of [2] as it works by absorbing all of the noise and possibly, some signal in $\mathcal{C}\eta_{\tau}$.

The following theorem shows that, if the signal is pure noise and the columns of \mathcal{C} are chosen independently and at random on the unit sphere $\mathbb{S}^{N-1} = \{x \in \mathbb{R}^N, \|x\|_{\ell_2} = 1\}$, then $\mathcal{C}\eta_{\tau} = e$ for any level of noise, with large probability.

Theorem 1 (No-Phantom Signal). *Suppose that $b_0 = 0$ and that $e/\|e\|_{\ell_2}$ is uniformly distributed on \mathbb{S}^{N-1} . Fix $\beta > 1$, and draw $\Sigma = N^{\beta}$ columns for \mathcal{C} independently from the uniform distribution on \mathbb{S}^{N-1} . For any $\kappa > 0$, there are constants $\tau = \tau(\kappa, \beta)$ and $N_0 = N_0(\kappa, \beta)$ such that, for all $N > N_0$, ρ_{τ} , the solution of Eq. 4, is 0 with probability $1 - 1/N^{\kappa}$.*

This theorem guarantees with large probability a zero false discovery rate in the absence of signals with meaningful

Significance

The ability to detect sparse signals from noisy, high-dimensional data is a top priority in modern science and engineering. For optimal results, current approaches need to tune parameters that depend on the level of noise, which is often difficult to estimate. We develop a parameter-free, computationally efficient, ℓ_1 -norm minimization approach that has a zero false discovery rate (no false positives) with high probability for any level of noise while it detects the exact location of sparse signals when the noise is not too large.

Author contributions: G.P. designed research; M.M., A.N., G.P., and C.T. performed research; and M.M., A.N., G.P., and C.T. wrote the paper.

Reviewers: A.C.F., University of California, Davis; and J.X., University of California, Irvine. The authors declare no competing interest.

This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: moscoso@math.uc3m.es or papanicolaou@stanford.edu.

First published May 11, 2020.

information. The key to a zero false discovery rate is the choice of a no-phantom weight τ . Next, we generalize this result for the case in which the recorded signals carry useful information.

Theorem 2 (Zero False Discoveries). *Let ρ be an M -sparse solution of the noiseless system $\mathcal{A}\rho = \mathbf{b}_0$. Assume that κ, β , the Noise Collector, and the noise are the same as in Theorem 1. In addition, assume that the columns of \mathcal{A} are incoherent in the sense that $|\langle \mathbf{a}_i, \mathbf{a}_j \rangle| \leq \frac{1}{3M}$. Then, there are constants $\tau = \tau(\kappa, \beta)$ and $N_0 = N_0(\kappa, \beta)$ such that $\text{supp}(\rho_\tau) \subseteq \text{supp}(\rho)$ for all $N > N_0$ with probability $1 - 1/N^\kappa$.*

This theorem holds for any level of noise and the same value of τ as in Theorem 1. The incoherence conditions in Theorem 2 are needed to guarantee that the true signal does not create false positives elsewhere. Theorem 2 guarantees that the support of ρ_τ is inside the support of ρ . The next theorem shows that, if the noise is not too large, then ρ_τ and ρ have exactly the same support.

Theorem 3 (Exact Support Recovery). *Keep the same assumptions as in Theorem 2. Let $\gamma = \min_{i \in \text{supp}(\rho)} |\rho_i| / \|\rho\|_{\ell_\infty}$. There are constants $\tau = \tau(\kappa, \beta)$, $c_1 = c_1(\kappa, \beta, \gamma)$, and $N_0 = N_0(\kappa, \beta)$ such that, if the noise level satisfies $\|e\|_{\ell_2} \leq c_1 \|\mathbf{b}_0\|_{\ell_2}^2 \|\rho\|_{\ell_1}^{-1} \sqrt{N} / \sqrt{\ln N}$, then for all $N > N_0$, $\text{supp}(\rho_\tau) = \text{supp}(\rho)$ with probability $1 - 1/N^\kappa$.*

To elucidate an interpretation of the last theorem, consider a model case where \mathcal{A} is the identity matrix and all coefficients of $\mathbf{b}_0 = \rho$ are either one or zero. Then, $\|\mathbf{b}_0\|_{\ell_2}^2 = \|\rho\|_{\ell_1} = M$. In this case, an acceptable relative level of noise is

$$\|e\|_{\ell_2} / \|\mathbf{b}_0\|_{\ell_2} \lesssim \sqrt{N} / \sqrt{M \ln N}. \quad [5]$$

This means that $\|e\|_{\ell_2} \lesssim \sqrt{N} / \sqrt{\ln N}$, and it implies that each coefficient of \mathbf{b}_0 may be corrupted by $O(1/\sqrt{\ln N})$ on average and that some coefficients of \mathbf{b}_0 may be corrupted by $O(1)$.

Motivation

We are interested in imaging sparse scenes accurately using limited and noisy data. Such imaging problems arise in many areas, such as medical imaging (12), structural biology (13), radar (14), and geophysics (15). In imaging, the ℓ_1 -norm minimization method in Eq. 2 is often used (16–21) as it has the desirable property of superresolution: that is, the enhancement of the fine-scale details of the images. This has been analyzed in different settings by Donoho (22), Candès and Fernandez-Granda (23), Fannjiang and Liao (24), and Borcea and Kocyigit (25) among others. We want to retain this property in our method when the data are corrupted by additive noise.

However, noise fundamentally limits the quality of the images formed with almost all computational imaging techniques. Specifically, ℓ_1 -norm minimization produces images that are unstable for low SNR due to the ill conditioning of superresolution reconstruction schemes. The instability emerges as clutter noise in the image, or grass, that degrades the resolution. Our initial motivation to introduce the Noise Collector matrix \mathcal{C} was to regularize the matrix \mathcal{A} and thus, to suppress the clutter in the images. We proposed in ref. 26 to seek the minimal ℓ_1 -norm solution of the augmented linear system $\mathcal{A}\rho + \mathcal{C}\eta = \mathbf{b}$. The idea was to choose the columns of \mathcal{C} almost orthogonal to those of \mathcal{A} . Indeed, the condition number of $[\mathcal{A} | \mathcal{C}]$ becomes $O(1)$ when $O(N)$ columns of \mathcal{C} are taken at random. This essentially follows from the bounds on the largest and the smallest nonzero singular values of random matrices (theorem 4.6.1 in ref. 27).

The idea to create a dictionary for noise is not new. For example, the work by Laska et al. (28) considers a specific version of the measurement noise model so that $\mathbf{b} = \mathcal{A}\rho + \mathcal{C}e$, where \mathcal{C} is a matrix with fewer (orthonormal) columns than rows and the

noise vector e is sparse. \mathcal{C} represents the basis in which the noise is sparse and it is assumed to be known. Then, they show that it is possible to recover sparse signals and sparse noise exactly. We stress that we do not assume here that the noise is sparse. In our work, the noise is large (SNR can be small) and is evenly distributed across the data, and therefore, it cannot be sparsely accommodated.

To suppress the clutter, our theory in ref. 26 required exponentially many columns, and therefore, $\Sigma \lesssim e^N$. This seemed to make the Noise Collector impractical, but the numerical experiments suggested that $O(N)$ columns were enough to obtain excellent results. We address this issue here and explain why the Noise Collector matrix \mathcal{C} only needs algebraically many columns. Moreover, to absorb the noise completely and thus, improve the algorithm in ref. 26, we introduce now the no-phantom weight τ in Eq. 4. Indeed, by weighting the columns of the Noise Collector matrix \mathcal{C} with respect to those in the model matrix \mathcal{A} , the algorithm now produces images with no clutter at all regardless of how much noise is added to the data.

Finally, we want the Noise Collector to be efficient, with almost no extra computational cost with respect to the Lasso problem in Eq. 3. To this end, the Noise Collector is constructed using circulant matrices that allow for efficient matrix vector multiplications using fast Fourier transforms (FFTs).

We now explain how the Noise Collector works and reduce our theorems to basic estimates in high-dimensional probability.

The Noise Collector

The method has two main ingredients: the Noise Collector matrix \mathcal{C} and the no-phantom weight τ . The construction of the Noise Collector matrix \mathcal{C} starts with the following three key properties. First, its columns should be sufficiently orthogonal to the columns of \mathcal{A} , and therefore, it does not absorb signals with “meaningful” information. Second, the columns of \mathcal{C} should be uniformly distributed on the unit sphere \mathbb{S}^{N-1} so that we could approximate well a typical noise vector. Third, the number of columns in \mathcal{C} should grow slower than exponential with N ; otherwise, the method is impractical.

One way to guarantee all three properties is to impose

$$|\langle \mathbf{a}_i, \mathbf{c}_j \rangle| < \frac{\alpha}{\sqrt{N}} \quad \forall i, j \quad \text{and} \quad |\langle \mathbf{c}_i, \mathbf{c}_j \rangle| < \frac{\alpha}{\sqrt{N}} \quad \forall i \neq j \quad [6]$$

with $\alpha > 1$ and fill out \mathcal{C} drawing \mathbf{c}_i at random with rejections until the rejection rate becomes too high. Then, by construction, the columns of \mathcal{C} are almost orthogonal to the columns of \mathcal{A} , and when the rejection rate becomes too high, this implies that we cannot pack more N -dimensional unit vectors into \mathcal{C} ; thus, we can approximate well a typical noise vector. Finally, the Kabatjanskii–Levenstein inequality (discussed in ref. 29) implies that the number Σ of columns in \mathcal{C} grows at most polynomially: $\Sigma \leq N^{\alpha^2}$. The first estimate in Eq. 6 implies that any solution $\mathcal{C}\eta = \mathbf{a}_i$ satisfies, for any $i \leq K$, $\|\eta\|_{\ell_1} \gtrsim \sqrt{N}$. This estimate measures how expensive it is to approximate columns of \mathcal{A} (i.e., the meaningful signal) with the Noise Collector. In turn, the no-phantom weight τ should be chosen so that it is expensive to approximate noise using columns of \mathcal{A} . It cannot be taken too large, however, because we may lose the signal. In fact, one can prove that, if $\tau \geq \sqrt{N}/\alpha$, then $\rho_\tau \equiv 0$ for any ρ and any level of noise. Intuitively, τ characterizes the rate at which the signal is lost as the noise increases. The most important property of the no-phantom weight τ is that it does not depend on the level of noise, and therefore, it is chosen before we start using the Noise Collector.

It is, however, more convenient for the proofs to use a probabilistic version of Eq. 6. Suppose that the columns of \mathcal{C} are drawn independently at random. Then, the dot product of any two

random unit vectors is still typically of order $1/\sqrt{N}$ (27). If the number of columns grows polynomially, we only have to sacrifice an asymptotically negligible event where our Noise Collector does not satisfy the three key properties, and the decoherence constraints in Eq. 6 are weakened by a logarithmic factor only. This follows from basic estimates in high-dimensional probability. We will state them in the next lemma after we interpret problem [4] geometrically.

Consider the convex hulls

$$H_1 = \left\{ x \in \mathbb{R}^N \mid x = \sum_{i=1}^{\Sigma} \xi_i \mathbf{c}_i, \sum_{i=1}^{\Sigma} |\xi_i| \leq 1 \right\}, \quad [7]$$

$$H_2 = \left\{ x \in \mathbb{R}^N \mid x = \sum_{i=1}^K \xi_i \mathbf{a}_i, \sum_{i=1}^{\Sigma} |\xi_i| \leq 1 \right\}, \quad [8]$$

and $H(\tau) = \{\xi h_1/\tau + (1-\xi)h_2, 0 \leq \xi \leq 1, h_i \in H_i\}$. Theorem 1 states that, for a typical noise vector \mathbf{e} , we can find $\lambda_0 > 0$ such that $\mathbf{e} \in \lambda_0 \partial H_1$ and $\mathbf{e} \notin \lambda \partial H(\tau)$ for any $\lambda < \lambda_0$.

Lemma 1 (Typical Width of Convex Hulls H_i). *Suppose that $\Sigma = N^\beta$, $\beta > 1$; vectors $\mathbf{c}_i \in \mathbb{S}^{N-1}$, $i = 1, 2, \dots, \Sigma$, are drawn at random and independently; and $\mathbf{e} \in \mathbb{S}^{N-1}$. Then, for any $\kappa > 0$, there are constants $c_0 = c_0(\kappa, \beta)$, $\alpha = \sqrt{(\beta-1)/2}$, and $N_0 = N_0(\kappa, \beta)$ such that, for all $N \geq N_0$,*

$$\max \left(\max_{i \leq K} (|\langle \mathbf{a}_i, \mathbf{e} \rangle|), \max_{i \leq \Sigma} (|\langle \mathbf{c}_i, \mathbf{e} \rangle|) \right) < c_0 \sqrt{\ln N} / \sqrt{N} \quad [9]$$

and

$$\alpha \sqrt{\ln N} \mathbf{e} / \sqrt{N} \in H_1 \quad [10]$$

with probability $1 - 1/N^\kappa$.

We sketch the proof of estimates [9] and [10] in *Proofs*. Estimate [9] can also be derived from Milman's version (30) of Dvoretzky's theorem. Informally, inequality [9] states that H_1 and H_2 are contained in the ℓ_2 ball of radius $c_0 \sqrt{\ln N} / \sqrt{N}$ except for a few spikes in statistically insignificant directions (Fig. 1, *Left*). Inequality [10] states that H_1 contains an ℓ_2 ball of radius $\alpha \sqrt{\ln N} / \sqrt{N}$ except for a few statistically insignificant directions.

These inequalities immediately imply Theorem 1. We just need to explain how to choose the no-phantom weight τ . There will be no phantoms if H_2/τ is strictly inside the ℓ_2 ball of radius $\alpha \sqrt{\ln N} / \sqrt{N}$. This could be done if $\tau > c_0/\alpha$.

If columns of \mathcal{A} are orthogonal to each other, then Theorem 2 follows from Theorem 1. We just need to project the linear system in Eq. 4 on the span of \mathbf{a}_i , $i \notin \text{supp}(\rho)$,

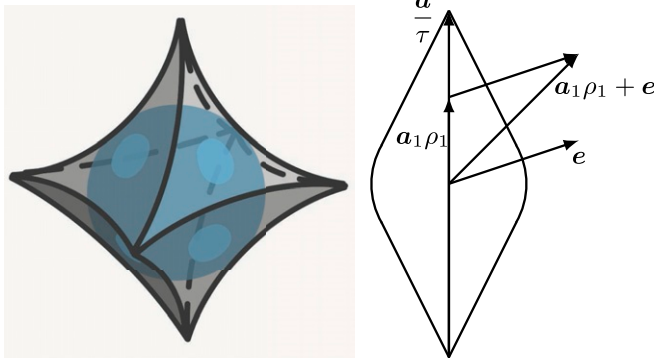


Fig. 1. (*Left*) A convex hull H_1 is an ℓ_2 ball of radius $O(\sqrt{\ln N}/\sqrt{N})$ with few spikes. (*Right*) An intersection of $H(\tau)$ with the span $(\mathbf{a}_1, \mathbf{e})$ is a rounded rhombus.

and apply Theorem 1 to the projections. If, in addition, we assume that $\mathbf{b}_0 = \mathbf{a}_1 \rho_1$, then Proof of Theorem 3 is illustrated in Fig. 1, *Right*. In detail, a typical intersection of $V = \text{span}(\mathbf{a}_1, \mathbf{e})$, and $H(\tau)$ is a rounded rhombus because it is the convex hull of \mathbf{a}_1/τ and the ℓ_2 ball of radius $c_0 \sqrt{\ln N} / \sqrt{N}$. If $\mathbf{a}_1 \rho_1 + \mathbf{e} \in \lambda_0 \partial H(\tau)$, then there are two options: 1) $\mathbf{a}_1 \rho_1 + \mathbf{e}$ lies on the curved boundary of the rounded rhombus, and then, $\text{supp}(\rho_\tau) = \emptyset$; or 2) $\mathbf{a}_1 \rho_1 + \mathbf{e}$ lies on the flat boundary of the rounded rhombus, and then, $\text{supp}(\rho_\tau) = \text{supp}(\rho)$. The second option happens if the vector $\mathbf{a}_1 \rho_1 + \mathbf{e}$ intersects the flat boundary of $\partial H(\tau)$. This gives the support recovery estimate in Theorem 3.

In the general case, the columns of the combined matrix $[\mathcal{A} \mid \mathcal{C}]$ are incoherent. This property allows us to prove Theorems 2 and 3 in *Proofs* using known techniques (26). In particular, we automatically have exact recovery using ref. 2 applied to $[\mathcal{A} \mid \mathcal{C}]$ if the data are noiseless.

Lemma 2 (Exact Recovery). *Suppose that ρ is an M -sparse solution of $\mathcal{A}\rho = \mathbf{b}$ and that there is no noise so that $\mathbf{e} = 0$. In addition, assume that the columns of \mathcal{A} are incoherent: $|\langle \mathbf{a}_i, \mathbf{a}_j \rangle| \leq \frac{1}{3M}$. Then, the solution to Eq. 4 satisfies $\rho_\tau = \rho$ for all*

$$M < \frac{2\sqrt{N}}{3c_0\tau\sqrt{\ln N}} \text{ with probability } 1 - \frac{1}{N^\kappa}. \quad [11]$$

Fast Noise Collector Algorithm

To find the minimizer in Eq. 4, we consider a variational approach. We define the function

$$F(\rho, \eta, \mathbf{z}) = \lambda (\tau \|\rho\|_{\ell_1} + \|\eta\|_{\ell_1}) + \frac{1}{2} \|\mathcal{A}\rho + \mathcal{C}\eta - \mathbf{b}\|_{\ell_2}^2 + \langle \mathbf{z}, \mathbf{b} - \mathcal{A}\rho - \mathcal{C}\eta \rangle \quad [12]$$

for a no-phantom weight τ and determine the solution as

$$\max_{\mathbf{z}} \min_{\rho, \eta} F(\rho, \eta, \mathbf{z}). \quad [13]$$

The key observation is that this variational principle finds the minimum in Eq. 4 exactly for all values of the regularization parameter λ . Hence, the method has no tuning parameters. To determine the exact extremum in Eq. 13, we use the iterative soft thresholding algorithm GeLMA (generalized Lagrangian multiplier algorithm) (31) that works as follows.

First, pick a value for β and τ . For optimal results, one can calibrate τ to be the smallest constant such that Theorem 1 holds: that is, we see no-phantom signals when the algorithm is fed with pure noise. In our numerical experiments, we use $\beta = 1.5$ and $\tau = 2$.

Second, pick a value for the regularization parameter λ (e.g., $\lambda = 1$). Choose step sizes $\Delta t_1 < 2/\|[\mathcal{A} \mid \mathcal{C}]\|^2$ and $\Delta t_2 < \lambda/\|\mathcal{A}\|$. Set $\rho_0 = \mathbf{0}$, $\eta_0 = \mathbf{0}$, and $\mathbf{z}_0 = \mathbf{0}$, and iterate for $k \geq 0$:

$$\begin{aligned} \mathbf{r} &= \mathbf{b} - \mathcal{A}\rho_k - \mathcal{C}\eta_k, \\ \rho_{k+1} &= \mathcal{S}_{\tau\lambda\Delta t_1}(\rho_k + \Delta t_1 \mathcal{A}^*(\mathbf{z}_k + \mathbf{r})), \\ \eta_{k+1} &= \mathcal{S}_{\lambda\Delta t_1}(\eta_k + \Delta t_1 \mathcal{C}^*(\mathbf{z}_k + \mathbf{r})), \\ \mathbf{z}_{k+1} &= \mathbf{z}_k + \Delta t_2 \mathbf{r}, \end{aligned} \quad [14]$$

where $\mathcal{S}_r(y_i) = \text{sign}(y_i) \max\{0, |y_i| - r\}$.

The Noise Collector matrix \mathcal{C} is computed by drawing $N^{\beta-1}$ normally distributed N -dimensional vectors normalized to unit length. These are the generating vectors of the Noise Collector. From each of them, a circulant $N \times N$ matrix \mathcal{C}_i ,

*Choosing two step sizes instead of the smaller one Δt_1 improves the convergence speed.




Fig. 2. General setup for passive array imaging. The source at \vec{z}_j emits a signal that is recorded at all array elements $\vec{x}_r, r = 1, \dots, N_r$.

$i = 1, \dots, N^{\beta-1}$, is constructed. The Noise Collector matrix is obtained by concatenation, and therefore, $\mathcal{C} = [\mathcal{C}_1 | \mathcal{C}_2 | \dots | \mathcal{C}_{N^{\beta-1}}]$. Exploiting the circulant structure of the matrices \mathcal{C}_i , we perform the matrix vector multiplications $\mathcal{C}\eta_k$ and $\mathcal{C}^*(z_k + r)$ in Eq. 14 using the FFT (32). This makes the complexity associated with the Noise Collector $O(N^\beta \log(N))$. Note that only the $N^{\beta-1}$ generating vectors are stored and not the entire $N \times N^\beta$ Noise Collector matrix. In practice, we use $\beta \approx 1.5$, which makes the cost of using the Noise Collector negligible as typically, $K \gg N^{\beta-1}$. The columns of the Noise Collector matrix \mathcal{C} with this circulant structure are uniformly distributed on \mathbb{S}^{N-1} , and they satisfy Lemma 1. This implies that the theorems of this paper are still valid for such \mathcal{C} .

Application to Imaging

We consider passive array imaging of point sources. The problem consists of determining the positions \vec{z}_j and the complex[†] amplitudes $\alpha_j, j = 1, \dots, M$, of a few point sources from measurements of polychromatic signals on an array of receivers (Fig. 2). The imaging system is characterized by the array aperture a , the distance L to the sources, the bandwidth B , and the central wavelength λ_0 .

The sources are located inside an image window (IW), which is discretized with a uniform grid of points $\vec{y}_k, k = 1, \dots, K$. The unknown is the source vector $\rho = [\rho_1, \dots, \rho_K]^T \in \mathbb{C}^K$, with components ρ_k that correspond to the complex amplitudes of the M sources at the grid points $\vec{y}_k, k = 1, \dots, K$, with $K \gg M$. For the true source vector, we have $\rho_k = \alpha_j$ if $\vec{y}_k = \vec{z}_j$ for some $j = 1, \dots, M$, while $\rho_k = 0$ otherwise.

Denoting by $G(\vec{x}; \vec{y}; \omega)$ Green's function for the propagation of a signal of angular frequency ω from point \vec{y} to point \vec{x} , we define the single-frequency Green's function vector that connects a point \vec{y} in the IW with all points $\vec{x}_r, r = 1, \dots, N_r$, on the array as

$$\mathbf{g}(\vec{y}; \omega) = [G(\vec{x}_1, \vec{y}; \omega), G(\vec{x}_2, \vec{y}; \omega), \dots, G(\vec{x}_N, \vec{y}; \omega)]^T \in \mathbb{C}^{N_r}.$$

In three dimensions, $G(\vec{x}, \vec{y}; \omega) = \frac{\exp\{i\omega|\vec{x} - \vec{y}|\}/c_0}{4\pi|\vec{x} - \vec{y}|}$ if the medium is homogeneous. The data for the imaging problem are the signals $b(\vec{x}_r, \omega_l) = \sum_{j=1}^M \alpha_j G(\vec{x}_r, \vec{z}_j; \omega_l)$ recorded at receiver

locations $\vec{x}_r, r = 1, \dots, N_r$, at frequencies $\omega_l, l = 1, \dots, S$. These data are stacked in a column vector

$$\mathbf{b} = [b(\omega_1)^T, b(\omega_2)^T, \dots, b(\omega_S)^T]^T \in \mathbb{C}^N; \quad N = N_r S, \quad [15]$$

with $b(\omega_l) = [b(\vec{x}_1, \omega_l), b(\vec{x}_2, \omega_l), \dots, b(\vec{x}_N, \omega_l)]^T \in \mathbb{C}^{N_r}$. Then, $\mathcal{A}\rho = \mathbf{b}$, with \mathcal{A} being the $N \times K$ measurement matrix with columns \mathbf{a}_k that are the multiple-frequency Green's function vectors

$$\mathbf{a}_k = [g(\vec{y}_k; \omega_1)^T, g(\vec{y}_k; \omega_2)^T, \dots, g(\vec{y}_k; \omega_S)^T]^T \in \mathbb{C}^N \quad [16]$$

normalized to have length 1. The system $\mathcal{A}\rho = \mathbf{b}$ relates the unknown vector $\rho \in \mathbb{C}^K$ to the data vector $\mathbf{b} \in \mathbb{C}^N$.

Next, we illustrate the performance of the Noise Collector in this imaging setup. The most important features are that 1) no calibration is necessary with respect to the level of noise, that 2) exact support recovery is obtained for relatively large levels of noise [i.e., $\|e\|_{\ell_2} \leq c_1 \|b_0\|_{\ell_2}^2 \sqrt{N} / (\|\rho\|_{\ell_1} \sqrt{\ln N})$], and that 3) we have zero false discovery rates for all levels of noise with high probability.

We consider a high-frequency microwave imaging regime with central frequency $f_0 = 60$ GHz corresponding to $\lambda_0 = 5$ mm. We make measurements for $S = 25$ equally spaced frequencies spanning a bandwidth $B = 20$ GHz. The array has $N = 25$ receivers and an aperture $a = 50$ cm. The distance from the array to the center of the imaging window is $L = 50$ cm. Then, the resolution is $\lambda_0 L/a = 5$ mm in the cross-range (direction parallel to the array) and $c_0/B = 15$ mm in range (direction of propagation). These parameters are typical in microwave scanning technology (33).

We seek to image a source vector with sparsity $M = 12$ (Fig. 3, Left). The size of the imaging window is 20×60 cm, and the pixel spacing is 5×15 mm. The number of unknowns is, therefore, $K = 1,681$, and the number of data is $N = 625$. The size of the Noise Collector is taken to be $\Sigma = 10^4$, and therefore, $\beta \approx 1.5$. When the data are noiseless, we obtain exact recovery as expected (Fig. 3, Right).

In Fig. 4, we display the imaging results with and without the Noise Collector when the data are corrupted by additive noise. The SNR = 1, and therefore, the ℓ_2 norms of the signals and the noise are equal. In column 1 of Fig. 4, we show the recovered image using ℓ_1 -norm minimization without the Noise Collector. There is a lot of grass in this image, with many nonzero values outside the true support. When the Noise Collector is used, the level of the grass is reduced, and the image improves (column 2 of Fig. 4). Still, there are several false discoveries because we use $\tau = 1$ in algorithm [14].

In column 3 of Fig. 4, we show the image obtained with a weight $\tau = 2$ in algorithm [14]. With this weight, there are no false discoveries, and the recovered support is exact. This simplifies the imaging problem dramatically as we can now restrict the




Fig. 3. Noiseless data. The exact solution is recovered for any value of λ in algorithm [14]. (Left) The true image. (Right) The recovered solution vector, ρ_r , is plotted with red stars, and the true solution vector, ρ , is plotted with green circles.

[†]We chose to work with real numbers in the previous sections for ease of presentation, but the results also hold with complex numbers.

inverse problem to the true support just obtained and then, solve an overdetermined linear system using a classical ℓ_2 approach. The results are shown in column 4 of Fig. 4. Note that this second step largely compensates for the signal that was lost in the first step due to the high level of noise.

In Fig. 5, we illustrate the performance of the Noise Collector for different sparsity levels M and $\|e\|_{\ell_2}/\|b_0\|_{\ell_2}$ values. Success in recovering the true support of the unknown corresponds to a value of one (yellow in Fig. 5), and failure corresponds to a value of zero (blue in Fig. 5). The small phase transition zone (green in Fig. 5) contains intermediate values. The black lines in Fig. 5 are the theoretical prediction Eq. 5. These results are obtained by averaging over 10 realizations of noise. We show results for three values of data sizes $N = 342$, $N = 625$, and $N = 961$. In our experiments, the nonzero components of the unknown ρ take values in $[0.6, 0.8]$, and therefore, $\|b_0\|_{\ell_2}/\|\rho\|_{\ell_1} = cst/\sqrt{M}$.

Remark 1: We considered passive array imaging for ease of presentation. The same results hold for active array imaging with or without multiple scattering; ref. 34 discusses the detailed analytical setup.

Remark 2: We have considered a microwave imaging regime. Similar results can be obtained in other regimes.

Proofs

Proof of Lemma 1: Using the rotational invariance of all of our probability distributions, inequality [9] is true if

$$\mathbb{P}(\max_i |\langle d_i, e \rangle| \geq c_0 \sqrt{\ln N} / \sqrt{N}) \leq 1/N^\kappa,$$

where $d_i, i = 1, 2, \dots, K + \Sigma$ are (possibly dependent) uniformly distributed on \mathbb{S}^{N-1} , and we can assume that $e = (1, 0, \dots, 0)$. Denote the event

$$\Omega_t = \left\{ \max_i |\langle d_i, e \rangle| \geq t/\sqrt{N} \right\}.$$

$\mathbb{P}(|\langle d_i, e \rangle| \geq t/\sqrt{N}) \leq 2 \exp(-t^2/2)$ for each d_i . We obtain $\mathbb{P}(\Omega_t) \leq 2(K + \Sigma) \exp(-t^2/2)$ using the union bound. Choosing $t = c_0 \sqrt{\ln N}$ for sufficiently large c_0 , we get $\mathbb{P}(\Omega_t) \leq$

$CN^\beta N^{-c_0^2/2} \leq N^{-\kappa}$, where $c_0^2 > 2(\beta + \kappa)$ and $N \geq N_0$. Hence, Eq. 9 holds with probability $1 - N^{-\kappa}$.

If N columns $c_j, j \in S$ of \mathcal{C} satisfy

$$\min_{j \in S} |\langle c_j, e \rangle| \geq \theta, \theta = \alpha \sqrt{\ln N} / \sqrt{N}, \quad [17]$$

then their convex hull will contain θe with probability $(1/2)^N$. Therefore, inequality [10] follows if [17] holds with probability $1 - 1/N^\kappa$. Using the rotational invariance of all of our probability distributions, we can assume that $e = (1, 0, \dots, 0)$. For each c_i ,

$$\mathbb{P}\left(|\langle c_i, e \rangle| \geq \frac{t}{\sqrt{N}}\right) = \frac{2}{\sqrt{2\pi}} \int_t^\infty e^{-\frac{x^2}{2}} dx \geq \frac{1}{2} e^{-t^2}.$$

Split the index set $1, 2, \dots, \Sigma$ into N nonoverlapping subsets $S_k, k = 1, 2, \dots, N$ of size $N^{\beta-1}$. For each S_k ,

$$\mathbb{P}\left(\max_{i \in S_k} |\langle c_i, e \rangle| \leq \frac{\alpha \sqrt{\ln N}}{\sqrt{N}}\right) \leq \left(1 - \frac{1}{2N^{\alpha^2}}\right)^{N^{\beta-1}} \leq e^{-\frac{1}{2}N^{\frac{\beta-1}{2}}}$$

for $\alpha = \sqrt{(\beta-1)/2}$. By independence,

$$\mathbb{P}([17] \text{ holds}) \geq \prod_{k=1}^N \mathbb{P}(\max_{i \in S_k} |\langle c_i, e \rangle| \geq \alpha \sqrt{\ln N} / \sqrt{N}).$$

Then, $\mathbb{P}([17] \text{ holds}) \geq (1 - e^{-\frac{1}{2}N^{\frac{\beta-1}{2}}})^N \geq 1 - Ne^{-\frac{1}{2}N^{\frac{\beta-1}{2}}}$. Choosing N_0 sufficiently large, we obtain [10]. \square

Proof of Theorem 2: When columns of \mathcal{A} are not orthogonal, we will choose a τ smaller than that in Theorem 1 by a factor of two. Suppose that the M -dimensional space V is the span of the column vectors a_j , with j in the support of ρ . Say that V is spanned by $a_1 \dots a_M$. Let $W = V^\perp$ be the orthogonal complement to V . Consider the orthogonal decomposition $a_i = a_i^v + a_i^w$ for all $i \geq M+1$. Incoherence of a_i implies that $\|a_i^w\|_{\ell_2} \geq 1/2$ for all $i \geq M+1$. Indeed, fix any $i \geq M+1$. Suppose that $a_i^v = \sum_{k=1}^M \xi_k a_k$ and that $|\xi_j| = \max_{k \leq M} |\xi_k| = \|\xi\|_{l_\infty}$. Thus, $\frac{1}{3M} \geq |\langle a_j, a_i^v \rangle| \geq |\langle a_j, \sum_{k=1}^M \xi_k a_k \rangle| \geq \|\xi\|_{l_\infty} (1 - \frac{M-1}{3M})$. Then,

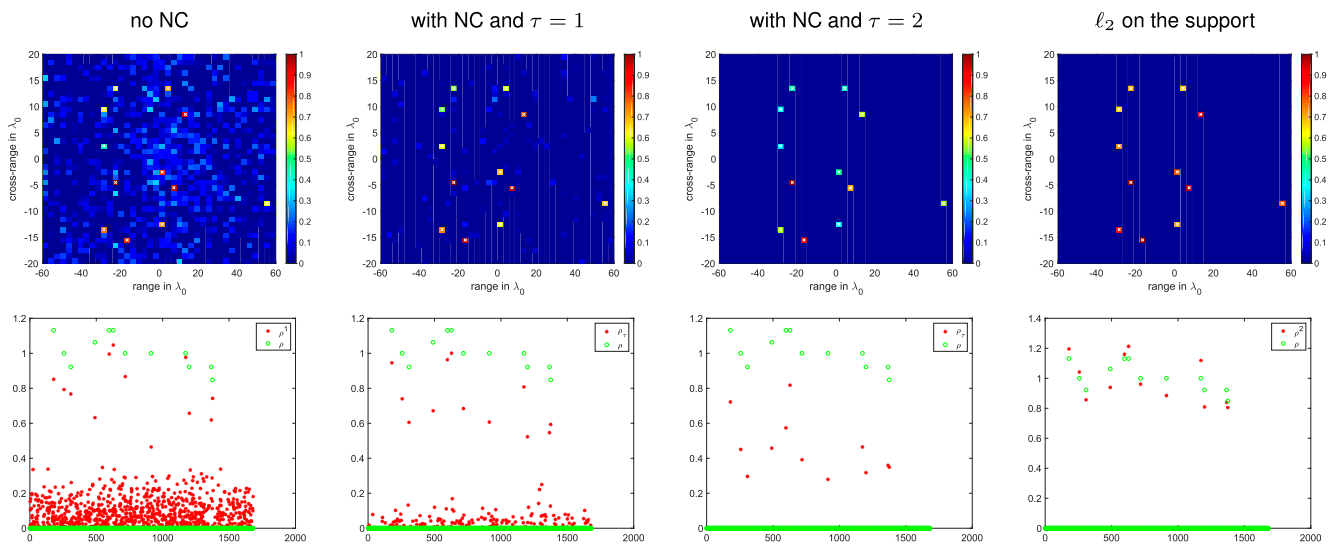


Fig. 4. High level of noise; SNR = 1. (Column 1) ℓ_1 -norm minimization without the noise collector. (Column 2) ℓ_1 -norm minimization with a noise collector with $\Sigma = 10^4$ columns and $\tau = 1$ in algorithm [14]. (Column 3) ℓ_1 -norm minimization with a noise collector and the correct $\tau = 2$ in algorithm [14]. (Column 4) ℓ_2 -norm solution restricted to the support. In *Upper*, we show the images. In *Lower*, we show the solution vector with red stars and the true solution vector with green circles. NC = noise collector.




Fig. 5. (Left) Algorithm performance for exact support recovery. Success corresponds to a value of one (yellow), and failure corresponds to a value of zero (blue). The small phase transition zone (green) contains intermediate values. The black lines are the theoretical estimate $\sqrt{N}/\sqrt{M \ln N}$. Ordinate and abscissa are the sparsity M and $\|e\|_2/\|b_0\|_2$. The data sizes are (Left) $N = 342$, (Center) $N = 625$, and (Right) $N = 961$.

$\|\xi\|_{l_\infty} \leq 1/(2M)$. Therefore, $\|a_i^v\|_{l_2} \leq \|\xi\|_{l_1} \leq M\|\xi\|_{l_\infty} \leq 1/2$, and $\|a_i^w\|_{l_2} \geq \|a_i\|_{l_2} - \|a_i^v\|_{l_2} \geq 1/2$.

Project system [4] on W . Then, we obtain a new system [4]. The l_2 norms of the columns of new \mathcal{A} are at least $1/2$. Otherwise, the new system satisfies all conditions of Theorem 1. Indeed, b_0 is projected to zero. All c_i and $e/\|e\|_{l_2}$ are projected to vectors uniformly distributed on S^{N-M-1} by the concentration of measure (27). If any a_i , $i \geq M+1$, was used in an optimal approximation of $b_0 + e$, then its projection a_i^w is used in an optimal approximation of the projection of $b_0 + e$ on W . This is a contradiction to Lemma 1 if we choose $\tau < c_0/(2\alpha)$ and recall that $\|a_i^w\|_{l_2} \geq 1/2$. \square

Proof of Theorem 3: Choose τ as in Theorem 2. Incoherence of \mathcal{A} implies that we can argue as in Proof of Theorem 2 and assume that $\langle a_i, a_j \rangle = 0$ for $i \neq j$, $i, j \in \text{supp}(\rho)$. Suppose that V^i are the two-dimensional (2D) spaces spanned by e and a_i for $i \in \text{supp}(\rho)$. By Lemma 1, all $\lambda H(\tau) \cap V^i$ look like the rounded rhombi depicted in Fig. 1, Right, and $\lambda H_1 \cap V^i \subset B_\lambda^i$ with probability $1 - N^{-\kappa}$, where B_λ^i is a 2D l_2 ball of radius $\lambda c_0 \sqrt{\ln N}/\sqrt{N}$. Thus, $\lambda H(\tau) \cap V^i \subset H_\lambda^i$ with probability $1 - N^{-\kappa}$, where H_λ^i is the convex hull of B_λ^i and a vector λf_i , $f_i = \rho_i \|\rho\|_{l_1}^{-1} \tau^{-1} a_i$. Then, $\text{supp}(\rho_\tau) = \text{supp}(\rho)$ if there exists λ_0 so that $\rho_i a_i + e$ lies on the flat boundary of $H_{\lambda_0}^i$ for all $i \in \text{supp}(\rho)$.

If $\min_{i \in \text{supp}(\rho)} |\rho_i| \geq \gamma \|\rho\|_\infty$, then there is a constant $c_2 = c_2(\gamma)$ such that, if $\rho_i a_i + e$ lies on the flat boundary of H_λ^i for

some i and some λ , then there exists λ_0 so that $\rho_i a_i + c_2 e$ lies on the flat boundary of $H_{\lambda_0}^i$ for all $i \in \text{supp}(\rho)$. Suppose that V is spanned by e and b_0 , $H_\lambda \subset V$ is the convex hull of B_λ and λf , and $f = b_0 \|\rho\|_{l_1}^{-1} \tau^{-1}$ where $B_\lambda \subset V$ is an l_2 ball of radius $\lambda c_0 \sqrt{\ln N}/\sqrt{N}$. If $b_0 + c_2 e$ lies on the flat boundary of H_λ , then there must be an $i \in \text{supp}(\rho)$ such that $\rho_i a_i + c_2 e$ lies on the flat boundary of H_λ^i . If

$$\frac{|\langle b_0, b_0 + c_2 e \rangle|}{\|b_0\|_{l_2} \|b_0 + c_2 e\|_{l_2}} \geq \frac{c_0 \sqrt{\ln N}}{\sqrt{N} \|f\|_{l_2}}, \quad [18]$$

then $b_0 + c_2 e$ lies on the flat boundary of H_λ . Since $|\langle b_0, e \rangle| \leq c_0 \|e\|_{l_2} \|b_0\|_{l_2}/\sqrt{N}$ with probability $1 - N^{-\kappa}$, Eq. 18 holds if $\|e\|_{l_2}/\|b_0\|_{l_2} \leq \|f\|_{l_2} \sqrt{N}/(c_2 c_0 \sqrt{\ln N}) \leq c_1 \|b_0\|_{l_2} \|\rho\|_{l_1}^{-1} \sqrt{N}/\sqrt{\ln N}$. \square

Data Availability Statement. There are no data in this paper; we present an algorithm, its theoretical analysis, and some numerical simulations.

ACKNOWLEDGMENTS. The work of M.M. was partially supported by Spanish Ministerio de Ciencia e Innovación Grant FIS2016-77892-R. The work of A.N. was partially supported by NSF Grants DMS-1515187 and DMS-1813943. The work of G.P. was partially supported by Air Force Office of Scientific Research (AFOSR) Grant FA9550-18-1-0519. The work of C.T. was partially supported by AFOSR Grants FA9550-17-1-0238 and FA9550-18-1-0519. We thank Marguerite Novikov for drawing Fig. 1, Left.

- S. S. Chen, D. L. Donoho, M. A. Saunders, Atomic decomposition by basis pursuit. *SIAM Rev.* **43**, 12–159 (2001).
- D. L. Donoho, M. Elad, Optimally sparse representation in general (nonorthogonal) dictionaries via l_1 minimization. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 2197–2202 (2003).
- E. J. Candès, T. Tao, Decoding by linear programming. *IEEE Trans. Inf. Theory* **51**, 4203–4215 (2005).
- R. Tibshirani, Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B* **58**, 267–288 (1996).
- J. J. Fuchs, Recovery of exact sparse representations in the presence of bounded noise. *IEEE Trans. Inf. Theory* **51**, 3601–3608 (2005).
- J. A. Tropp, Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE Trans. Inf. Theory* **52**, 1030–1051 (2006).
- M. J. Wainwright, Sharp thresholds for high-dimensional and noisy sparsity recovery using l_1 -constrained quadratic programming (Lasso). *IEEE Trans. Inf. Theory* **55**, 2183–2202 (2009).
- A. Maleki, L. Anitori, Z. Yang, R. Baraniuk, Asymptotic analysis of complex Lasso via complex approximate message passing (CAMP). *IEEE Trans. Inf. Theory* **59**, 4290–4308 (2013).
- H. Zou, The adaptive Lasso and its oracle properties. *J. Am. Stat. Assoc.* **101**, 1418–1429 (2006).
- J. N. Sampson, N. Chatterjee, R. J. Carroll, S. Müller, Controlling the local false discovery rate in the adaptive Lasso. *Biostatistics* **14**, 653–666 (2013).
- A. Belloni, V. Chernozhukov, L. Wang, Square-root Lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika* **98**, 791–806 (2011).
- J. Trzasko, A. Manduca, Highly undersampled magnetic resonance image reconstruction via homotopic l_0 -minimization. *IEEE Trans. Med. Imag.* **28**, 106–121 (2009).
- M. AlQuraishi, H. H. McAdams, Direct inference of protein dna interactions using compressed sensing methods. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 14819–14824 (2011).
- R. Baraniuk, P. Steeghs, “Compressive radar imaging” in *2007 IEEE Radar Conference* (IEEE, 2007), pp. 128–133.
- H. L. Taylor, S. C. Banks, J. F. McCoy, Deconvolution with the l_1 norm. *Geophysics* **44**, 39–52 (1979).
- D. Malioutov, M. Cetin, A. S. Willsky, A sparse signal reconstruction perspective for source localization with sensor arrays. *IEEE Trans. Signal Process.* **53**, 3010–3022 (2005).
- J. Romberg, Imaging via compressive sampling. *IEEE Signal Process. Mag.* **25**, 14–20 (2008).
- M. A. Herman, T. Strohmer, High-resolution radar via compressed sensing. *IEEE Trans. Signal Process.* **57**, 2275–2284 (2009).
- J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, R. G. Baraniuk, Beyond Nyquist: Efficient sampling of sparse bandlimited signals. *IEEE Trans. Inf. Theory* **56**, 520–544 (2010).
- A. C. Fannjiang, T. Strohmer, P. Yan, Compressed remote sensing of sparse objects. *SIAM J. Imag. Sci.* **3**, 595–618 (2010).
- A. Chai, M. Moscoso, G. Papanicolaou, Robust imaging of localized scatterers using the singular value decomposition and l_1 optimization. *Inverse Probl.* **29**, 025016 (2013).
- D. L. Donoho, Superresolution via sparsity constraints. *SIAM J. Math. Anal.* **23**, 1303–1331 (1992).
- E. J. Candès, C. Fernandez-Granda, Towards a mathematical theory of super-resolution. *Comm. Pure Appl. Math.* **67**, 906–956 (2014).
- A. C. Fannjiang, W. Liao, Coherence pattern-guided compressive sensing with unresolved grids. *SIAM J. Imag. Sci.* **5**, 179–202 (2012).
- L. Borcea, I. Kocyyigit, Resolution analysis of imaging with l_1 optimization. *SIAM J. Imag. Sci.* **8**, 3015–3050 (2015).
- M. Moscoso, A. Novikov, G. Papanicolaou, C. Tsogka, Imaging with highly incomplete and corrupted data. *Inverse Probl.* **36**, 035010 (2020).
- R. Vershynin, *High-Dimensional Probability: An Introduction with Applications in Data Science* (Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 2018).

28. J. N. Laska, M. A. Davenport, R. G. Baraniuk, "Exact signal recovery from sparsely corrupted measurements through the pursuit of justice" in *2009 Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers* (IEEE, Pacific Grove, CA, 2009), pp. 1556–1560.
29. T. Tao, "A cheap version of the Kabatjanskii-Levenstein bound for almost orthogonal vectors." [https://terrytao.wordpress.com/2013/07/18/a-cheap-version-of-the-kabatjanskii-levenstein-bound-for-almost-orthogonal-vectors.](https://terrytao.wordpress.com/2013/07/18/a-cheap-version-of-the-kabatjanskii-levenstein-bound-for-almost-orthogonal-vectors/) Accessed 23 January 2019.
30. V. D. Milman, A new proof of Dvoretzky's theorem on cross-sections of convex bodies. *Funkcional. Anal. i Priložen.* **5**, 28–37 (1971).
31. M. Moscoso, A. Novikov, G. Papanicolaou, L. Ryzhik, A differential equations approach to l_1 -minimization with applications to array imaging. *Inverse Probl.* **28**, 10 (2012).
32. R. M. Gray, Toeplitz and circulant matrices: A review. *Commun. Inf. Theory* **2**, 155–239 (2006).
33. J. Laviada, A. Arboleya-Arboleya, Y. Alvarez-Lopez, C. Garcia-Gonzalez, F. Las-Heras, Phaseless synthetic aperture radar with efficient sampling for broadband near-field imaging: Theory and validation. *IEEE Trans. Antennas Propag.* **63**, 573–584 (2015).
34. A. Chai, M. Moscoso, G. Papanicolaou, Imaging strong localized scatterers with sparsity promoting optimization. *SIAM J. Imaging Sci* **7**, 1358–1387 (2014).