

Dense Error Correction via ℓ^1 -Minimization

John Wright, *Member*, and Yi Ma, *Senior Member*.

Abstract—This paper studies the problem of recovering a sparse signal $x \in \mathbb{R}^n$ from highly corrupted linear measurements $y = Ax + e \in \mathbb{R}^m$, where e is an unknown error vector whose nonzero entries may be unbounded. Motivated by an observation from face recognition in computer vision, this paper proves that for highly correlated (and possibly overcomplete) dictionaries A , any sufficiently sparse signal x can be recovered by solving an ℓ^1 -minimization problem:

$$\min \|x\|_1 + \|e\|_1 \quad \text{subject to} \quad y = Ax + e.$$

More precisely, if the fraction of the support of the error e is bounded away from one and the support of x is a very small fraction of the dimension m , then as m becomes large the above ℓ^1 -minimization succeeds for all signals x and almost all sign-and-support patterns of e . This result suggests that accurate recovery of sparse signals is possible and computationally feasible even with nearly 100% of the observations corrupted. The proof relies on a careful characterization of the faces of a convex polytope spanned together by the standard crosspolytope and a set of iid Gaussian vectors with nonzero mean and small variance, dubbed the “cross-and-bouquet” model. Simulations and experiments corroborate the findings, and suggest extensions to the result.

Index Terms—Dense error correction, sparse signal recovery, measure concentration, Gaussian matrices, ℓ^1 -minimization, polytope neighborliness.

I. INTRODUCTION

Recovery of high-dimensional sparse signals or errors has been one of the fastest growing research areas in signal processing in the past few years. At least two factors have contributed to this explosive growth. On the theoretical side, the progress has been propelled by powerful tools and results from multiple mathematical areas such as measure concentration [1]–[3], statistics [4]–[6], combinatorics [7], and coding theory [8]. On the practical side, a lot of excitement has been generated by remarkable successes in real-world applications in areas such as signal (image or speech) processing [9], communications [10], computer vision and pattern recognition [11]–[13] and more.

A. A motivating example

One notable, and somewhat surprising, successful application of sparse representation is automatic face recognition. The

basic setting of face recognition is as follows. For each person, a set of training images are taken under different illuminations. We can view each image as a vector by stacking its columns and put all the training images as column vectors of a matrix, say $A \in \mathbb{R}^{m \times n}$. Then, m is the number of pixels in an image and n is the total number of images for all the subjects of interest. Given a new query image, again we can stack it as a vector $y \in \mathbb{R}^m$. To identify which subject the image belongs to, we can try to represent y as a linear combination of all the images, i.e., $y = Ax$ for some $x \in \mathbb{R}^n$. Since in practice n can potentially be larger than m , the equations can be underdetermined and the solution x may not be unique. In this context, it is natural to seek the sparsest solution for x whose large non-zero coefficients then provide information about the subject’s true identity. This can be done by solving the typical ℓ^1 -minimization problem:

$$\min_x \|x\|_1 \quad \text{subject to} \quad y = Ax. \quad (1)$$

The problem becomes more interesting if the query image y is severely occluded or corrupted, as shown in Figure 1 left, column (a). In this case, one needs to solve a corrupted set of linear equations $y = Ax + e$, where $e \in \mathbb{R}^m$ is an unknown vector whose nonzero entries correspond to the corrupted pixels. For sparse errors e and tall matrices A ($m > n$), Candes and Tao [14] proposed to multiply the equation $y = Ax + e$ with a matrix B such that $BA = 0$, and then use ℓ^1 -minimization to recover the error vector e from the new linear equation $By = Be$.

As we mentioned earlier, in face recognition (and many other applications), n can be larger than m and the matrix A can be full rank. One cannot directly apply the above technique even if the error e is known to be very sparse. To resolve this difficulty, in [11], the authors proposed to instead seek $[x, e]$ together as the sparsest solution to the extended equation $y = [A \ I]w$ with $w = \begin{bmatrix} x \\ e \end{bmatrix} \in \mathbb{R}^{m+n}$, by solving the extended ℓ^1 -minimization problem:

$$\min_{x,e} \|x\|_1 + \|e\|_1 \quad \text{subject to} \quad y = Ax + e. \quad (2)$$

This seemingly minor modification to the previous error correction approach has drastic consequences for the performance of robust face recognition. Solving the modified ℓ^1 -minimization enables almost perfect recognition even with more than 60% of the pixels of the query image arbitrarily corrupted (see Figure 1 for an example), far beyond the amount of error that can theoretically be corrected by the previous error correction method [14].

Although ℓ^1 -minimization is expected to recover sufficiently sparse solutions with overwhelming probability for general systems of linear equations (see [16]), it is rather surprising that it works for the equation $y = Ax + e$ at all. In the

Manuscript received September 1, 2008; revised July 27, 2009; accepted September 29, 2009. J. Wright is with the Visual Computing Group, Microsoft Research Asia. Y. Ma is with the Visual Computing Group, Microsoft Research Asia, as well as the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. Corresponding author: John Wright, 146 Coordinated Science Laboratory, 1308 West Main Street, Urbana, Illinois 61801, USA. Email: jnwright@uiuc.edu. This work was partially supported by grants NSF CRS-EHS-0509151, NSF CCF-TF-0514955, ONR YIP N00014-05-1-0633, and NSF IIS 07-03756. John Wright was also supported by a Microsoft Fellowship (sponsored by Microsoft Live Labs, Redmond).


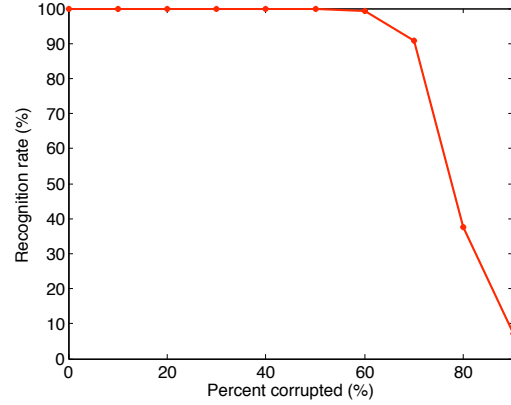


Fig. 1. **Face recognition under random corruption.** Left: (a) Test images \mathbf{y} with random corruption from the database presented in [15]. Top row: 30% of pixels are corrupted, Middle row: 50% corrupted, Bottom row: 70% corrupted. (b) Estimated errors $\hat{\mathbf{e}}$. (c) Estimated sparse coefficients $\hat{\mathbf{x}}$. (d) Reconstructed images $\mathbf{y}_r = A\hat{\mathbf{x}}$. The extended ℓ^1 -minimization (2) correctly recovers and identifies all three corrupted face images. Right: The recognition rate across the entire range of corruption for all the 38 subjects in the database. It performs almost perfectly upto 60% random corruption.



application described above, the columns of A are highly correlated. As m becomes large (i.e., the resolution of the image becomes high), the convex hull spanned by all face images of all subjects is only an extremely tiny portion of the unit sphere \mathbb{S}^{m-1} . For example, the images in Figure 1 lie on $\mathbb{S}^{8,063}$. The smallest inner product with their normalized mean is 0.723; they are contained within a spherical cap of volume $\leq 1.47 \times 10^{-229}$. These vectors are tightly bundled together as a “bouquet,” whereas the vectors associated with the identity matrix and its negative $\pm \mathbf{I}$ together¹ form a standard “cross” in \mathbb{R}^m , as illustrated in Figure 2. Notice that such a “cross-and-bouquet” matrix $[A \ \mathbf{I}]$ is neither incoherent nor (restrictedly) isometric, at least not uniformly. Also, the density of the desired solution \mathbf{w} is not uniform either. The \mathbf{x} part of \mathbf{w} is usually a very sparse² vector, but the \mathbf{e} part can be very dense. Existing results for recovering sparse signals suggest that ℓ^1 -minimization may have difficulty in dealing with such signals, contrary to its empirical success in face recognition.

We have experimented with similar cross-and-bouquet type models where the matrix A is a random matrix with highly correlated column vectors. The simulation results in Section IV indicate that what we have seen in face recognition is not an isolated phenomenon. In fact, the simulations reveal something even more striking and puzzling: As the dimension m increases (and the sample size n grows in proportion), the percentage of errors that the ℓ^1 -minimization (2) can correct seems to approach 100%. At first sight, this may seem rather surprising, but this paper explains why this should be expected when the space dimension is high.

B. Assumptions and the main result

Motivated by the face recognition example above, this paper aims to resolve this apparent discrepancy between theory and

¹Here, we allow the entries of the error \mathbf{e} to assume either positive or negative signs.

²Ideally, \mathbf{x} is concentrated only on images of the same subject. In the standard face recognition scenario where the number of subjects in the database grows, but the number of images per subject is fixed, the desired sparsity of \mathbf{x} is actually a constant.




Fig. 2. **The “cross-and-bouquet” model.** Left: the bouquet A and the crosspolytope spanned by the matrix $\pm \mathbf{I}$. Right: the tip of the bouquet magnified; it is a collection of i.i.d. Gaussian vectors with small variance σ^2 and common mean vector $\boldsymbol{\mu}$. The cross-and-bouquet polytope is spanned by vertices from both the bouquet A and the cross $\pm \mathbf{I}$.

practice of ℓ^1 -minimization. To this end, we need to give a more careful characterization of its behavior in recovering a sparse signal $\mathbf{x}_0 \in \mathbb{R}^n$ from highly corrupted linear measurements $\mathbf{y} \in \mathbb{R}^m$:

$$\mathbf{y} = A\mathbf{x}_0 + \mathbf{e}_0,$$

where $\mathbf{e}_0 \in \mathbb{R}^m$ is a vector of errors of arbitrarily large magnitude. We are interested in the situation when $[A, \mathbf{I}]$ forms a cross-and-bouquet (CAB) model, for which we give a precise definition below.

The model for $A \in \mathbb{R}^{m \times n}$ should capture the idea that it consists of small deviations about a mean, hence a “bouquet”:

Assumption 1 (Bouquet Model): We consider the case where the columns of A are i.i.d. samples from a Gaussian distribution: $A = [\mathbf{a}_1, \dots, \mathbf{a}_n] \in \mathbb{R}^{m \times n}$, with

$$\mathbf{a}_i \sim_{iid} \mathcal{N}\left(\boldsymbol{\mu}, \frac{\nu^2}{m} \mathbf{I}_m\right), \quad \|\boldsymbol{\mu}\|_2 = 1, \quad \|\boldsymbol{\mu}\|_\infty \leq C_\mu m^{-1/2}, \quad (3)$$

for some constant $C_\mu \geq 1$.

Together, the two conditions on the variance and mean force the bouquet to remain both tight and incoherent with the standard basis (or “cross”) as the dimension m increases.

We assume little prior knowledge in the signal \mathbf{x}_0 and error \mathbf{e}_0 to be recovered. In this paper, we assume:

Assumption 2 (Signs and support of \mathbf{x}_0 and \mathbf{e}_0): The signs σ_x and support I of \mathbf{x}_0 are arbitrary. The support J of \mathbf{e}_0 is chosen uniformly at random, and the signs σ of \mathbf{e}_0 are random (iid symmetric ± 1 on J).

Geometrically, for \mathbf{x}_0 and \mathbf{e}_0 to be recoverable from the ℓ^1 minimization (2), it would require vertices from the “bouquet” A to “see” through the crosspolytope $\pm I$ to vertices of the “negative bouquet” $-A$ that are nearly antipodal to them. In a low-dimensional space, it is difficult to imagine this is even possible at all. However, our analysis will show that in very high-dimensional spaces, as long as \mathbf{x}_0 is sparse enough, this seemingly-impossible event occurs with probability approaching one.

We study the behavior of the solution to the ℓ^1 -minimization (2) in the following growth framework:

Assumption 3 (Proportional Growth): We say a sequence of signal-error problems exhibits proportional growth with parameters $\delta > 0$, $\rho \in (0, 1)$, $\alpha > 0$ if

$$n = \lfloor \delta m \rfloor, \quad \|\mathbf{e}_0\|_0 = \lfloor \rho m \rfloor, \quad \|\mathbf{x}_0\|_0 = \lfloor \alpha m \rfloor. \quad (4)$$

In the above assumption, while the support size of the error \mathbf{e} can be any fraction $\rho < 1$ of the dimension m , the support size of the sparse signal \mathbf{x} is only a very small fraction of m : $\|\mathbf{x}\|_0 = \alpha m$ for some small constant $\alpha > 0$ (which may depend on δ and ρ). This condition differs from the typical assumption in the sparse representation literature in the sense that the signals of interest \mathbf{x} and \mathbf{e} have different support density. We will discuss in more detail below why the support density α of the signal \mathbf{x} must be very small when ρ is close to one (in fact, a small fraction of $1 - \rho$). However, such linear sparsity of \mathbf{x} is more than adequate for signals arising in many practical problems where the support size is sublinear in the dimension or even bounded. For example in face recognition, the support size of \mathbf{x} is bounded by a constant – the number of images per subject.

Under Assumptions 1, 2, and 3, we have:

Theorem 1 (Dense Error Correction in the CAB): Fix any $\delta > 0$, $\rho < 1$. Suppose that A is distributed according to (3) with ν sufficiently small, that $J \subset [m]$ is a uniform random subset of size ρm , and that $\sigma \in \{\pm 1\}$ is iid symmetric (independent of J). If m is sufficiently large, then with probability at least $1 - C \exp(-\varepsilon^* m)$ in A, J, σ , for all \mathbf{x}_0 with $\|\mathbf{x}_0\|_0 \leq \alpha^* m$ and any \mathbf{e}_0 with $\text{supp}(\mathbf{e}_0) = J$ and $\text{sign}(\mathbf{e}_0(J)) = \sigma$,

$$(\mathbf{x}_0, \mathbf{e}_0) = \arg \min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{s.t.} \quad A\mathbf{x} + \mathbf{e} = A\mathbf{x}_0 + \mathbf{e}_0,$$

and the minimizer is uniquely defined.

Here, C is numerical, α^* and ε^* are positive constants (wrt m), which depend on δ, ρ, ν . By “ ν sufficiently small” and “ m sufficiently large” we mean that $0 < \nu \leq \nu^*$ and $m \geq m^*$, where $\nu^* > 0$ and $m^* > 0$ are constant wrt m .

Below, we comment further on the nature of the constants in Theorem 1. We find it convenient to introduce one special function, which is defined as follows. Let $f : (0, 1/e) \rightarrow (0, 1/e)$ be defined by

$$f(x) = x \log(1/x).$$

It is easy to verify that f is strictly increasing on $(0, 1/e)$, and hence has a well-defined inverse. We let

$$\Upsilon : (0, 1/e) \rightarrow (0, 1/e), \quad \Upsilon[t] = f^{-1}(t)$$

denote this inverse.³ Because f is strictly increasing, Υ is also strictly increasing. Since $\forall x \in (0, 1/e)$, $f(x) \geq x$,

$$\Upsilon[t] \leq t \quad \forall t \in (0, 1/e). \quad (5)$$

Finally, notice that for any fixed a and $t < a/e$,

$$x \log(a/x) \leq t \iff x \leq a\Upsilon(t/a), \quad (6)$$

a property that will be useful in taking union bounds.

Remark 2: One can give explicit expressions for the constants in Theorem 1. We first introduce

$$\begin{aligned} \tau &= \exp(-48^{-2} \delta^{-1} \nu^{-2}), \\ C_s &= \Upsilon[2^{-17} C_\mu^{-4} \delta^{-1} (1 - \rho)^2]. \end{aligned} \quad (7)$$

These two quantities will play important roles in the proof, but for now we just record them to keep the notation compact. One can then choose $C = 31$, and set

$$\begin{cases} \nu_{\delta, \rho}^* &= 2^{-12} \sqrt{C_s / \delta}, \\ \varepsilon_{\delta, \rho, \nu}^* &= \nu^2 (1 - \rho) \tau^2 / 2, \\ \alpha_{\delta, \rho, \nu}^* &= \nu^2 \Upsilon[\varepsilon^* / \delta], \\ m_{\delta, \rho, \nu}^* &= 2^{17} / (1 - \rho) C_s. \end{cases} \quad (9)$$

Above, we have added subscripts to emphasize the dependence of these quantities on the properties of the bouquet (δ, ν) and the density of the error, ρ . In calculations below, we occasionally drop these subscripts for the sake of compactness; e.g., α^* will refer to $\alpha_{\delta, \rho, \nu}^*$.

One immediately notices that the given constants are extremely small (or, in the case of m^* , extremely large). We did not attempt to obtain the best possible constants. Indeed, it is reasonable to believe that the best constants that can be obtained via our analysis will be quite pessimistic. Our analysis uses linear algebraic manipulations to reduce the quantity of interest to a function of sparse singular values and then applies standard concentration results to argue that they are well-behaved for all sparse vectors. This is reminiscent of (and inspired by) the RIP analysis of the standard compressed sensing setup. That analysis obtains, in a very clean and simple manner, guarantees for sparse recovery that are order-optimal. However, in a finite- m setting RIP results are pessimistic compared to those obtained with tools from geometric functional analysis [18] or combinatorics [19].

It should also be noted that if the support of \mathbf{x}_0 is sublinear (i.e., $\|\mathbf{x}_0\|_0 = o(m)$), exact recovery holds for sufficiently

³One can express Υ in terms of the Lambert W function [17].

large m . In other words, as long as the bouquet is sufficiently tight, asymptotically ℓ^1 -minimization recovers any sufficiently sparse signal from almost any error with support fraction bounded away from 100%. Notice that we make no assumption on the signal \mathbf{x} , except that its support is extremely small. Existing results in the sparse representation literature typically either pertain to all sparse signals (i.e., the strong threshold of [20]) or almost all sparse signals, in terms of signs and support (i.e., the weak threshold of [20]). Our result has a hybrid flavor: it holds for all signals \mathbf{x} , but only for almost all errors \mathbf{e} , in terms of signs and support.

C. Relations to previous results

a) *Restricted isometry and incoherence of the cross-and-bouquet model:* As mentioned earlier, typical results in the literature for sparse signal recovery do not apply to equations of the type $\mathbf{y} = A\mathbf{x} + \mathbf{e}$. The cross-and-bouquet matrix $[A \ \mathbf{I}]$ is neither restrictedly isometric nor incoherent. As a result, greedy algorithms such as Orthogonal Matching Pursuit [21], [22] succeed only when the error \mathbf{e} is very sparse (see Section IV) for the simulation results and comparison with our method). However, this does not mean that the restricted isometry property is irrelevant to the new problem. On the contrary, similar ideas play an important role in our proof. Moreover, unlike the typical compressed sensing setting, the solution $[\mathbf{x}, \mathbf{e}]$ sought has very uneven density (or sparsity). This is reminiscent of the block sparsity studied in [23]. However, as we will see, the special block structure of the cross-and-bouquet model enables sparse recovery far beyond the breakdown point for general sparse (or block sparse) signals.

b) *Error correction:* From an error correction viewpoint, the above result seems surprising: One can correctly solve a set of linear equations with *almost all* the equations randomly corrupted. This is perhaps surprising considering that the best error-correcting codes (in the binary domain \mathbb{Z}_2), constructed based on expander graphs, normally correct a fixed fraction of errors [24]–[26] (although such results typically pertain to worst-case, rather than random errors). The exact counterpart of our result in the binary domain is not clear.⁴ While there are superficial similarities between our result and [25], [27] in the use of linear programming for decoding and analysis via polytope geometry, those works do not consider real valued signals. In particular, the negative result of [27] for specific families of binary codes admitting linear programming decoders does not apply here.

We can, however, draw the following comparisons with existing error correction methods in the domain of real numbers:

- When $n < m$, the range of A is a subspace in \mathbb{R}^m . In such an overdetermined case, one could directly apply the method of Candes and Tao [14] mentioned earlier. However, that result requires the error vector \mathbf{e} to be relatively sparse, and treats \mathbf{x} as a generic vector. In fact, as we will see in simulations in Section IV, even in the overdetermined case, solving (2) has advantages

for coherent matrices A , if the signal \mathbf{x} to be recovered is known to be sparse.

- The small linear growth of the support of \mathbf{x} in m is the best one can hope for in the regime of dense errors. In general, we need at least $\|\mathbf{x}\|_0$ uncorrupted linear measurements to recover \mathbf{x} uniquely. If a fraction $\rho < 1$ of the m equations is totally corrupted by \mathbf{e} , the support of \mathbf{x} has to be smaller than the number of remaining $(1 - \rho)m$ uncorrupted measurements. On the other hand, the sparser the error \mathbf{e} is, the larger support \mathbf{x} can be. Simulation results in Section IV also confirm this phenomenon. However, in this paper, we are mainly interested in how ℓ^1 -minimization behaves with dense errors, for ρ very close to 1.
- When $n \geq m$, the matrix A is full rank and the method of Candes and Tao [14] no longer applies.⁵ Our result suggests that as long as A is highly correlated, the extended ℓ^1 -minimization (2) can still recover the sparse signal \mathbf{x} exactly even if almost all the equations might be corrupted. This is verified by the simulation results in Section IV.

c) *Polytope geometry:* The success of ℓ^1 -minimization in recovering sparse solutions \mathbf{x} from underdetermined systems of linear equations $\mathbf{y} = A\mathbf{x}$ can be viewed as a consequence of a surprising property of high-dimensional polytopes. If the column vectors of A are random samples from a zero-mean Gaussian $\mathcal{N}(0, \mathbf{I})$, and m and n are allowed to grow proportionally, then with overwhelming probability the convex polytope spanned by the columns of A is highly centrally symmetric neighborly [20], [28]. Neighborliness provides the necessary and sufficient condition for uniform sparse recovery: the ℓ^1 -minimization (1) correctly recovers \mathbf{x} if and only if the columns associated with the nonzero entries of \mathbf{x} span a face of the polytope $\text{conv}(A)$.

In our case, the columns of the matrix A are i.i.d. Gaussian vectors with nonzero mean $\boldsymbol{\mu}$ and small variance σ^2 , whereas the vectors of the identity matrix \mathbf{I} are completely fixed. To characterize when the extended ℓ^1 -minimization (2) is able to recover the solution $[\mathbf{x}, \mathbf{e}]$ correctly, we need to examine the geometry of the peculiar centrally symmetric convex polytope $\text{conv}(\pm A, \pm \mathbf{I})$ spanned together by columns of A and \mathbf{I} . In a sense, our result implies that the “cross-and-bouquet” polytope is neighborly⁶. Or in other words, the vertices associated with the nonzero entries of \mathbf{x} and \mathbf{e} form a face of the polytope with probability approaching one as the dimension m becomes large. Precisely due to this kind of neighborliness of the cross-and-bouquet polytope, the extended ℓ^1 -minimization (2) is able to correctly recover the desired solution, even though the

⁵One could choose to premultiply the equation $\mathbf{y} = A\mathbf{x} + \mathbf{e}$ with an “approximate orthogonal complement” of A , say the orthogonal complement of the mean vector $\boldsymbol{\mu}$, which is an $(m - 1) \times m$ matrix B . This gives $B\mathbf{y} = B\mathbf{e} + \mathbf{z}$ where $\mathbf{z} = BA\mathbf{x}$. If the norm of \mathbf{x} is bounded, then \mathbf{z} is a signal with small magnitude due to the near-orthogonality of B and A . In this case, one can view \mathbf{z} as a noise term and try to recover \mathbf{e} as a sparse signal via ℓ^1 -minimization. However, the recovery will not be exact, and it is then unclear how to subsequently recover the desired sparse signal \mathbf{x} .

⁶Although, the precise notion of neighborliness is not the same - our result can be viewed as a hybrid of the “weak” neighborliness of [28] (since the signs and support of \mathbf{e} are random) and “strong” (or central) neighborliness [28] (since the signs and support of \mathbf{x} are arbitrary).

⁴It is possible that under an analogous growth model (see Section I-B), the LP decoder of [25] could also correct large fractions of binary errors.

part of the solution corresponding to e might be dense.

D. Implications for applications

a) Robust reconstruction, classification, and source separation: The new result about the cross-and-bouquet model has strong implications for robust reconstruction, classification, and separation of highly correlated classes of signals such as faces or voices, despite severe corruption. It helps explain the surprising performance of robust face recognition that we discussed earlier. It further suggests that if the resolution of the image increases in proportion with the size of the database (say, due to the increasing number of subjects), the ℓ^1 -minimization would tolerate an even higher level of corruption, far beyond the 60% at the 80×60 resolution experimented with in [11]. Other applications where this kind of model could be useful and effective include speech recognition/imputation, audio source separation, video segmentation, or activity recognition from motion sensors.

b) Communication through an almost random channel: The result suggests that we can use the cross-and-bouquet model to accurately send information through a highly corrupting channel. Hypothetically, we can imagine a channel through which we can send one real number at a time, say as one packet of binary bits, and each packet has a high probability $\rho < 1$ of being totally corrupted. One can use the sparse vector \mathbf{x} (or its support) to represent useful information, and use a set of highly correlated high-dimensional vectors as the encoding transformation A . The high correlation in A ensures that there is sufficient redundancy built in the encoded message $A\mathbf{x}$ so that the information about \mathbf{x} will not be lost even if most entries of $A\mathbf{x}$ can be corrupted while being sent through such a channel (although the coding “rate” for \mathbf{x} is pessimistically small with the current estimates, roughly characterized by α which is a small fraction of $1 - \rho$). Our result suggests that the decoding can be done correctly and efficiently using linear programming.

c) Encryption and information hiding: One can potentially use the cross-and-bouquet model for encryption. For instance, if both the sender and receiver share the same encoding matrix A (say a randomly chosen Gaussian matrix), the sender can deliberately corrupt the message $A\mathbf{x}$ with arbitrary random errors e before sending it to the receiver. The receiver can use linear programming to decode the information \mathbf{x} , whereas any eavesdropper will not be able to make much sense out of the highly corrupted message $\mathbf{y} = A\mathbf{x} + e$. The higher ρ is, the more “secured” the encrypted message would be (in exchange for a lower coding rate for \mathbf{x}). Of course, the long-term security of such an encryption scheme relies on the difficulty of learning the encoding matrix A after gathering many instances of corrupted message. It is not even clear whether it is easy to learn A from instances of uncorrupted message $\mathbf{y} = A\mathbf{x}$ – existing algorithms do not guarantee global optimality [9]. Presumably, in the presence of corruption the problem becomes even more daunting.

II. ROADMAP OF PROOF

In this section, we will lay out the roadmap of the proof for our main result. After introducing some basic notation,

we will prove two key lemmas. In subsection II-B, we first introduce a certificate \mathbf{q} whose existence guarantees the uniqueness and optimality of (\mathbf{x}_0, e_0) . In subsection II-C, the second lemma gives a set of sufficient conditions under which the certificate of optimality can be constructed via a simple iterative procedure. We postpone to Section III the detailed probabilistic analysis that proves that these conditions hold with overwhelming probability under the assumptions of the theorem. The proof will borrow many standard measure concentration results that are summarized in the appendix for completeness. We put such technical details in a separate section and the appendix so that the reader might want to skip them in a first reading, and directly jump to Section IV for empirical verification and experiments.

A. Notation

Before stating our main result, we fix some additional notation. For any $n \in \mathbb{Z}_+$, $[n]$ denotes the set $\{1, \dots, n\}$. For $n, k \in \mathbb{Z}_+$, $\binom{[n]}{k}$ denotes the set of subsets $S \subset [n]$ of size k . We will let $I = \text{supp}(\mathbf{x}_0) \subset [n]$, $J = \text{supp}(e_0) \subset [m]$. Let $k_1 = |I|$ be the support size of the signal \mathbf{x}_0 and $k_2 = |J|$ the support size of the error e_0 , and let $\boldsymbol{\sigma} = \text{sign}(e_0(J)) \in \{\pm 1\}^{k_2}$. For $M \in \mathbb{R}^{m \times n}$, M_i will denote the i -th row of M ; for $I \subseteq [m]$, M_I will denote the submatrix of M containing the rows indexed by I . If further $J \subseteq [n]$, $M_{I,J} \in \mathbb{R}^{|I| \times |J|}$ will denote the square submatrix of M indexed by I and J . To reduce confusion between the index set I and the identity matrix, we use \mathbf{I} to denote the latter.

We will use several norms on vectors and matrices. For $\mathbf{x} \in \mathbb{R}^m$, $\|\mathbf{x}\|_1 = \sum_i |x_i|$ is the ℓ^1 -norm, $\|\mathbf{x}\|$ is the ℓ^2 norm. For matrices M , $\|M\| = \sigma_{\max}(M)$ denotes the ℓ^2 operator norm, $\|M\|_F = \sqrt{\text{tr}[M^*M]}$ denotes the Frobenius norm, $\|M\|_{2,\infty} = \max_i \|M_i\|$ denotes the $\ell^2 \rightarrow \ell^\infty$ operator norm. For an $m \times n$ matrix M with ordered singular values $\sigma_1 \geq \sigma_2 \geq \dots$, we will use $\sigma_{\min}(M)$ to denote $\sigma_{\min\{m,n\}}(M)$, i.e., the smallest singular value. For a linear subspace $V \subseteq \mathbb{R}^n$, $P_V \in \mathbb{R}^{n \times n}$ will denote the projection operator onto V .

We will make extensive use of the shrinkage operator S with threshold $1/2$, defined for scalars t as

$$S[t] = \text{sign}(t) \max\{t - 1/2, 0\}, \quad (10)$$

and extended to vectors by applying it elementwise.

B. Optimality conditions

Notice that the matrix A is distributed as

$$\begin{aligned} A &= [\mathbf{a}_1 \mid \dots \mid \mathbf{a}_n] \\ &= [\boldsymbol{\mu} + \nu \mathbf{z}_1 \mid \dots \mid \boldsymbol{\mu} + \nu \mathbf{z}_n] \doteq \boldsymbol{\mu} \mathbf{1}^* + \nu Z, \end{aligned}$$

where $Z \in \mathbb{R}^{m \times n}$ has elements i.i.d. $\mathcal{N}(0, 1/m)$ and represents the deviations of the elements \mathbf{a}_i of the bouquet about the common mean $\boldsymbol{\mu}$. The following lemma expresses the optimality conditions for (\mathbf{x}_0, e_0) in terms of properties of $\boldsymbol{\mu}$ and submatrices of Z .

Lemma 1: Fix $\mathbf{x}_0 \in \mathbb{R}^n$ and $e_0 \in \mathbb{R}^m$. Let $I = \text{supp}(\mathbf{x}_0)$, $\boldsymbol{\sigma}_x = \text{sign}(\mathbf{x}_0(I))$, $J = \text{supp}(e_0)$, $\boldsymbol{\sigma} = \text{sign}(e_0(J))$. Suppose

that the following system of equations has a solution $\mathbf{q} \in \mathbb{R}^{m-k_2+n-k_1}$ with $\|\mathbf{q}\|_\infty < 1$:

$$\begin{cases} \begin{bmatrix} \nu(Z_{J^c, I^c})^* & \mathbf{I} \\ \nu(Z_{J^c, I})^* & 0 \end{bmatrix} \mathbf{q} = \begin{bmatrix} -\nu(Z_{J, I^c})^* \boldsymbol{\sigma} \\ \boldsymbol{\sigma}_x - \nu(Z_{J, I})^* \boldsymbol{\sigma} \end{bmatrix}, \\ \langle \boldsymbol{\mu}_{J^c}, \mathbf{q}_{1\dots m-k_2} \rangle = -\langle \boldsymbol{\mu}_J, \boldsymbol{\sigma} \rangle, \end{cases} \quad (11)$$

and that

$$A_{J^c, I} \text{ has full rank } k_1. \quad (12)$$

Then $(\mathbf{x}_0, \mathbf{e}_0)$ is the unique optimal solution to

$$\min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad A\mathbf{x} + \mathbf{e} = A\mathbf{x}_0 + \mathbf{e}_0. \quad (13)$$

Proof: Duality (e.g., Theorem 4 of [29]) shows that if there exists $\mathbf{h} \in \mathbb{R}^m$ satisfying

$$\begin{cases} (A^*)_I \mathbf{h} = \boldsymbol{\sigma}_x, \quad \mathbf{h}_J = \boldsymbol{\sigma}, \\ \|(A^*)_{I^c} \mathbf{h}\|_\infty < 1, \quad \|\mathbf{h}_{J^c}\|_\infty < 1, \end{cases} \quad (14)$$

and

$$B \doteq [A_{[m], I} \quad (\mathbf{I})_{[m], J}] \text{ has full rank } k_1 + k_2. \quad (15)$$

Then $(\mathbf{x}_0, \mathbf{e}_0)$ is the unique optimal solution to (13).

Notice that $\text{rank}(B) = k_2 + \text{rank}(A_{J^c, I})$, and so (12) implies (15). Consider the system

$$\begin{bmatrix} (A_{J^c, I^c})^* & \mathbf{I} \\ (A_{J^c, I})^* & 0 \end{bmatrix} \mathbf{q} = \begin{bmatrix} -(A_{J, I^c})^* \boldsymbol{\sigma} \\ \boldsymbol{\sigma}_x - (A_{J, I})^* \boldsymbol{\sigma} \end{bmatrix}, \quad \|\mathbf{q}\|_\infty < 1, \quad (16)$$

in unknown $\mathbf{q} \in \mathbb{R}^{m+n-k_1-k_2}$. If there exists \mathbf{q} satisfying (16), then setting $\mathbf{h}_J = \boldsymbol{\sigma}$ and $\mathbf{h}_{J^c} = \mathbf{q}_{1\dots m-k_2}$ gives a dual vector \mathbf{h} satisfying (14). Since any \mathbf{q} satisfying (11) also satisfies (16), this completes the proof. \blacksquare

Thus, we merely need to construct a vector \mathbf{q} of small ℓ^∞ norm satisfying (11). Notice that if ν is small, the $n+1$ equality constraints in (11) may have very different scale. We introduce an equivalent set of constraints that is better conditioned, and also introduce notation for the various components.

In the top component of (11), involving $[\nu(Z_{J^c, I^c})^* \mathbf{I}]\mathbf{q}$, the rows have norm ≈ 1 . For this component, let $U \doteq (Z_{J^c, I^c})^* \in \mathbb{R}^{(n-k_1) \times (m-k_2)}$ and

$$R \doteq \begin{bmatrix} -(Z_{J, I^c})^* \\ 0 \\ 0 \end{bmatrix} \in \mathbb{R}^{n+1 \times k_2}.$$

In the middle component, involving $[\nu(Z_{J^c, I})^* \mathbf{0}]\mathbf{q}$ the rows have norm $\approx \nu\sqrt{1-\rho}$. We rescale each row by $\nu^{-1}(1-\rho)^{-1/2}$. Finally, the last constraint $[\boldsymbol{\mu}_{J^c}^* \mathbf{0}]\mathbf{q} = -\langle \boldsymbol{\mu}_J, \boldsymbol{\sigma} \rangle$, has norm $\approx \sqrt{1-\rho}$. We rescale this constraint by $(1-\rho)^{-1/2}$. Write

$$\begin{aligned} W &\doteq (1-\rho)^{-1/2} \begin{bmatrix} (Z_{J^c, I^c})^* \\ \boldsymbol{\mu}_{J^c}^* \end{bmatrix} \in \mathbb{R}^{k_1+1 \times m-k_2}, \\ T &\doteq (1-\rho)^{-1/2} \begin{bmatrix} 0 \\ -(Z_{J, I})^* \\ -\boldsymbol{\mu}_J^* \end{bmatrix} \in \mathbb{R}^{n+1 \times k_2}, \\ \mathbf{z} &\doteq (1-\rho)^{-1/2} \nu^{-1} \begin{bmatrix} 0 \\ \boldsymbol{\sigma}_x \\ 0 \end{bmatrix} \in \mathbb{R}^{n+1}. \end{aligned}$$

Further set

$$\Phi \doteq \begin{bmatrix} \nu U & \mathbf{I} \\ W & 0 \end{bmatrix} \in \mathbb{R}^{n+1 \times m+n-k_1-k_2}. \quad (17)$$

Then \mathbf{q} is a solution to (11) if and only if \mathbf{q} satisfies

$$\Phi \mathbf{q} = \nu R \boldsymbol{\sigma} + T \boldsymbol{\sigma} + \mathbf{z}. \quad (18)$$

In general, (18) has many solutions. To certify optimality, we need only produce a solution \mathbf{q}^* to (18) with ℓ^∞ norm smaller than one. As we will see later, Φ is overwhelmingly likely to be well-conditioned, so that we can construct a close guess at such a certificate \mathbf{q} using least squares. The next section will show that such a guess is actually sufficient.

C. Iterative construction of separator

The following lemma allows us to relax our requirements on \mathbf{q} : rather than directly producing a solution with $\|\mathbf{q}\|_\infty < 1$, we can instead produce a solution of small ℓ^2 norm that does not too severely violate the ℓ^∞ constraint. More precisely, it suffices to construct a solution with $\|\mathcal{S}[\mathbf{q}]\|_2$ sufficiently small, where \mathcal{S} is the soft-thresholding operator with threshold $1/2$, defined in (10). For simplicity we often write $\mathcal{S}[\mathbf{q}]$ as $\mathcal{S}\mathbf{q}$.

Lemma 2: Consider an underdetermined system of equations $\Phi \mathbf{q} = \mathbf{w}$. Let P_Φ denote the projection operator onto $\text{range}(\Phi^*)$, and let $\xi_s(\Phi)$ denote the norm P_Φ , restricted to sparse vectors:

$$\xi_s(\Phi) \doteq \sup_{\substack{\|\mathbf{x}\|_0 \leq s, \\ \|\mathbf{x}\|_2 \leq 1}} \|P_\Phi \mathbf{x}\|_2. \quad (19)$$

Suppose $\xi_s(\Phi) < 1$ and there exists a solution \mathbf{q}_0 with $\Phi \mathbf{q}_0 = \mathbf{w}$ satisfying

$$\|\mathbf{q}_0\| + \frac{\xi_s(\Phi)}{1 - \xi_s(\Phi)} \|\mathcal{S}\mathbf{q}_0\| \leq \sqrt{s}/2. \quad (20)$$

Then there exists a solution \mathbf{q}^* satisfying $\|\mathbf{q}^*\|_\infty \leq 1/2$.

Proof: We construct a convergent sequence of vectors $\mathbf{q}_0, \mathbf{q}_1, \dots$ whose limit \mathbf{q}^* satisfies $\Phi \mathbf{q}^* = \mathbf{w}$, $\|\mathbf{q}^*\|_\infty \leq 1/2$. Set

$$\mathbf{q}_i = \mathbf{q}_{i-1} - (\mathbf{I} - P_\Phi) \mathcal{S}[\mathbf{q}_{i-1}], \quad (21)$$

and notice that for all i , $\Phi \mathbf{q}_i = \Phi \mathbf{q}_0 = \mathbf{w}$.

We will show by induction on i that the sequence (\mathbf{q}_i) obeys the following estimates

$$\begin{aligned} \text{(i)} \quad \|\mathbf{q}_i\| &\leq \|\mathbf{q}_0\| + \|\mathcal{S}\mathbf{q}_0\| \sum_{j=1}^i \xi_s(\Phi)^j, \\ \text{(ii)} \quad \|\mathcal{S}\mathbf{q}_i\| &\leq \xi_s(\Phi)^i \|\mathcal{S}\mathbf{q}_0\|, \\ \text{(iii)} \quad \|\mathcal{S}\mathbf{q}_i\|_0 &\leq s. \end{aligned}$$

For the base case $i = 0$, (i),(ii) hold trivially. For (iii), notice that since each nonzero element of $\mathcal{S}[\mathbf{q}_0]$ corresponds to an element of \mathbf{q}_0 of size at least $1/2$, $\|\mathcal{S}[\mathbf{q}_0]\|_0 \times 1/4 \leq \|\mathbf{q}_0\|^2$. By (20), $\|\mathbf{q}_0\|^2 \leq s/4$, establishing (iii).

Now, suppose (i)-(iii) hold for $0, \dots, i-1$. For any scalar t , $|t - \mathcal{S}(t)| \leq |t|$, $\|\mathbf{q}_{i-1} - \mathcal{S}[\mathbf{q}_{i-1}]\| \leq \|\mathbf{q}_{i-1}\|$. Hence, using that $\mathcal{S}[\mathbf{q}_{i-1}]$ is s -sparse,

$$\begin{aligned} \|\mathbf{q}_i\| &= \|\mathbf{q}_{i-1} - \mathcal{S}[\mathbf{q}_{i-1}] + P_\Phi \mathcal{S}[\mathbf{q}_{i-1}]\|, \\ &\leq \|\mathbf{q}_{i-1} - \mathcal{S}[\mathbf{q}_{i-1}]\| + \|P_\Phi \mathcal{S}[\mathbf{q}_{i-1}]\|, \\ &\leq \|\mathbf{q}_{i-1}\| + \xi_s(\Phi) \|\mathcal{S}[\mathbf{q}_{i-1}]\|, \\ &\leq \|\mathbf{q}_0\| + \|\mathcal{S}[\mathbf{q}_0]\| \sum_{j=1}^i \xi_s(\Phi)^j, \end{aligned}$$

and so (i) holds for i . Bounding the summation by $\sum_{j=1}^{\infty} \xi_s(\Phi)^j = \frac{\xi_s}{1-\xi_s}$ and applying (20), $\|\mathbf{q}_i\|^2 \leq s/4$, and hence $\|\mathcal{S}[\mathbf{q}_i]\|_0 \leq s$, establishing (iii) for i .

Finally, notice that for any vectors \mathbf{p}, \mathbf{q} ,

$$\|\mathcal{S}[\mathbf{q} + \mathbf{p}]\| \leq \|\mathcal{S}[\mathbf{q}]\| + \|\mathbf{p}\|,$$

and further, that $\mathcal{S}[\mathbf{q} - \mathcal{S}[\mathbf{q}]] = 0$. Hence,

$$\begin{aligned} \|\mathcal{S}[\mathbf{q}_i]\| &= \|\mathcal{S}[\mathbf{q}_{i-1} - \mathcal{S}[\mathbf{q}_{i-1}] + P_{\Phi}\mathcal{S}[\mathbf{q}_{i-1}]\| \\ &\leq \|\mathcal{S}[\mathbf{q}_{i-1} - \mathcal{S}[\mathbf{q}_{i-1}]\| + \|P_{\Phi}\mathcal{S}[\mathbf{q}_{i-1}]\| \\ &\leq \xi_s(\Phi)\|\mathcal{S}[\mathbf{q}_{i-1}]\|. \end{aligned}$$

Applying the inductive hypothesis establishes (ii), and so the three conditions hold for all i .

Since $\|\mathbf{q}_i - \mathbf{q}_{i-1}\| \leq \|\mathcal{S}[\mathbf{q}_{i-1}]\| \leq \xi_s^i \|\mathcal{S}[\mathbf{q}_0]\|$, the sequence (\mathbf{q}_i) converges to some limit \mathbf{q}^* . Moreover, $\|\mathcal{S}[\mathbf{q}^*]\| = \lim_{i \rightarrow \infty} \|\mathcal{S}[\mathbf{q}_i]\| = 0$, and hence $\mathcal{S}[\mathbf{q}^*] = 0$ and so all of the elements of \mathbf{q}^* have magnitude $\leq 1/2$. ■

D. Ideas for the remainder of the proof.

Lemmas 1 and 2 have shown that to verify that a given $\mathbf{x}_0, \mathbf{e}_0$ is the optimal solution to the ℓ^1 -minimization (2), it is sufficient to produce a solution \mathbf{q}_0 to (18) satisfying (20) (and further to verify that (12) holds). We consider a very natural candidate \mathbf{q}_0 , the minimum ℓ^2 -norm solution to (18), and show that with overwhelming probability it satisfies (20). To do this, we must bound three quantities: the restricted operator norm ξ_s , the norm of the initial guess $\|\mathbf{q}_0\|_2$, and the norm of the initial violations $\|\mathcal{S}\mathbf{q}_0\|_2$. The next section gives probabilistic bounds on these three quantities that show that which overwhelming probability they hold simultaneously for m sufficiently large. As we will see in Section III-D, the failure probability is small enough to allow a union bound over all sufficiently sparse signals \mathbf{x}_0 . Readers who are less interested in the details of the analysis can skip the next section in the first reading, and directly jump to the empirical verification in Section IV.

III. PROBABILISTIC ANALYSIS AND FINAL PROOF

In this section, we complete the proof of the main result. In Section III-A, we first introduce several standard concentration results that we will use throughout. In Section III-B, we then show that with overwhelming probability, the range of the transpose of the matrix Φ in (18) is not too aligned with any sparse vectors, i.e., the quantity $\xi_s(\Phi)$ in Lemma 2) is bounded away from one. In Section III-C, we then show that with overwhelming probability, the guess \mathbf{q}_0 constructed by least squares has small norm, and further that $\|\mathcal{S}\mathbf{q}_0\|$ is small. Finally, Section III-D combines these observations to prove Theorem 1.

A. Tools and preliminaries

Our probabilistic analysis depends on measure concentration results for Lipschitz functions of random matrices, in particular their norms and singular values. In the following,

let M be an $am \times bm$, $a > b$, random matrix with entries iid $\mathcal{N}(0, 1/m)$. Then (see [30] and the references therein)

$$\mathbb{P}[\|M\| \geq \sqrt{a} + \sqrt{b} + \sqrt{2\varepsilon}] < \exp(-\varepsilon m), \quad (22)$$

$$\mathbb{P}[\sigma_{\min}(A) \leq \sqrt{a} - \sqrt{b} - \sqrt{2\varepsilon}] < \exp(-\varepsilon m). \quad (23)$$

Similarly (e.g., via [3] Equation (2.35))

$$\mathbb{P}[\|M\|_F \geq \sqrt{abm} + \sqrt{2\varepsilon}] \leq \exp(-\varepsilon m). \quad (24)$$

We will be interested in the action of random matrices on sparse vectors, which can be studied via the following extremal submatrix singular values⁷:

$$\begin{aligned} \beta_s(M) &\doteq \sup_{\|\mathbf{x}\|_2 \leq 1, \|\mathbf{x}\|_0 \leq s} \|M\mathbf{x}\|_2, \\ \gamma_s(M) &\doteq \inf_{\|\mathbf{x}\|_2 = 1, \|\mathbf{x}\|_0 \leq s} \|M\mathbf{x}\|_2, \\ \phi_s(M) &\doteq \sup_{\substack{\|\mathbf{x}\|_2 \leq 1, \|\mathbf{x}\|_0 \leq s \\ \|\mathbf{y}\|_2 \leq 1, \|\mathbf{y}\|_0 \leq s}} \mathbf{x}^* M \mathbf{y}. \end{aligned}$$

Now let $a > 0$ be arbitrary and suppose that $c < \min(a, b)$. Since $\binom{m}{k} < (me/k)^k$, the $am \times bm$ matrix M has $\binom{bm}{cm} \leq \exp(mc \log \frac{b}{c})$ submatrices $M_{[am], I}$ of size $am \times cm$. Union bounds then show that

$$\mathbb{P} \left[\beta_{cm}(M) \geq \sqrt{a} + \sqrt{c} + \sqrt{2c \log \left(\frac{b}{c} \right) + 2\varepsilon} \right] < \exp(-\varepsilon m). \quad (25)$$

$$\mathbb{P} \left[\gamma_{cm}(M) \leq \sqrt{a} - \sqrt{c} - \sqrt{2c \log \left(\frac{b}{c} \right) + 2\varepsilon} \right] < \exp(-\varepsilon m). \quad (26)$$

$$\begin{aligned} \mathbb{P} \left[\phi_{cm}(M) \geq 2\sqrt{c} + \sqrt{2c \log \left(\frac{a}{c} \right) + 2c \log \left(\frac{b}{c} \right) + 2\varepsilon} \right] \\ < \exp(-\varepsilon m). \end{aligned} \quad (27)$$

We also note that if S is (possibly random) subspace of dimension $a'm < am$ that is probabilistically independent of M , then by the rotational invariance of the Gaussian distribution, the above estimates hold for $P_S M$, with a replaced by a' . We will apply the above lemmas to bound the norms of the components of (18).

B. Projection of sparse vectors

We next show that sparse vectors cannot be *too* aligned with the range of Φ^* . In this lemma (and also in Section III-C), it is convenient to make assumptions that might appear to make the error correction problem more difficult. First, we will assume that $\delta \geq 1$; i.e., the matrix A is a square or wide matrix. Second, we assume that the fraction of errors is *large*. These two assumptions are technical in nature, and allow a somewhat more uniform argument. In the proof of Theorem 1 (Section III-D), we will first prove the result for this difficult case, and then show that it implies correct recovery for arbitrary $\delta > 0, \rho < 1$.

⁷The quantities β, γ, ϕ have a strong relationship to the restricted isometry property of the matrices in question. In particular, our argument suggests the possibility of a deterministic analogue of Theorem 1 under RIP-like conditions on the matrix Z

Lemma 3: Fix any $\delta \geq 1$, $\rho \geq 15/16$, and suppose that $\nu \leq \nu_{\rho, \delta}^*$, and that $m \geq m_{\rho, \delta, \nu}^*$. Fix any support $I \subset [n]$ of size no larger than

$$|I| \leq \alpha_{\delta, \rho, \nu}^* m, \quad (28)$$

and set

$$s = C_s(\delta, \rho) m. \quad (29)$$

Then with probability at least $1 - 9 \exp(-2\varepsilon_{\rho, \delta, \nu}^* m)$ in A, J , the matrix Φ satisfies

$$1 - \xi_s(\Phi) \geq \nu^4 \delta (1 - \rho) / 512. \quad (30)$$

Here, $C_s > 0$ is given in (8), and does not depend on m or ν .

To avoid the possibility for confusion, please notice that the quantity $s = C_s m$ in Lemma 3 represents our bound on the number of elements of $\mathcal{S}[\mathbf{q}]$ (i.e., the number of violations of the box constraint), rather than the sparsity of the signal \mathbf{x} itself.

Proof: (of Lemma 3) The proof has two parts: we first reduce the quantity of interest to a function of the singular values of several submatrices of U and W , and then apply Gaussian measure concentration to bound these quantities.

a) *Lower bounding* $1 - \xi$: Notice that

$$\begin{aligned} 1 - \xi_s(\Phi) &\geq \frac{1}{2}(1 - \xi_s(\Phi)^2) = \frac{1}{2} \left(1 - \sup_{\substack{\|\mathbf{w}\| \leq 1, \\ \|\mathbf{w}\|_0 \leq s}} \|P_\Phi \mathbf{w}\|^2 \right) \\ &= \frac{1}{2} \left(1 - \sup_{\substack{\|\mathbf{w}\| \leq 1, \\ \|\mathbf{w}\|_0 \leq s}} \|\mathbf{w}\|^2 - \|\mathbf{w} - P_\Phi \mathbf{w}\|^2 \right) \\ &= \frac{1}{2} \left(1 - \sup_{\substack{\|\mathbf{w}\| = 1, \\ \|\mathbf{w}\|_0 \leq s}} 1 - \|\mathbf{w} - P_\Phi \mathbf{w}\|^2 \right) \\ &= \frac{1}{2} \inf_{\substack{\|\mathbf{w}\| = 1, \\ \|\mathbf{w}\|_0 \leq s}} \|\mathbf{w} - P_\Phi \mathbf{w}\|^2 = \frac{1}{2} \inf_{\substack{\|\mathbf{w}\| = 1, \\ \|\mathbf{w}\|_0 \leq s}} \min_{\mathbf{r}} \|\mathbf{w} - \Phi^* \mathbf{r}\|^2. \end{aligned}$$

We can write this more explicitly as

$$\begin{aligned} &\frac{1}{2} \inf_{\substack{\|\mathbf{w}_1\|^2 + \|\mathbf{w}_2\|^2 = 1 \\ \|\mathbf{w}_1\|_0 + \|\mathbf{w}_2\|_0 \leq s}} \min_{\mathbf{r}} \left\| \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix} - \begin{bmatrix} \nu U^* & W^* \\ \mathbf{I} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix} \right\|^2 \\ &= \frac{1}{2} \inf_{\mathbf{w}_1, \mathbf{w}_2} \min_{\mathbf{u}_1, \mathbf{u}_2} \left\| \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix} - \begin{bmatrix} \nu U^* & W^* \\ \mathbf{I} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 + \mathbf{w}_2 \\ \mathbf{u}_2 \end{bmatrix} \right\|^2, \\ &= \frac{1}{2} \inf_{\mathbf{w}} \min_{\mathbf{u}} \|\mathbf{w}_1 - \nu U^* \mathbf{u}_1 - \nu U^* \mathbf{w}_2 - W^* \mathbf{u}_2\|^2 + \|\mathbf{u}_1\|^2. \end{aligned} \quad (31)$$

Notice that $W \in \mathbb{R}^{(k_1+1) \times (1-\rho)m}$ has a nontrivial nullspace; let $\Gamma \subset \mathbb{R}^{(1-\rho)m} = \text{null}(W)$, let P_Γ denote the projection onto Γ , and P_W denote the projection onto $\text{range}(W^*)$, so that

$$P_W + P_\Gamma = \mathbf{I}, \quad P_W P_\Gamma = 0, \quad \text{and} \quad P_\Gamma W^* = 0.$$

Then in terms of these operators, the first term in (31) can be written as

$$\begin{aligned} &\|\mathbf{w}_1 - \nu U^* \mathbf{u}_1 - \nu U^* \mathbf{w}_2 - W^* \mathbf{u}_2\|^2 \\ &= \|P_\Gamma(\mathbf{w}_1 - \nu U^* \mathbf{u}_1 - \nu U^* \mathbf{w}_2 - W^* \mathbf{u}_2)\|^2 \\ &\quad + \|P_W(\mathbf{w}_1 - \nu U^* \mathbf{u}_1 - \nu U^* \mathbf{w}_2 - W^* \mathbf{u}_2)\|^2, \\ &\geq \|P_\Gamma(\mathbf{w}_1 - \nu U^* \mathbf{u}_1 - \nu U^* \mathbf{w}_2 - W^* \mathbf{u}_2)\|^2, \\ &= \|\nu P_\Gamma U^* \mathbf{u}_1 - (P_\Gamma \mathbf{w}_1 - \nu P_\Gamma U^* \mathbf{w}_2)\|^2. \end{aligned}$$

Plugging into (31), $1 - \xi(\Phi)$ is bounded below by

$$\frac{1}{2} \inf_{\mathbf{w}_1, \mathbf{w}_2} \min_{\mathbf{u}_1} \|\nu P_\Gamma U^* \mathbf{u}_1 - (P_\Gamma \mathbf{w}_1 - \nu P_\Gamma U^* \mathbf{w}_2)\|^2 + \|\mathbf{u}_1\|^2. \quad (32)$$

For any matrix M and vector and \mathbf{y} , the unique solution to

$$\min_{\mathbf{x}} \|M\mathbf{x} - \mathbf{y}\|^2 + \|\mathbf{x}\|^2$$

is given by $\mathbf{x}^*(\mathbf{y}) = (\mathbf{I} + M^* M)^{-1} M^* \mathbf{y}$. Hence, in (32), for any $\mathbf{w} = [\mathbf{w}_1^*, \mathbf{w}_2^*]^*$, the optimizer $\mathbf{u}_1^*(\mathbf{w})$ has the explicit form

$$\begin{aligned} \mathbf{u}_1^*(\mathbf{w}) &= \nu (\mathbf{I} + \nu^2 U P_\Gamma U^*)^{-1} U P_\Gamma (\mathbf{w}_1 - \nu U^* \mathbf{w}_2) \\ &= \nu (\mathbf{I} + \nu^2 U P_\Gamma U^*)^{-1} U \Psi \mathbf{w}, \end{aligned} \quad (33)$$

where Ψ denotes the matrix $[P_\Gamma \mid -\nu P_\Gamma U^*]$. Notice that for any \mathbf{w} ,

$$\|\mathbf{u}^*(\mathbf{w})\| \geq \frac{\nu \sigma_{\min}(U)}{\|\mathbf{I} + \nu^2 U P_\Gamma U^*\|} \|\Psi \mathbf{w}\| \geq \frac{\nu \sigma_{\min}(U)}{1 + \nu^2 \|U\|^2} \|\Psi \mathbf{w}\|.$$

Hence, for any \mathbf{w} satisfying $\|\mathbf{w}\| = 1$ and $\|\mathbf{w}\|_0 \leq s$,

$$\|\mathbf{u}^*(\mathbf{w})\| \geq \frac{\nu \sigma_{\min}(U)}{1 + \nu^2 \|U\|^2} \gamma_s(\Psi). \quad (34)$$

Applying (32), we have that

$$1 - \xi_s(\Phi) \geq \frac{1}{2} \inf_{\substack{\|\mathbf{w}\| = 1 \\ \|\mathbf{w}\|_0 \leq s}} \|\mathbf{u}^*(\mathbf{w})\|^2 \geq \frac{\nu^2 \sigma_{\min}^2(U)}{2(1 + \nu^2 \|U\|^2)^2} \gamma_s^2(\Psi). \quad (35)$$

We apply a sequence of straightforward algebraic manipulations to lower bound $\gamma_s(\Psi)$ in terms of the singular values of submatrices of U and W . Each step below should be intuitive; their correctness is proved in Lemmas 6-9 of the appendix:

$$\begin{aligned} \gamma_s(\Psi) &\geq \nu \gamma_s([P_\Gamma \mid P_\Gamma U^*]), && \text{by (59),} \\ &= \nu \gamma_s([\mathbf{I} - P_W \mid (\mathbf{I} - P_W) U^*]), \\ &\geq \nu (\gamma_s([\mathbf{I} \ U^*]) - \beta_s(P_W) - \beta_s(P_W U^*)) && \text{by (60),} \end{aligned}$$

Above, in invoking (59), we have used that

$$\nu \leq \nu^* = 2^{-12} \sqrt{C_s / \delta} \ll 1.$$

Finally, by (63),

$$\gamma_s([\mathbf{I} \ U^*]) \geq \sqrt{\min\{1, \gamma_s^2(U^*)\} - \phi_s(U^*)}.$$

b) *Concentration of measure:* We will prove that when C_s is chosen according to (8), the following bounds hold:

- (i) $\|U\| \leq 3\sqrt{\delta}$ w.p. $\geq 1 - \exp(-2\varepsilon^*m)$,
- (ii) $\sigma_{\min}(U) \geq \sqrt{\delta}/2$ w.p. $\geq 1 - \exp(-2\varepsilon^*m)$,
- (iii) $\gamma_s(U^*) \geq \sqrt{(1-\rho)/2}$ w.p. $\geq 1 - \exp(-2\varepsilon^*m)$,
- (iv) $\phi_s(U^*) \leq \frac{1-\rho}{4}$ w.p. $\geq 1 - \exp(-2\varepsilon^*m)$,
- (v) $\beta_s(P_W U^*) \leq \frac{\sqrt{1-\rho}}{8}$ w.p. $\geq 1 - \exp(-2\varepsilon^*m)$,
- (vi) $\bar{\beta}_s(P_W) \leq \frac{\sqrt{1-\rho}}{8}$ w.p. $\geq 1 - 4\exp(-2\varepsilon^*m)$.

When each of the above bounds holds, then

$$\gamma_s(\Psi) \geq \frac{\nu\sqrt{1-\rho}}{4},$$

and by (35)

$$1 - \xi_s(\Phi) \geq \frac{\nu^2\delta}{8(1+9\nu^2\delta)^2} \gamma_s^2(\Psi) \geq \frac{\nu^2\delta}{32} \gamma_s^2(\Psi),$$

where we have used that

$$\nu \leq \nu^* = 2^{-12}\sqrt{C_s/\delta} < 2^{-12}/\sqrt{\delta} \ll 1/3\sqrt{\delta}.$$

Combining two bounds establishes that $1-\xi \geq \nu^4\delta(1-\rho)/512$ on the intersection of events (i)-(vi); summing the failure probabilities shows that the bound holds with probability at least $1 - 9\exp(-2\varepsilon^*m)$.

We conclude the proof by establishing the probabilistic bounds (i)-(vi). It is convenient to note that since $\Upsilon(t) \leq t$, the choice of C_s in (8) and the fact that $C_\mu \geq 1$, $\delta \geq 1$ imply that

$$\sqrt{C_s} \leq C_\mu\sqrt{C_s} < 2^{-8}(1-\rho). \quad (36)$$

Furthermore, using (6) and the fact that $C_\mu \geq 1$, $\delta \geq 1$,

$$C_s \log\left(\frac{\delta}{C_s}\right) \leq 2^{-17}C_\mu^{-4}\delta^{-1}(1-\rho)^2 \leq 2^{-17}(1-\rho)^2$$

and so

$$\sqrt{2C_s \log\left(\frac{\delta}{C_s}\right)} \leq 2^{-8}(1-\rho). \quad (37)$$

We note three remaining bounds of interest. First, notice from (9) that $m^* = 2^{17}/(1-\rho)C_s$, and so using (36),

$$m^* \geq 2^{33}(1-\rho)^{-3} > 2^{33}(1-\rho)^{-2}. \quad (38)$$

Furthermore, from (9), and (36)

$$\nu^* \leq 2^{-12}\sqrt{C_s/\delta} \leq 2^{-12}\sqrt{C_s} \leq 2^{-20}(1-\rho),$$

Since in (9), $\varepsilon^* < \nu^2$, we have that $\varepsilon^* < 2^{-40}(1-\rho)^2$. In (9), $\alpha^* = \Upsilon[\varepsilon^*/\delta] \leq \varepsilon^*/\delta \leq \varepsilon^*$. Hence, $\alpha^* \leq 2^{-40}(1-\rho)^2$, and for any $m \geq m^*$,

$$\begin{aligned} \sqrt{\alpha^* + \frac{1}{m}} &\leq \sqrt{\alpha^*} + \sqrt{\frac{1}{m^*}} \\ &\leq 2^{-20}(1-\rho) + 2^{-33/2}(1-\rho) < 2^{-8}(1-\rho). \end{aligned} \quad (39)$$

Finally, it is convenient to record the weak bound

$$\sqrt{4\varepsilon^*} < 2^{-8}(1-\rho). \quad (40)$$

(i). Notice that U is an $(n-k_1) \times (m-k_2)$ iid $\mathcal{N}(0, 1/m)$ matrix. Applying (22), with probability at least $1 - \exp(-2\varepsilon^*m)$,

$$\|U\| \leq \sqrt{\delta - \alpha^*} + \sqrt{1-\rho} + \sqrt{4\varepsilon^*}.$$

Since $\sqrt{1-\rho} < 1 \leq \sqrt{\delta}$ and (applying (40)) $\sqrt{4\varepsilon^*} \leq 2^{-8}(1-\rho) \ll 1 \leq \sqrt{\delta}$, on the above good event $\|U\| \leq 3\sqrt{\delta}$.

(ii). Similarly, with probability at least $1 - \exp(-2\varepsilon^*m)$,

$$\begin{aligned} \sigma_{\min}(U) &\geq \sqrt{\delta - \alpha^*} - \sqrt{1-\rho} - \sqrt{4\varepsilon^*} \\ &\geq \sqrt{\delta} - \sqrt{\alpha^*} - \sqrt{1-\rho} - \sqrt{4\varepsilon^*} \end{aligned}$$

Since $\rho \geq 15/16$, $\sqrt{1-\rho} \leq 1/4 \leq \sqrt{\delta}/4$. Similarly, $\sqrt{4\varepsilon^*} \leq 2^{-8}(1-\rho) \leq 2^{-8}\sqrt{\delta}$. Finally, by (39), $\sqrt{\alpha^*} \leq 2^{-8} \leq 2^{-8}\sqrt{\delta}$ and so on the above good event,

$$\sigma_{\min}(U) \geq \sqrt{\delta} \times (1 - 2^{-8} - 1/4 - 2^{-8}) > \sqrt{\delta}/2.$$

(iii). Similarly, applying (26) to U^* gives that with probability at least $1 - \exp(-2\varepsilon^*m)$,

$$\begin{aligned} \gamma_s(U^*) &\geq \sqrt{1-\rho} - \sqrt{C_s} - \sqrt{2C_s \log\left(\frac{\delta}{C_s}\right)} - \sqrt{4\varepsilon^*} \\ &> \sqrt{1-\rho} - 3 \times 2^{-8} \times (1-\rho), \\ &> (1 - 3 \times 2^{-8})\sqrt{1-\rho} > \sqrt{(1-\rho)/2}, \end{aligned}$$

where above we have used (36), (37) and (40).

(iv). Using (27), with probability at least $1 - \exp(-2\varepsilon^*m)$,

$$\begin{aligned} \phi_s(U^*) &\leq 2\sqrt{C_s} + \sqrt{2C_s \log\left(\frac{\delta}{C_s}\right)} + 2C_s \log\left(\frac{1-\rho}{C_s}\right) + 4\varepsilon^* \\ &< 2\sqrt{C_s} + 2\sqrt{2C_s \log\left(\frac{\delta}{C_s}\right)} + \sqrt{4\varepsilon^*} \\ &\leq 5 \times 2^{-8}(1-\rho) < (1-\rho)/4. \end{aligned}$$

Above we have used $\log\left(\frac{1-\rho}{C_s}\right) \leq \log\left(\frac{\delta}{C_s}\right)$ (which follows since $\delta \geq 1 > 1-\rho$), and have applied (36), (37) and (40).

(v). Next consider the term $P_W U^*$. The matrix P_W projects the columns of U^* onto the independent $k_1 + 1$ -dimensional subspace $\text{range}(W^*)$. Since W and U are probabilistically independent, $\beta_s(P_W U^*)$ is distributed as $\beta_s(M)$, where M is a $(k_1 + 1) \times (n - k_1)$ iid $\mathcal{N}(0, 1/m)$ matrix. In particular, using (25), with probability at least $1 - \exp(-2\varepsilon^*m)$,

$$\begin{aligned} \beta_s(P_W U^*) &\leq \sqrt{\frac{k_1+1}{m}} + \sqrt{\frac{s}{m}} + \sqrt{2\frac{s}{m} \log\left(\frac{n-k_1}{s}\right)} + 4\varepsilon^* \\ &\leq \sqrt{\alpha^* + \frac{1}{m}} + \sqrt{C_s} + \sqrt{2C_s \log\left(\frac{\delta}{C_s}\right)} + \sqrt{4\varepsilon^*} \\ &\leq 4 \times 2^{-8}(1-\rho) = 2^{-6}(1-\rho) < 2^{-3}\sqrt{1-\rho}, \end{aligned}$$

where we have used (39), (36), (37) and (40) to bound the four terms.

(vi). Lemma 11, proved in the appendix, shows that with probability at least $1 - 3\exp(-2\varepsilon^*m)$, $\sigma_{\min}(W) > 1/4$. Whenever $\sigma_{\min}(W) > 0$, W has full row rank $k + 1$ and by Lemma 8,

$$\beta_s(P_W) \leq \beta_s(W)/\sigma_{\min}(W). \quad (41)$$

We will next show that $\beta_s(W) < \sqrt{1-\rho}/32$ with probability at least $1 - \exp(-2\varepsilon^*m)$; combining this with our bound on $\sigma_{\min}(W)$ will establish (vi). From the definition of W and (61),

$$\beta_s(W) \leq \frac{\beta_s((Z_{J^c, I})^*)}{\sqrt{1-\rho}} + \frac{\beta_s(\boldsymbol{\mu}_{J^c}^*)}{\sqrt{1-\rho}}. \quad (42)$$

With probability at least $1 - \exp(-2\varepsilon^*m)$,

$$\begin{aligned} \beta_s((Z_{J^c, I})^*) &\leq \sqrt{\frac{k_1}{m}} + \sqrt{\frac{s}{m}} + \sqrt{\frac{2s}{m} \log\left(\frac{(1-\rho)m}{s}\right)} + 4\varepsilon^* \\ &\leq \sqrt{\alpha^*} + \sqrt{C_s} + \sqrt{2C_s \log\left(\frac{1-\rho}{m}\right)} + \sqrt{4\varepsilon^*}, \\ &\leq 4 \times 2^{-8}(1-\rho) = 2^{-6}(1-\rho). \end{aligned}$$

For the second term in (42), notice that $\boldsymbol{\mu}_{J^c}^*$ is simply a row vector of length $m - k_2$, and so $\beta_s(\boldsymbol{\mu}_{J^c}^*)$ is simply the largest Euclidean norm of any vector formed from s elements of $\boldsymbol{\mu}_{J^c}$:

$$\begin{aligned} \beta_s(\boldsymbol{\mu}_{J^c}^*) &\leq \beta_s(\boldsymbol{\mu}^*) = \sup_{|L| \leq s} \|\boldsymbol{\mu}_L\| = C_\mu \sqrt{s/m} \\ &= C_\mu \sqrt{C_s} \leq 2^{-8}(1-\rho), \end{aligned}$$

where we have used that $\|\boldsymbol{\mu}\|_\infty \leq C_\mu m^{-1/2}$, and invoked (36) to bound $C_\mu \sqrt{C_s}$. Hence, with overall probability at least $1 - \exp(-2\varepsilon^*m)$,

$$\beta_s(W) \leq 2^{-6}\sqrt{1-\rho} + 2^{-8}\sqrt{1-\rho} < 2^{-5}\sqrt{1-\rho}.$$

This completes the proof. \blacksquare

C. Initial certificate

In this section, we analyze the guess \mathbf{q}_0 at a certificate of optimality, obtained as the minimum ℓ^2 -norm solution to the equation (18). We upper bound both $\|\mathbf{q}_0\|_2$ and $\|\mathcal{S}\mathbf{q}_0\|_2$. These bounds provide the second half of the conditions needed in Lemma 2 to show that \mathbf{q}_0 can be refined to a true separator via the construction in Lemma 2.

Lemma 4: Fix any $\delta \geq 1$, $\rho < 1$. Suppose that $\nu \leq \nu_{\delta, \rho}^*$ and $m \geq m_{\delta, \rho, \nu}^*$. Fix any subset $I \subset [n]$ of size

$$|I| \leq \alpha_{\delta, \rho, \nu}^* m. \quad (43)$$

Then the minimum ℓ^2 -norm solution \mathbf{q}_0 to the system of equations (18) satisfies

$$\begin{aligned} \|\mathbf{q}_0\|_2 &\leq C\nu\sqrt{\delta m} + O(1), \\ \|\mathcal{S}\mathbf{q}_0\|_2 &\leq C\nu\sqrt{\delta m} \exp(-C'\delta^{-1}\nu^{-2}) + O(1). \end{aligned} \quad (44)$$

simultaneously for all \mathbf{x}_0 with $\text{supp}(\mathbf{x}_0) \subseteq I$, with probability at least $1 - 22 \exp(-2\varepsilon_{\delta, \rho, \nu}^* m)$ in Z, J, σ .

Here, C and C' are numerical constants. In particular, the result holds with $C = 312$ and $C' = 1/48^2$. With these choices, the $O(1)$ term is upper bounded by $48(1-\rho)^{-1/2}$.

The proof of Lemma 4 is largely based on the following general concentration lemma for functions of iid Rademacher sequences. Its proof is straightforward; we delay it to the appendix.

Lemma 5: Fix any matrix $M \in \mathbb{R}^{m \times n}$ and let $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$ be a sequence of iid Rademacher (symmetric

± 1) random variables. Let \mathcal{S} denote the shrinkage operator with threshold $1/2$, defined in (10). Then for any $\varepsilon > 0$,

$$\mathbb{P}[\|\mathcal{S}M\boldsymbol{\sigma}\| > 2\|M\|_F + 4\|M\|\sqrt{\varepsilon n}] < 4\exp(-\varepsilon n). \quad (45)$$

$$\mathbb{P}[\|\mathcal{S}[M\boldsymbol{\sigma}]\| > 2\|M\|_F \exp(-1/4\|M\|_{2, \infty}^2) + 4\|M\|\sqrt{\varepsilon n}] < 4\exp(-\varepsilon n). \quad (46)$$

With Lemma 5 in hand, to establish Lemma 4 we need only estimate the relevant norms for our system of equations:

Proof: (of Lemma 4) We first note that by (9), $\alpha^* = \nu^2 \Upsilon[\varepsilon^*/\delta] \leq \nu^2 \varepsilon^*/\delta \leq \frac{1}{2} \nu^4 \tau^2 (1-\rho)/\delta$, where we have used that $\Upsilon[t] \leq t$ for all t . Using that $\delta \geq 1$, we record the slightly weaker bound

$$\alpha^* \leq \nu^4 \tau^2 (1-\rho). \quad (47)$$

Hence, it is easy to see that $\alpha^* \ll 1$.

Conditioning of Φ . Whenever $\Phi\Phi^*$ is invertible, the pseudoinverse operator

$$\Phi^\dagger \doteq \Phi^*(\Phi\Phi^*)^{-1}$$

is well-defined, and $\|\Phi^\dagger\| = 1/\sigma_{\min}(\Phi)$. We will let \mathcal{E}_Φ denote the event

$$\mathcal{E}_\Phi = \{\Phi\Phi^* \text{ is invertible, and } \|\Phi^\dagger\| \leq 8\}.$$

Notice that from (17),

$$\sigma_{\min}(\Phi) \geq \min(1, \sigma_{\min}(W)) - \nu\|U\|.$$

Lemma 11, proved in the appendix, shows that with probability at least $1 - 3 \exp(-2\varepsilon^*m)$, $\sigma_{\min}(W) \geq 1/4$. Applying (22) to U shows that with probability $1 - \exp(-2\varepsilon^*m)$,

$$\nu\|U\| \leq \nu(\sqrt{\delta} + \sqrt{1-\rho} + \sqrt{4\varepsilon^*}) \leq 3\nu\sqrt{\delta} \leq 1/8,$$

where we have used that

$$\nu^* < 1/24\sqrt{\delta}.$$

So, w.p. $\geq 1 - 4 \exp(-2\varepsilon^*m)$, $\sigma_{\min}(\Phi) \geq 1/8$, and

$$\mathbb{P}[\mathcal{E}_\Phi] \geq 1 - 4 \exp(-2\varepsilon^*m).$$

Bounding the solution. Hence, on \mathcal{E}_Φ , the minimum ℓ^2 norm solution \mathbf{q}_0 to (18) is given by

$$\mathbf{q}_0 = \nu\Phi^\dagger R\boldsymbol{\sigma} + \Phi^\dagger T\boldsymbol{\sigma} + \Phi^\dagger \mathbf{z}.$$

We define four bad events, three of which involve the norms of the elements of \mathbf{q}_0 deviating above their expectations, and the one of which involves the norm of a soft-thresholded version of the first component deviating above its expectation:

$$\begin{aligned} \mathcal{E}_1 &= \left\{ \sigma_{\min}(\Phi) = 0, \text{ or } \|\nu\Phi^\dagger R\boldsymbol{\sigma}\| > 128\nu\sqrt{\delta m} \right\}, \\ \mathcal{E}_2 &= \left\{ \sigma_{\min}(\Phi) = 0, \text{ or } \|\mathcal{S}[\nu\Phi^\dagger R\boldsymbol{\sigma}]\| > 128\nu\tau\sqrt{\delta m} \right\}, \\ \mathcal{E}_3 &= \left\{ \sigma_{\min}(\Phi) = 0, \text{ or } \|\Phi^\dagger T\boldsymbol{\sigma}\| > 176\nu\tau\sqrt{\delta m} + \frac{48}{\sqrt{1-\rho}} \right\}, \\ \mathcal{E}_4 &= \left\{ \sigma_{\min}(\Phi) = 0, \text{ or } \|\Phi^\dagger \mathbf{z}\| > 8\nu\tau\sqrt{\delta m} \right\}. \end{aligned}$$

We will see below that the forms of the bounds in these events arise naturally from the structure of (18); the numerical constants however are pessimistic. On $\cap_i \mathcal{E}_i^c$, using that $\tau \leq 1$

$$\begin{aligned} \|\mathbf{q}_0\| &\leq \|\nu\Phi^\dagger R\boldsymbol{\sigma}\| + \|\Phi^\dagger T\boldsymbol{\sigma}\| + \|\Phi^\dagger \mathbf{z}\| \\ &\leq 312\nu\sqrt{\delta m} + O(1). \end{aligned}$$

Similarly, noting that for any vectors \mathbf{x}, \mathbf{y} $\|\mathcal{S}[\mathbf{x} - \mathbf{y}]\| \leq \|\mathcal{S}[\mathbf{x}]\| + \|\mathbf{y}\|$, on $\cap_i \mathcal{E}_i^c$,

$$\begin{aligned} \|\mathcal{S}[\mathbf{q}_0]\| &\leq \|\mathcal{S}[\nu\Phi^\dagger R\boldsymbol{\sigma}]\| + \|\Phi^\dagger T\boldsymbol{\sigma}\| + \|\Phi^\dagger \mathbf{z}\| \\ &\leq 312\nu\tau\sqrt{\delta m} + O(1). \end{aligned}$$

We next bound the probability of these four bad events.

\mathcal{E}_1 : bounding $\|\nu\Phi^\dagger R\boldsymbol{\sigma}\|$. With probability at least $1 - 2\exp(-2\varepsilon^*m)$, the following two bounds hold:

$$\|R\| \leq \sqrt{\delta} + \sqrt{\rho} + \sqrt{4\varepsilon^*} \leq 3\sqrt{\delta}, \quad (48)$$

$$\|R\|_F \leq \sqrt{\delta\rho m} + \sqrt{4\varepsilon^*} \leq 2\sqrt{\delta m}, \quad (49)$$

where we have used the weak bound $\sqrt{4\varepsilon^*} < 1 \leq \sqrt{\delta}$ (established, e.g., in (40)) and $\sqrt{\rho} \leq 1 \leq \sqrt{\delta}$. Let \mathcal{E}_R denote the event that (48)-(49) hold. On $\mathcal{E}_\Phi \cap \mathcal{E}_R$,

$$\begin{aligned} 2\|\nu\Phi^\dagger R\|_F + 4\|\nu\Phi^\dagger R\|\sqrt{2\varepsilon^*m} \\ \leq 2\nu\|\Phi^\dagger\|\|R\|_F + 4\nu\|\Phi^\dagger\|\|R\|\sqrt{2\varepsilon^*m} \\ \leq 32\nu\sqrt{\delta m} + 96\nu\sqrt{2\varepsilon^*m} < 128\nu\sqrt{\delta m}, \end{aligned}$$

where again we have used that $\sqrt{2\varepsilon^*} < 1$. Then,

$$\begin{aligned} \mathbb{P}[\mathcal{E}_1] &= \mathbb{P}[\mathcal{E}_1 \mid \mathcal{E}_\Phi \cap \mathcal{E}_R] \mathbb{P}[\mathcal{E}_\Phi \cap \mathcal{E}_R] \\ &\quad + \mathbb{P}[\mathcal{E}_1 \mid (\mathcal{E}_\Phi \cap \mathcal{E}_R)^c] \mathbb{P}[(\mathcal{E}_\Phi \cap \mathcal{E}_R)^c] \\ &\leq \mathbb{P}[\mathcal{E}_1 \mid \mathcal{E}_\Phi \cap \mathcal{E}_R] + \mathbb{P}[\mathcal{E}_\Phi^c \cup \mathcal{E}_R^c] \\ &\leq \sup_{Z_0, J_0 \in \mathcal{E}_\Phi \cap \mathcal{E}_R} \mathbb{P}[\mathcal{E}_1 \mid Z_0, J_0] + \mathbb{P}[\mathcal{E}_\Phi^c] + \mathbb{P}[\mathcal{E}_R^c] \\ &\leq \sup_{Z_0, J_0 \in \mathcal{E}_\Phi \cap \mathcal{E}_R} \mathbb{P}[\mathcal{E}_1 \mid Z_0, J_0] + 5\exp(-2\varepsilon^*m). \quad (50) \end{aligned}$$

For any $(Z_0, J_0) \in \mathcal{E}_\Phi \cap \mathcal{E}_R$, applying Lemma 5 with $\Phi^\dagger R$ fixed gives

$$\begin{aligned} \mathbb{P}[\mathcal{E}_1 \mid Z_0, J_0] &= \mathbb{P}\left[\|\nu\Phi^\dagger R\boldsymbol{\sigma}\| > 128\nu\sqrt{\delta m} \mid Z_0, J_0\right] \\ &\leq \mathbb{P}\left[\|\nu\Phi^\dagger R\boldsymbol{\sigma}\| > 2\|\nu\Phi^\dagger R\|_F + 4\|\nu\Phi^\dagger R\|\sqrt{2\varepsilon^*m} \mid Z_0, J_0\right] \\ &\leq \exp(-2\varepsilon^*m), \end{aligned}$$

and $\mathbb{P}[\mathcal{E}_1] \leq 6\exp(-2\varepsilon^*m)$.

\mathcal{E}_2 : bounding $\|\mathcal{S}[\nu\Phi^\dagger R\boldsymbol{\sigma}]\|$. By the same calculation as (50)

$$\mathbb{P}[\mathcal{E}_2] \leq \sup_{Z_0, J_0 \in \mathcal{E}_\Phi \cap \mathcal{E}_R} \mathbb{P}[\mathcal{E}_2 \mid Z_0, J_0] + 5\exp(-2\varepsilon^*m).$$

Since

$$\|\nu\Phi^\dagger R\|_{2,\infty} \leq \|\nu\Phi^\dagger R\| \leq \nu\|\Phi^\dagger\|\|R\|,$$

for any fixed $Z_0, J_0 \in \mathcal{E}_\Phi \cap \mathcal{E}_R$,

$$\begin{aligned} \exp\left(-\frac{1}{4\|\nu\Phi^\dagger R\|_{2,\infty}^2}\right) &\leq \exp\left(-\frac{1}{4\nu^2\|\Phi^\dagger\|^2\|R\|^2}\right) \\ &\leq \exp\left(-\frac{1}{48^2\delta\nu^2}\right) = \tau, \end{aligned}$$

and further

$$\begin{aligned} 2\|\nu\Phi^\dagger R\|_F \exp\left(-\frac{1}{4\|\nu\Phi^\dagger R\|_{2,\infty}^2}\right) + 4\|\nu\Phi^\dagger R\|\sqrt{2\varepsilon^*m} \\ \leq 32\nu\tau\sqrt{\delta m} + 96\nu\sqrt{2\varepsilon^*m} \leq 128\nu\tau\sqrt{\delta m}, \end{aligned}$$

where in the last step we have used that $\delta \geq 1$ and that by (9) $\varepsilon^* < \tau^2/2$. Applying Lemma 5,

$$\begin{aligned} \mathbb{P}[\mathcal{E}_2 \mid Z_0, J_0] &= \mathbb{P}[\|\mathcal{S}[\nu\Phi^\dagger R\boldsymbol{\sigma}]\| > 128\nu\tau\sqrt{\delta m} \mid Z_0, J_0] \\ &\leq \mathbb{P}\left[\|\mathcal{S}[\nu\Phi^\dagger R\boldsymbol{\sigma}]\| > \left(2\|\nu\Phi^\dagger R\|_F e^{-\frac{1}{4\|\nu\Phi^\dagger R\|_{2,\infty}^2}} + 4\|\nu\Phi^\dagger R\|\sqrt{2\varepsilon^*m}\right) \mid Z_0, J_0\right] \\ &\leq \exp(-2\varepsilon^*m). \end{aligned}$$

We conclude that $\mathbb{P}[\mathcal{E}_2] < 6\exp(-2\varepsilon^*m)$.

\mathcal{E}_3 : bounding $\|\Phi^\dagger T\boldsymbol{\sigma}\|$. Notice that if $\|\cdot\|_\diamond$ is either the Frobenius or operator norm,

$$\|T\|_\diamond \leq \frac{\|Z_{J,I}\|_\diamond}{\sqrt{1-\rho}} + \frac{\|\boldsymbol{\mu}_J\|}{\sqrt{1-\rho}} \leq \frac{\|Z_{J,I}\|_\diamond}{\sqrt{1-\rho}} + \frac{1}{\sqrt{1-\rho}}.$$

Notice that with probability at least $1 - 2\exp(-2\varepsilon^*m)$,

$$\begin{aligned} \|Z_{J,I}\| &\leq \sqrt{\rho} + \sqrt{\alpha^*} + \sqrt{4\varepsilon^*} \leq 3, \\ \|Z_{J,I}\|_F &\leq \sqrt{\alpha^*\rho m} + \sqrt{4\varepsilon^*} \leq \sqrt{\alpha^*m} + 1. \end{aligned}$$

The final bounds above are loose, but only impact the numerical constants below. Since (47) implies that $\alpha^* \leq \nu^2\tau^2$, with probability at least $1 - 2\exp(-2\varepsilon^*m)$,

$$\|T\| \leq \frac{\|Z_{J,I}\|}{\sqrt{1-\rho}} + \frac{1}{\sqrt{1-\rho}} \leq \frac{5}{\sqrt{1-\rho}}, \quad (51)$$

$$\|T\|_F \leq \frac{\|Z_{J,I}\|_F}{\sqrt{1-\rho}} + \frac{1}{\sqrt{1-\rho}} \leq \nu\tau\sqrt{m} + \frac{3}{\sqrt{1-\rho}}. \quad (52)$$

Let \mathcal{E}_T denote the event that (51)-(52) hold. On $\mathcal{E}_T \cap \mathcal{E}_\Phi$,

$$\begin{aligned} 2\|\Phi^\dagger T\|_F + 4\|\Phi^\dagger T\|\sqrt{2\varepsilon^*m} \\ \leq 2\|\Phi^\dagger\|\|T\|_F + 4\|\Phi^\dagger\|\|T\|\sqrt{2\varepsilon^*m} \\ \leq 16\nu\tau\sqrt{m} + 48(1-\rho)^{-1/2} + 160(1-\rho)^{-1/2}\sqrt{2\varepsilon^*m}, \\ \leq 176\nu\tau\sqrt{\delta m} + 48(1-\rho)^{-1/2}, \end{aligned}$$

where we have used (9) to bound $2\varepsilon^* \leq \nu^2\tau^2(1-\rho)$. Then,

$$\begin{aligned} \mathbb{P}[\mathcal{E}_3] &\leq \sup_{Z_0, J_0 \in \mathcal{E}_\Phi \cap \mathcal{E}_T} \mathbb{P}[\mathcal{E}_3 \mid Z_0, J_0] + \mathbb{P}[\mathcal{E}_\Phi^c] + \mathbb{P}[\mathcal{E}_T^c] \\ &\leq \sup_{Z_0, J_0 \in \mathcal{E}_\Phi \cap \mathcal{E}_T} \mathbb{P}[\mathcal{E}_3 \mid Z_0, J_0] + 5\exp(-2\varepsilon^*m). \end{aligned}$$

Applying Lemma 5 with $\Phi^\dagger T$ fixed,

$$\begin{aligned} \mathbb{P}[\mathcal{E}_3 \mid Z_0, J_0] \\ \leq \mathbb{P}\left[\|\Phi^\dagger T\boldsymbol{\sigma}\| \geq 2\|\Phi^\dagger T\|_F + 4\|\Phi^\dagger T\|\sqrt{\varepsilon^*m} \mid Z_0, J_0\right] \\ \leq \exp(-2\varepsilon^*m). \end{aligned}$$

and $\mathbb{P}[\mathcal{E}_3] \leq 6\exp(-2\varepsilon^*m)$.

\mathcal{E}_4 : bounding $\|\Phi^\dagger \mathbf{z}\|$. Finally, on \mathcal{E}_Φ ,

$$\|\Phi^\dagger \mathbf{z}\| \leq \|\Phi^\dagger\|\|\mathbf{z}\| \leq \frac{8}{\nu}\sqrt{\frac{\alpha^*m}{1-\rho}} \leq 8\nu\tau\sqrt{m}.$$

Here we have used that by (47), $\alpha^* \leq \nu^4\tau^2(1-\rho)$. So, $P[\mathcal{E}_4] \leq P[\mathcal{E}_\Phi^c] \leq 4\exp(-2\varepsilon^*m)$. Summing the failure probabilities establishes the lemma. \blacksquare

D. Proof of Theorem 1

Proof: We first prove our result under the assumption that the matrix A is wide ($\delta \geq 1$) and the error e is not too sparse ($\rho \geq 15/16$). We then show that this more difficult case implies the result for arbitrary δ, ρ .

Underdetermined A , dense errors. Fix any $\delta \geq 1$ and $\rho \geq 15/16$. Fix any $\nu \leq \nu_{\delta, \rho}^*$, and set $\alpha^*, \varepsilon^*, m^*$ as specified in (9). Suppose that $m \geq m^*$, and let $k_1 = \lfloor \alpha^* m \rfloor$. Fix $C_s > 0$ according to (8) and set $s = \lfloor C_s m \rfloor$. Notice that these choices are consistent with assumptions and conclusions of Lemmas 3 and 4.

For each $I \in \binom{[n]}{k_1}$, let \mathcal{E}_I be the event that:

- The matrix Φ has $\xi_s(\Phi)$ bounded as in Lemma 3.
- For all \mathbf{x}_0 with $\text{supp}(\mathbf{x}_0) \subseteq I$, the minimum ℓ^2 norm solution \mathbf{q}_0 to (18) has $\|\mathbf{q}_0\|$ and $\|\mathcal{S}\mathbf{q}_0\|$ bounded as in Lemma 4.
- $\text{rank}(A_{J^c, I}) = k_1$.

Now, $A_{J^c, I}$ has full rank k_1 with probability one. Hence, by Lemmas 3 and 4

$$\forall I \in \binom{[n]}{k_1}, \quad \mathbb{P}[\mathcal{E}_I^c] \leq 31 \exp(-2\varepsilon^* m). \quad (53)$$

On \mathcal{E}_I , for all \mathbf{x}_0 supported on I , \mathbf{q}_0 satisfies

$$\|\mathbf{q}_0\| + \frac{\|\mathcal{S}\mathbf{q}_0\|}{1 - \xi_s(\Phi)} \leq 312 \nu \sqrt{\delta m} \left(1 + \frac{512 \tau}{\nu^4 \delta (1 - \rho)} \right) + \frac{48}{\sqrt{1 - \rho}}.$$

Below, we show that for any $\nu \leq \nu^*$,

- (i). $\frac{512 \tau}{\nu^4 \delta (1 - \rho)} \leq 1$,
- (ii). $2 \cdot 312 \cdot \nu \sqrt{\delta m} < \frac{1}{4} \sqrt{C_s m}$,
- (iii). $\frac{48}{\sqrt{1 - \rho}} < \frac{1}{4} \sqrt{C_s m}$.

Hence, on \mathcal{E}_I , for all \mathbf{x}_0 supported on I

$$\|\mathbf{q}_0\| + \frac{\|\mathcal{S}\mathbf{q}_0\|}{1 - \xi_s(\Phi)} < \sqrt{C_s m}/2 = \sqrt{s}/2. \quad (54)$$

So, by Lemma 2, there is a solution \mathbf{q}^* to (18) with $\|\mathbf{q}^*\|_\infty \leq 1/2 < 1$. Lemma 1 then implies that on \mathcal{E}_I , for all \mathbf{x}_0 with $\text{supp}(\mathbf{x}_0) \subseteq I$, $(\mathbf{x}_0, \mathbf{e}_0)$ is the unique solution to

$$\min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad A\mathbf{x} + \mathbf{e} = A\mathbf{x}_0 + \mathbf{e}_0.$$

Now, further notice that by our choice of α^* in (9),

$$\alpha^* < \Upsilon[\varepsilon^*/\delta] \leq \delta \Upsilon[\varepsilon^*/\delta].$$

Hence, by (6)

$$\alpha^* \log(\delta/\alpha^*) \leq \varepsilon^*,$$

and

$$\begin{aligned} \mathbb{P}[\cap_I \mathcal{E}_I] &\geq 1 - \sum_I \mathbb{P}[\mathcal{E}_I^c], \\ &\geq 1 - \# \binom{[n]}{k_1} \times 31 \exp(-2\varepsilon^* m), \\ &\geq 1 - 31 \exp(\alpha^* \log(\delta/\alpha^*) m - 2\varepsilon^* m), \\ &\geq 1 - 31 \exp(-\varepsilon^* m), \end{aligned}$$

Hence, to establish the result for $\delta \geq 1$ and $\rho \in [15/16, 1)$, it just remains to prove (i)-(iii).

(i). We establish (i) by proving two intermediate bounds (55) and (56) below hold whenever $\nu \leq \nu^*$.

$$\sqrt{\tau} \leq \nu^4. \quad (55)$$

To show (55), it is enough to argue that $\nu^{-1} \tau^{1/8} \leq 1$. Now,

$$\tau^{1/8} = \exp(-8^{-1} 48^{-2} \delta^{-1} \nu^{-2}).$$

By (9), for $\nu \leq \nu^*$,

$$\nu \leq 2^{-12} \sqrt{C_s/\delta} < 2^{-20} \delta^{-1} < 8^{-1} 48^{-2} \delta^{-1},$$

where in the second inequality we have used that $C_s \leq 2^{-17} C_\mu^{-2} \delta^{-2} (1 - \rho) \ll 2^{-16} \delta^{-1}$ (this follows since $\Upsilon[t] \leq t$). So $\tau^{1/8} < \exp(-\nu^{-1})$ and

$$\nu^{-1} \tau^{1/8} < \nu^{-1} \exp(-\nu^{-1}) \leq 1/e < 1,$$

showing (55). Above, we used that $\max_x x e^{-x} = e^{-1}$.

$$\sqrt{\tau} \leq \frac{\delta(1 - \rho)}{512}. \quad (56)$$

To show this, notice that $\frac{1}{2} 48^{-2} \geq 2^{-13}$, and so

$$\begin{aligned} \sqrt{\tau} &\leq \exp(-2^{-13} \delta^{-1} \nu^{-2}) \\ &\leq \exp(-2^{11}/C_s) \leq \exp(-2^{28}/(1 - \rho)), \end{aligned}$$

where we have used that by (9), $\nu^{-1} \geq 2^{12} \sqrt{\delta/C_s}$, and $C_s \leq 2^{-17} (1 - \rho)$. Consider

$$h(x) \doteq 512 x \exp(-2^{28} x).$$

It is easy to show that $\max_x h(x) = 2^{-19} e^{-1}$, and so

$$\frac{512}{1 - \rho} \sqrt{\tau} \leq h\left(\frac{1}{1 - \rho}\right) \leq 2^{-19} e^{-1} \ll \delta,$$

establishing (56).

(ii). For $\nu \leq \nu^*$, by (9)

$$2 \cdot 312 \cdot \nu \sqrt{\delta m} \leq 624 \cdot 2^{-12} \sqrt{C_s m} < \frac{1}{4} \sqrt{C_s m}. \quad (57)$$

(iii). For $m \geq m^*$, by (9),

$$\frac{1}{4} \sqrt{C_s m} \geq \frac{1}{4} \sqrt{\frac{2^{17}}{1 - \rho}} > \frac{48}{\sqrt{1 - \rho}}.$$

This establishes the theorem for $\delta \geq 1$, $\rho \in [15/16, 1)$.

Extension to arbitrary $\delta > 0, \rho \in (0, 1)$. Now, fix any $\delta > 0$. Replace δ in the above choice of constants with $\max(\delta, 1)$. Generate

$$\tilde{A} = \mu \mathbf{1}^* + \nu \tilde{Z} \in \mathbb{R}^{m \times m}$$

according to the bouquet model. Then with overwhelming probability in \tilde{Z}, J, σ for any \mathbf{e}_0 with signs and support J, σ and all sufficiently sparse $\tilde{\mathbf{x}}_0$,

$$(\tilde{\mathbf{x}}_0, \mathbf{e}_0) = \arg \min \|\tilde{\mathbf{x}}\|_1 + \|\mathbf{e}\|_1 \quad \text{s.t.} \quad \tilde{A} \tilde{\mathbf{x}} + \mathbf{e} = \tilde{A} \tilde{\mathbf{x}}_0 + \mathbf{e}_0. \quad (58)$$

and the minimizer is uniquely defined. Form $A \in \mathbb{R}^{m \times \delta m}$ by selecting the first δm columns of \tilde{A} . Clearly, A is distributed according to the bouquet model. Lemma 12 shows that for any $\mathbf{x}_0 \in \mathbb{R}^{\delta m}$, if $[\mathbf{x}_0^* \ 0 \ \dots \ 0]^*$, \mathbf{e}_0 exactly recovered by (58), $(\mathbf{x}_0, \mathbf{e}_0)$ is the unique optimal solution to

$$\min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{s.t.} \quad \tilde{A}\mathbf{x} + \mathbf{e} = \tilde{A}\tilde{\mathbf{x}}_0 + \mathbf{e}_0,$$

and so with at least the same large probability, ℓ^1 -minimization in the reduced system $A\mathbf{x} + \mathbf{e} = \mathbf{y}$ recovers all sparse \mathbf{x}_0 from a random error \mathbf{e}_0 .

Finally, a straightforward argument given in Lemma 13 shows if ℓ^1 -minimization uniquely recovers $(\mathbf{x}_0, \tilde{\mathbf{e}}_0)$, then for any \mathbf{e}_0 with $J \doteq \text{supp}(\mathbf{e}_0) \subseteq \text{supp}(\tilde{\mathbf{e}}_0)$ and $\text{sign}(\mathbf{e}_0(J)) = \text{sign}(\tilde{\mathbf{e}}_0(J))$, then ℓ^1 -minimization uniquely recovers $(\mathbf{x}_0, \mathbf{e}_0)$. Fix any $\rho \in (0, 1)$. Choose constants as above, but with ρ replaced by $\max\{\rho, 15/16\}$. For $\rho < 15/16$, generate the signs and support of \mathbf{e}_0 as follows: generate a random subset $\tilde{J} \subset [m]$ of size $15m/16$, and a random sign pattern $\tilde{\sigma}$ with $\tilde{\sigma}(\tilde{J}^c) = 0$, and $\tilde{\sigma}(\tilde{J})$ iid Rademacher. Then with overwhelming probability, ℓ^1 minimization uniquely recovers $(\mathbf{x}_0, \tilde{\mathbf{e}}_0)$ for all sufficiently sparse \mathbf{x}_0 . Generate J as a random subset of J' of size ρm , and let $\sigma = \tilde{\sigma}(J)$. By Lemma 13, ℓ^1 -minimization, with the same overwhelmingly large probability ℓ^1 -minimization also uniquely recovers $(\mathbf{x}_0, \mathbf{e}_0)$ for all sufficiently sparse \mathbf{x}_0 . ■

IV. SIMULATIONS AND EXPERIMENTS

In this section, we perform simulations verifying the conclusions of Theorem 1, and investigating the effect of various model parameters on the error correction capability of the ℓ^1 -minimization (2). In the simulations below we use the publicly available ℓ^1 -magic package [31], except for one (higher-dimensional) face recognition example, which requires a customized interior point method. Since ℓ^1 -recoverability depends only on the signs and support of $(\mathbf{x}_0, \mathbf{e}_0)$, in the simulations below we choose $\mathbf{x}_0(i) \in \{0, 1\}$ and $\mathbf{e}_0(i) \in \{-1, 0, 1\}$. We will judge an output $(\hat{\mathbf{x}}, \hat{\mathbf{e}})$ to be correct if $\max(\|\hat{\mathbf{x}} - \mathbf{x}_0\|_\infty, \|\hat{\mathbf{e}} - \mathbf{e}_0\|_\infty) < 0.01$.

a) *Comparison with alternative approaches:* We first compare the performance of the extended ℓ^1 -minimization

$$\min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad \mathbf{y} = A\mathbf{x} + \mathbf{e}$$

to two alternative approaches. The first is the error correction approach of [14], which multiplies by a full rank matrix B such that $BA = 0$,⁸ solves

$$\min \|\mathbf{e}\|_1 \quad \text{subject to} \quad B\mathbf{e} = B\mathbf{y},$$

and then subsequently recovers \mathbf{x} from the clean system of equations $A\mathbf{x} = \mathbf{y} - \mathbf{e}$. The second is the Regularized Orthogonal Matching Pursuit (ROMP) algorithm [32], a state-of-the-art greedy method for recovering sparse signals.⁹ For this algorithm, we use the implementation from <http://math.ucdavis.edu/~dneedell/>.

⁸This comparison requires $n \ll m$ although our method is not limited to this case.

⁹For the models considered here, less sophisticated greedy methods such as the standard orthogonal matching pursuit fail even for small error fractions.




Fig. 3. **Comparison with alternative approaches.** Here, we fix $m = 500$, $\delta = 0.25$, $\nu = 0.05$, and $k_1 = 15$, and compare three approaches to recovering the sparse signal \mathbf{x}_0 from error \mathbf{e}_0 . The first, denoted “ $L^1 - [A \ I]$ ” solves the extended ℓ^1 minimization advocated in this work. The second, denoted “ $L^1 - \perp$ comp” premultiplies by the orthogonal complement of A , and then solves an underdetermined system of linear equations for the sparse error \mathbf{e} [14]. The final approach is the greedy Regularized Orthogonal Matching Pursuit (ROMP) [32].

For this experiment, the ambient dimension is $m = 500$; the parameters of the CAB model are $\nu = 0.05$ and $\delta = 0.25$. We fix the signal support to be $k_1 = 15$, and vary the fraction of errors from 0 to 0.95. For each error fraction, we generate 500 independent problems. Figure 3 plots the fraction of successes for each of the three algorithms, as a function of error density ρ . There the extended ℓ^1 -minimization is denoted “ $L^1 - [A \ I]$ ” (red curve), while the alternative approach of [14] is denoted “ $L^1 - \perp$ comp” (blue curve). Whereas both ROMP and the ℓ^1 approach of [14] break down around 40% corruption, the extended ℓ^1 -minimization continues to succeed with high probability even beyond 60% corruption.

b) *Error correction capacity:* While the previous experiment demonstrates the advantages of the extended ℓ^1 -minimization (2) for the CAB model, Theorem 1 suggests that more is true: If the support of \mathbf{x} is a diminishing fraction of the dimension, as the dimension increases, the fraction of errors that the extended ℓ^1 -minimization can correct should approach one. We generate problem instances with $\delta = 0.25$, $\nu = 0.05$, for varying $m = 100, 200, 400, 800, 1600$. For each problem size, and for each error fraction $\rho = 0.05, 0.1, \dots, 0.95$, we generate 500 random problems, and plot the fraction of correct recoveries in Figure 4. At left, we fix $k_1 = 1$, while at right, k_1 grows as $k_1 = m^{1/2}$. In both cases, as m increases, the fraction of errors that can be corrected also increases, with no sign of stopping anywhere before $\rho = 1$.

c) *Varying model parameters:* We next investigate the effect of varying δ (Figure 5 left) and ν (Figure 5 right). We first fix $m = 400$, $\nu = .3$, and consider different bouquet sizes $n = 100, 200, 300, 400, 500$. Figure 5 left plots the fraction of correct trials for varying error densities ρ , for each of these bouquet sizes. For this fixed m , the error correction capability decreases only slightly as n increases.

We next fix $m = 400$, $n = 200$, and consider the effect of varying ν . Figure 5 plots the result for $\nu = .1, .3, .5, .7, .9$. Notice that as ν decreases (i.e., the bouquet becomes tighter), the error correction capacity increases: for any fixed fraction of successful trials, the fraction of error that can be corrected increases by approximately 15% as ν decreases from .9 to .5.




Fig. 4. **Error correction for \mathbf{x} with sublinear support.** We fix $\delta = 0.25$, $\nu = 0.05$, and plot the fraction of successful recoveries as a function of the error density ρ , for each $m = 100, 200, 400, 800, 1600$. At left, k_1 is fixed at 1; at right, $k_1 = m^{1/2}$. In both cases, as m increases, the fraction of errors that can be corrected approaches 1.




Fig. 5. **Effect of varying n and ν .** At left, we fix $m = 400$, $\nu = .3$, and consider varying $n = 100, 200, \dots, 500$. For each of these model settings, we plot the fraction of correct recoveries as a function of the fraction of errors. Notice that the error correction capacity decreases only slightly as n increases. At right, we fix $m = 400$, $n = 200$, and vary ν from .1 to .9. Again, we plot the fraction of correct recoveries for each error fraction. As expected from Theorem 1, as ν decreases, the error correction capacity of ℓ^1 increases.

d) Phase transition in proportional growth: Theorem 1 claims that for any $\rho < 1$, the ℓ^1 -minimization works even when the signal support k_1 grows proportionally to m : $k_1 \leq \alpha m$. However, it gives rather pessimistic estimate on how large the fraction can be. Based on intuition from more homogeneous polytopes (especially the work of Donoho and Tanner on Gaussian matrices [20]), we might expect that when k_1 exhibits proportional growth, an asymptotically sharp phase transition between guaranteed recovery and guaranteed failure will occur at some critical error fraction $\rho^* \in (0, 1)$. We investigate this empirically here by again setting $\delta = 0.25$, $\nu = 0.05$, but this time allowing $k_1 = 0.05m$. Figure 6 plots the fraction of correct recovery for varying error fractions ρ , as m grows: $m = 100, 200, 400, 800, 1600$. In this proportional growth setting, we see an increasingly sharp phase transition, near $\rho = 0.6$.

e) Error correction with real face images: Finally, we return to the motivating example of face recognition under varying illumination and random corruption. For this experiment, we use the Extended Yale B face database [15], which tests illumination sensitivity of face recognition algorithms. As in [11], we form the matrix A from images in Subsets 1 and 2, which contain mild-to-moderate illumination variations. Each column of the matrix A is a $w \times h$ face image, stacked as a vector in \mathbb{R}^m ($m = w \times h$). Here, the weak proportional growth setting corresponds to the case when the total number of image




Fig. 6. **Phase transition in proportional growth.** When the signal support grows in proportion to the dimension ($k_1/m \rightarrow \alpha \in (0, 1)$), we observe an asymptotically sharp phase transition in the probability of correct recovery, similar to that investigated in [20]. Here, for $\delta = 0.25$, $\nu = 0.05$, $k_1 = 0.05m$, we indeed see a sharp phase transition at $\rho = 0.6$.

pixels grows proportionally to the number n of face images. Since the number of images per subject is fixed, this is the same as the total image resolution growing proportionally to the number of subjects. We vary the image resolutions through the range $34 \times 30, 48 \times 42, 68 \times 60, 96 \times 84$.¹⁰ The matrix A is formed from images of 4, 9, 19, 38 subjects, respectively, corresponding to $\delta \approx 0.09$. Here, $\nu \approx 0.3$. In face recognition, the sublinear growth of $\|\mathbf{x}_0\|_0$ comes from the fact that the observation should ideally be a linear combination of only images of the same subject. Various estimates of the required number of images, k_1 , appear in the literature, ranging from 5 to 9. Here, we fix $k_1 = 7$, and generate the (clean) test image synthetically as a linear combination of k_1 training images from a single subject. The reason for using synthetic linear combinations as opposed to real test images is simply that it allows us to verify whether \mathbf{x}_0 was correctly recovered; in the real data experiments of the introduction of this work and of [11], success could only be judged in terms of the recognition rate of the entire classification pipeline.

For each resolution considered, and for each error fraction, we generate 75 trials. Figure 7 (left) plots the fraction of successes as a function of the fraction of corruption. Notice that as predicted by Theorem 1, the fraction of errors that can be corrected again approaches 1 as the data size increases. Figure 7 (right) gives a visual demonstration of the algorithm's capability. In the test images in Figure 7 (right, top), the amount of corruption is chosen to correspond to a 50% probability of success according to the plots in Figure 7 (left). Below each corrupted test image, the "clean" image recovered by our method is shown.

V. DISCUSSIONS AND FUTURE WORK

a) Compressed sensing for signals with varying sparsity:

In the conventional setting for recovering a sparse signal, one often implicitly assumes that each entry of the signal has an equal probability of being nonzero. As a result, one typically requires that the incoherence (or coherence) of the dictionary is somewhat uniform. In this paper, we saw quite a different

¹⁰Thus, the total dimension $m = 1020, 2016, 4080, 8064$ grows roughly by a factor of 2 from one curve to the next, similar to the simulations above.

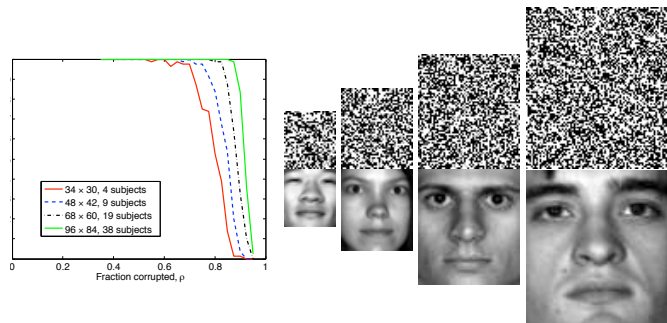


Fig. 7. **Error correction with real face images.** We simulate weak proportional growth in the Extended Yale B face database, with the resolution of the images growing in proportion to the number of subjects. Left: fraction of correct recoveries for varying levels of occlusion. Right: examples of correct recovery for each resolution considered. Top: corrupted test image. The fraction of corruption is chosen so that the probability of correct recovery is 50%. Bottom: clean image, from correctly recovered \mathbf{x}_0 .

example. If we view both \mathbf{x} and \mathbf{e} as the signal that we want to recover, then the sparsity or density of the combined signal is quite uneven – \mathbf{x} is very sparse but \mathbf{e} can be very dense. Nevertheless, our result suggests that if the incoherence of the dictionary is adaptive to the distribution of the density – more coherent for the sparse part and less for the dense part, then ℓ^1 -minimization will be able to recover such uneven signals even if bounds based on the even sparsity assumption suggest otherwise. Thus, if one has some prior knowledge about which part of the signal is likely to be more sparse or more dense, one can achieve much better performance with ℓ^1 -minimization by using a dictionary with matching incoherence. More generally, for any given distribution of sparsity, one may ask the question whether there exists an optimal dictionary with matching incoherence such that ℓ^1 -minimization has the highest chance of success.

b) Stability with respect to noise: Although in our model we do not explicitly consider any noise, say $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e} + \mathbf{z}$, where \mathbf{z} adds small Gaussian noise to each entry, we have observed empirically in our simulations and also in experiments with face images that the ℓ^1 -minimization (2) for the cross-and-bouquet model is surprisingly stable with respect to measurement or numerical noise. In fact, as the method is able to deal with dense errors regardless of their magnitude, large noisy entries in \mathbf{z} might be treated like errors and be absorbed into \mathbf{e} . This is especially the case if the signal \mathbf{x} is extremely sparse – in the case of face recognition, the support of \mathbf{x} is nearly constant. Nevertheless, for less sparse signals \mathbf{x} , the noise tolerance can likely be further improved by relaxing the equality constraint in (2), in a similar spirit to the Lasso [4]–[6], [33] or Dantzig selector [34]:

$$\min_{\mathbf{x}, \mathbf{e}} \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad \|\mathbf{y} - \mathbf{A}\mathbf{x} + \mathbf{e}\| \leq \varepsilon^2.$$

To our best knowledge, a more precise characterization of the effect of entry-wise small noise or perturbation \mathbf{z} on the estimate of \mathbf{x} and \mathbf{e} through the above convex program remains an open problem.

c) Neighborliness of polytopes: As we have seen in this paper, a precise characterization of the performance of ℓ^1 -

minimization requires us to analyze the geometry of polytopes associated with the specific dictionaries in question. In practice, we often use ℓ^1 -minimization for purposes other than signal reconstruction or error correction. For instance, using machine learning techniques, we can learn from exemplars a dictionary that is optimal for certain tasks such as data classification [13]. The polytope associated with such a dictionary may be very different from those that are normally studied in signal processing or coding theory or error correction, leading to qualitatively different behavior of the ℓ^1 -minimization. Thus, we should expect that in the coming years, many new classes of high-dimensional polytopes with even more interesting properties may arise from other applications and practical problems.

d) Dense error correction for low-rank matrices: Recent studies of low-rank matrix completion have shown that it is possible to exactly complete a low-rank matrix X from an observation of only a diminishing fraction of its entries [35]. Moreover, one can exactly recover a low-rank matrix X with grossly corrupted entries $Y = X + E$ by solving a similar convex program:

$$\min \|X\|_* + \lambda\|E\|_1 \quad \text{subject to} \quad Y = X + E,$$

as long as the matrix E is sufficiently sparse [36]. Moreover, empirically one can observe a similar *dense error correction* phenomenon: when the rank of the matrix X is sublinear in its dimension, as the dimension increases, the fraction of entries corrupted by random errors with random signs that the convex program can correct also approaches to 100%. As an extension to the analysis of [35], one can prove this is indeed the case if the signs of the errors in E are chosen at random [37]. It seems that the dense error correction property discovered in this work may be an instance of a more general property of error correction problems involving structured high-dimensional signals and random errors.

e) Modifications of the error model: We have assumed the signs and support of the error \mathbf{e}_0 are random, which plays a crucial role in the proof of the main theorem. In many applications, the errors may not be completely random and have additional structure, or may even be chosen in an adversarial manner. To reduce the randomness in our error model, one can invoke similar derandomization arguments of [35] (Theorem 2.3) and show ℓ^1 minimization successfully recovers \mathbf{e}_0 with random support and signs selected from any pre-specified sign pattern. In many applications, we may also have certain prior knowledge about the support of the errors. For example, in image processing, corruption incurred by occlusion normally has a spatially contiguous support. Recent empirical studies in face recognition [38] have demonstrated that such prior knowledge can significantly improve the error correction capability of the ℓ^1 -minimization. Giving a rigorous mathematical justification for this success remains an open problem.

ACKNOWLEDGMENTS

The authors would like to acknowledge helpful conversations with and useful comments from Prof. Robert Fossum

(UIUC Math), Prof. Olgica Milenkovic (UIUC ECE), Prof. Sean Meyn (UIUC ECE), Prof. Emmanuel Candès (Stanford Statistics), and Dr. Gang Hua. The authors would also like to thank the anonymous reviewers for helping to improve the clarity and presentation of the results. Finally, Yi Ma would like to thank Microsoft Research Asia, Beijing, China, for its hospitality during his visit there in Summer 2008.

APPENDIX TECHNICAL LEMMAS

This appendix proves several lemmas used in the main text.

OBSERVATIONS ON SUBMATRIX SINGULAR VALUES

Lemma 6: For any block matrix $[M \ N]$, and scalar t with $|t| \leq 1$,

$$\gamma_s([M \ tN]) \geq |t|\gamma_s([M \ N]). \quad (59)$$

Proof:

$$\begin{aligned} \gamma_s([M \ tN]) &= \inf_{\substack{\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 \geq 1 \\ \|\mathbf{x}\|_0 + \|\mathbf{y}\|_0 \leq s}} \left\| [M \ tN] \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \right\| \\ &= \inf_{\substack{\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 \geq 1, \\ \|\mathbf{x}\|_0 + \|\mathbf{y}\|_0 \leq s.}} \left\| [M \ N] \begin{bmatrix} \mathbf{x} \\ t\mathbf{y} \end{bmatrix} \right\| \\ &= \inf_{\substack{\|t\mathbf{x}\|^2 + \|\mathbf{y}\|^2 \geq 1, \\ \|\mathbf{x}\|_0 + \|\mathbf{y}\|_0 \leq s.}} \left\| [M \ N] \begin{bmatrix} t\mathbf{x} \\ \mathbf{y} \end{bmatrix} \right\| \\ &\geq |t| \inf_{\substack{\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 \geq 1, \\ \|\mathbf{x}\|_0 + \|\mathbf{y}\|_0 \leq s.}} \left\| [M \ N] \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \right\| = |t|\gamma_s([M \ N]). \end{aligned}$$

Lemma 7: For any matrices M, N ,

$$\gamma_s(M + N) \geq \gamma_s(M) - \beta_s(N), \quad (60)$$

$$\text{and } \beta_s(M + N) \leq \beta_s(M) + \beta_s(N). \quad (61)$$

Proof: Notice that

$$\begin{aligned} \gamma_s(M + N) &= \inf_{\|\mathbf{x}\|=1, \|\mathbf{x}\|_0 \leq s} \|M\mathbf{x} + N\mathbf{x}\| \\ &\geq \inf_{\|\mathbf{x}\|=1, \|\mathbf{x}\|_0 \leq s} (\|M\mathbf{x}\| - \|N\mathbf{x}\|) \\ &\geq \inf_{\|\mathbf{x}\|=1, \|\mathbf{x}\|_0 \leq s} \|M\mathbf{x}\| - \sup_{\|\mathbf{y}\|=1, \|\mathbf{y}\|_0 \leq s} \|N\mathbf{y}\| \\ &= \gamma_s(M) - \beta_s(N). \end{aligned}$$

The calculation for (61) is identical. \blacksquare

Lemma 8: Let $W \in \mathbb{R}^{m \times n}$, ($m < n$) be any matrix with full row rank m , and let $P_W \in \mathbb{R}^{n \times n}$ denote projection operator onto the range of W^* . Then

$$\beta_s(P_W) \leq \beta_s(W) / \sigma_{\min}(W). \quad (62)$$

Proof:

$$\begin{aligned} \beta_s(P_W) &= \sup_{\|\mathbf{x}\| \leq 1, \|\mathbf{x}\|_0 \leq s} \|W^*(WW^*)^{-1}W\mathbf{x}\| \\ &\leq \|W^*(WW^*)^{-1}\| \sup_{\substack{\|\mathbf{x}\| \leq 1, \\ \|\mathbf{x}\|_0 \leq s}} \|W\mathbf{x}\| = \sigma_{\min}^{-1}(W) \beta_s(W). \end{aligned}$$

Lemma 9: For any matrix M

$$\gamma_s^2([\mathbf{I} \ M]) \geq \min\{1, \gamma_s^2(M)\} - \phi_s(M). \quad (63)$$

Proof:

$$\begin{aligned} \gamma_s^2([\mathbf{I} \ M]) &= \inf_{\substack{\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 = 1 \\ \|\mathbf{x}\|_0 + \|\mathbf{y}\|_0 \leq s}} \|\mathbf{x} + M\mathbf{y}\|^2 \\ &= \inf_{\mathbf{x}, \mathbf{y}} \|\mathbf{x}\|^2 + \|M\mathbf{y}\|^2 + 2\mathbf{x}^*M\mathbf{y} \\ &\geq \inf_{\mathbf{x}, \mathbf{y}} \|\mathbf{x}\|^2 + \gamma_s^2(M)\|\mathbf{y}\|^2 - 2\|\mathbf{x}\|\|\mathbf{y}\|\phi_s(M) \\ &\geq \inf_{\mathbf{x}, \mathbf{y}} \min\{1, \gamma_s^2(M)\}(\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2) - \phi_s(M) \\ &= \min\{1, \gamma_s^2(M)\} - \phi_s(M), \end{aligned}$$

where we have used that for $t \in [0, 1]$, $\sqrt{t(1-t)} \leq 1/2$. \blacksquare

TWO CONCENTRATION RESULTS

Lemma 10: Suppose $\boldsymbol{\mu} \in \mathbb{R}^m$, $\|\boldsymbol{\mu}\|_\infty \leq C_\mu m^{-1/2}$, $\|\boldsymbol{\mu}\| = 1$, and let $J \subset [m]$ be uniformly distributed amongst all subsets of size ρm . Then for any $\varepsilon > 0$, with probability at least $1 - \exp(-\varepsilon m)$ in J ,

$$\|\boldsymbol{\mu}_{J^c}\|^2 > 1 - \rho - 8C_\mu^2\sqrt{2\varepsilon}. \quad (64)$$

In particular, if $\delta \geq 1$ then with probability at least $1 - \exp(-2\varepsilon_{\delta, \rho, \nu}^* m)$,

$$\|\boldsymbol{\mu}_{J^c}\|^2 > (1 - \rho)/4, \quad (65)$$

where $\varepsilon_{\delta, \rho, \nu}^*$ is specified in (9).

Proof: Consider the group of permutations of $[m]$ with distance $d(\pi_1, \pi_2) = m^{-1} \#\{i \mid \pi_1(i) \neq \pi_2(i)\}$. Let π be a random permutation sampled according to the uniform measure. Then random variable $\|\boldsymbol{\mu}_{J^c}\|^2$ is equal in distribution to $f(\pi) = \sum_{i=1}^{(1-\rho)m} \boldsymbol{\mu}(\pi(i))^2$. Since $\|\boldsymbol{\mu}\|_\infty < C_\mu m^{-1/2}$,

$$\begin{aligned} |f(\pi_1) - f(\pi_2)| &\leq 4C_\mu^2 m^{-1} \#\{\pi_1(i) \neq \pi_2(i)\} \\ &= 4C_\mu^2 d(\pi_1, \pi_2). \end{aligned}$$

Corollary 4.3 of [3] shows that

$$\mathbb{P}[f(\pi) \leq \mathbb{E}f(\pi) - t] \leq \exp\left(-\frac{mt^2}{27C_\mu^4}\right).$$

Noticing that $\mathbb{E}f = 1 - \rho$ establishes (64).

For (65), notice that by (9),

$$\begin{aligned} \sqrt{\varepsilon^*} &\leq \nu\sqrt{1 - \rho} \leq 2^{-12}\sqrt{C_s(1 - \rho)/\delta} \\ &\leq 2^{-12}\sqrt{2^{-17}C_\mu^{-4}\delta^{-2}(1 - \rho)^3} \\ &\leq 2^{-20}C_\mu^{-2}(1 - \rho), \end{aligned}$$

where in the last step we have used that $\delta \geq 1$. Hence, $8C_\mu^2\sqrt{2} \times 2\varepsilon^* \leq 2^{-16}(1 - \rho) \ll 3(1 - \rho)/4$, and by (64) with probability at least $1 - \exp(-2\varepsilon^* m)$, $\|\boldsymbol{\mu}_{J^c}\|^2 \geq (1 - \rho)/4$. \blacksquare

Lemma 11: Fix any $\delta \geq 1, \rho < 1$. Suppose that $\nu \leq \nu^*$, and $\varepsilon^*, m^*, \alpha^*$ are fixed according to (9), and that $m \geq m^*$. With probability at least $1 - 3\exp(-2\varepsilon^* m)$, the matrix W defined by

$$W^* = \left[\begin{array}{c|c} \frac{\boldsymbol{\mu}_{J^c}}{\sqrt{1-\rho}} & \frac{Z_{J^c, I}}{\sqrt{1-\rho}} \end{array} \right] \in \mathbb{R}^{(m-k_2) \times (k_1+1)}$$

satisfies $\sigma_{\min}(W) \geq 1/4$.

Proof: For compactness, let

$$\boldsymbol{\vartheta} \doteq \boldsymbol{\mu}_{J^c} / \sqrt{1-\rho}$$

and let $P_{\boldsymbol{\vartheta}}$ be the matrix that projects onto the span of $\boldsymbol{\vartheta}$. Then

$$W = \begin{bmatrix} \boldsymbol{\vartheta} & (\mathbf{I} - P_{\boldsymbol{\vartheta}}) \frac{Z_{J^c, I}}{\sqrt{1-\rho}} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & P_{\boldsymbol{\vartheta}} \frac{Z_{J^c, I}}{\sqrt{1-\rho}} \end{bmatrix},$$

and so

$$\sigma_{\min}(W) \geq \sigma_{k_1+1} \left[\boldsymbol{\vartheta} \quad \frac{(\mathbf{I} - P_{\boldsymbol{\vartheta}}) Z_{J^c, I}}{\sqrt{1-\rho}} \right] - \left\| \frac{P_{\boldsymbol{\vartheta}} Z_{J^c, I}}{\sqrt{1-\rho}} \right\|.$$

For any matrix $M = [A \ B] \in \mathbb{R}^{p \times (q_1+q_2)}$ with $A \in \mathbb{R}^{p \times q_1}$ and $B \in \mathbb{R}^{p \times q_2}$ and $A^*B = 0$, $\sigma_{q_1+q_2}(M) = \min\{\sigma_{q_1}(A), \sigma_{q_2}(B)\}$, and so

$$\begin{aligned} \sigma_{\min}(W) &\geq \min \left\{ \|\boldsymbol{\vartheta}\|, \frac{\sigma_{k_1}[(\mathbf{I} - P_{\boldsymbol{\vartheta}}) Z_{J^c, I}]}{\sqrt{1-\rho}} \right\} - \left\| \frac{P_{\boldsymbol{\vartheta}} Z_{J^c, I}}{\sqrt{1-\rho}} \right\| \\ &\geq \min \left\{ \|\boldsymbol{\vartheta}\|, \frac{\sigma_{k_1}[Z_{J^c, I}]}{\sqrt{1-\rho}} \right\} - 2 \left\| \frac{P_{\boldsymbol{\vartheta}} Z_{J^c, I}}{\sqrt{1-\rho}} \right\|. \end{aligned}$$

Further define the events

$$\begin{aligned} \mathcal{E}_1 &= \left\{ \|\boldsymbol{\vartheta}\| \leq 1/2 \right\}, \\ \mathcal{E}_2 &= \left\{ \|\mathbb{P}_{\boldsymbol{\vartheta}} Z_{J^c, I}\| \geq \sqrt{\alpha^*} + \sqrt{\frac{1}{m}} + \sqrt{4\varepsilon^*} \right\}, \\ \mathcal{E}_3 &= \left\{ \sigma_{k_1}[Z_{J^c, I}] \leq \sqrt{1-\rho} - \sqrt{\alpha^*} - \sqrt{4\varepsilon^*} \right\}. \end{aligned}$$

Notice that by Lemma 10, $\mathbb{P}[\mathcal{E}_1] \leq \exp(-2\varepsilon^*m)$. Similarly, by (22), noticing that $\mathbb{P}_{\boldsymbol{\vartheta}}$ projects $Z_{J^c, I}$ onto an independently distributed subspace of dimension at most one, $\mathbb{P}[\mathcal{E}_2] < \exp(-2\varepsilon^*m)$. Finally, by (23), $\mathbb{P}[\mathcal{E}_3] < \exp(-2\varepsilon^*m)$. So, $P[\cap_{i=1}^3 \mathcal{E}_i^c] > 1 - 3 \exp(-2\varepsilon^*m)$.

Now, from (9), notice that

$$\sqrt{\varepsilon^*/(1-\rho)} \leq \nu\tau < \nu < 2^{-12},$$

and so $\sqrt{4\varepsilon^*/(1-\rho)} < 2^{-11}$. Similarly, since

$$\alpha^* < \Upsilon[\varepsilon^*/\delta] \leq \varepsilon^*/\delta \leq \varepsilon^*,$$

$\sqrt{\alpha^*/(1-\rho)} < 2^{-12}$. Finally, by (9), for $m \geq m^*$,

$$\sqrt{\frac{1}{m(1-\rho)}} \leq \sqrt{2^{-17}} < 2^{-8}.$$

So, on $\cap_i \mathcal{E}_i$,

$$\begin{aligned} \sigma_{\min}(W) &\geq \min(1/2, 1 - 2^{-12} - 2^{-11}) \\ &\quad - 2 \times (2^{-12} + 2^{-8} + 2^{-11}) \\ &> 1/4, \end{aligned}$$

the desired result. \blacksquare

PROOF OF LEMMA 5

Proof: The lemma is an application of a concentration result on the cube due to [39]. We use the version stated in Corollary 4.10 and Equation (4.10) of [3]. The result states that if $\boldsymbol{\sigma} \in \mathbb{R}^n$ is distributed according to any product measure on $[-1, 1]^n$, $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex, $\|f\|_{\text{lip}}$ -Lipschitz function, and m_f is any median of f , then

$$\mathbb{P}[|f(\boldsymbol{\sigma}) - m_f| \geq t] \leq 4 \exp(-t^2/16\|f\|_{\text{lip}}^2). \quad (66)$$

To show (45), notice that

$$\begin{aligned} \mathbb{E}[\|M\boldsymbol{\sigma}\|^2] &= \mathbb{E}\left[\sum_i \left(\sum_j M_{ij}\sigma_j\right)^2\right] \\ &= \sum_i \mathbb{E}\left[\sum_{j,k} M_{ij}M_{ik}\sigma_j\sigma_k\right] = \sum_i \sum_j M_{ij}^2 = \|M\|_F^2. \end{aligned}$$

So, $\mathbb{E}\|M\boldsymbol{\sigma}\| \leq \sqrt{\mathbb{E}[\|M\boldsymbol{\sigma}\|^2]} \leq \|M\|_F$. Since $f(\cdot) \doteq \|M \cdot\|$ is nonnegative, by the Markov inequality, f has a median no larger than $2\mathbb{E}f \leq 2\|M\|_F$. Since f is $\|M\|$ -Lipschitz, plugging into (66) completes the proof of (45).

For (46), let $g: \mathbb{R} \rightarrow \mathbb{R}$ with $g(t) = |\mathcal{S}(t)| = \max(t - 1/2, -t - 1/2, 0)$. As the pointwise maximum of three linear functions, g is convex. Let $h: \mathbb{R}^n \rightarrow \mathbb{R}$ as $h(x_1, \dots, x_n) = \|(g(x_1), \dots, g(x_n))\|$. Notice that g is nonnegative, and further that for nonnegative \mathbf{z} , $\|\mathbf{z}\|$ is nondecreasing in each z_i . Hence, $h(\alpha\mathbf{x} + (1-\alpha)\mathbf{y}) = \|(g(\alpha x_1 + (1-\alpha)y_1), \dots, g(\alpha x_n + (1-\alpha)y_n))\| \leq \|\alpha(g(x_1), \dots, g(x_n)) + (1-\alpha)(g(y_1), \dots, g(y_n))\| \leq \alpha h(\mathbf{x}) + (1-\alpha)h(\mathbf{y})$. Thus h is a convex function. $f(\boldsymbol{\sigma}) \doteq \|\mathcal{S}[M\boldsymbol{\sigma}]\| = h(M\boldsymbol{\sigma})$ is the composition of a convex function with a linear function and hence is also convex.

Moreover, since for any scalars a, b , $|\mathcal{S}(a) - \mathcal{S}(b)| \leq |a - b|$, for vectors \mathbf{x}, \mathbf{y} , $\|\mathcal{S}(\mathbf{x}) - \mathcal{S}(\mathbf{y})\| \leq \|\mathbf{x} - \mathbf{y}\|$, and so g is 1-Lipschitz. As the composition of a 1-Lipschitz function with an $\|M\|$ -Lipschitz function, f is $\|M\|$ -Lipschitz. Hence, (66) applies, (again noticing that there is a median no larger than $2\mathbb{E}f$), giving

$$\mathbb{P}[\|\mathcal{S}[M\boldsymbol{\sigma}]\| > 2\mathbb{E}\|\mathcal{S}[M\boldsymbol{\sigma}]\| + 4\|M\|\sqrt{\varepsilon n}] < 4 \exp(-\varepsilon n). \quad (67)$$

It remains to estimate the expectation. Notice that $(M\boldsymbol{\sigma})_i = M_{ij}\sigma_j = \sum_j M_{ij}\sigma_j$ is a sum of independent, bounded random variables. Applying Hoeffding's inequality,

$$\mathbb{P}[(M\boldsymbol{\sigma})_i > t] < 2 \exp(-2t^2/\|M_i\|^2). \quad (68)$$

To simplify notation, let X be an arbitrary random variable satisfying a normal tail bound $\mathbb{P}[|X| > t] < C_1 \exp(-C_2 t^2)$. Then

$$\begin{aligned} \mathbb{E}[\mathcal{S}[X]^2] &= 2 \int_{t=0}^{\infty} t \mathbb{P}[|\mathcal{S}[X]| > t] dt \\ &= 2 \int_{t=1/2}^{\infty} (t - 1/2) \mathbb{P}[|X| > t] dt \\ &\leq C_1 \int_{t=1/2}^{\infty} 2t \exp(-C_2 t^2) dt = \frac{C_1}{C_2} \exp(-C_2/4). \end{aligned}$$

Hence,

$$\begin{aligned} \mathbb{E}[\|\mathcal{S}[M\boldsymbol{\sigma}]\|^2] &= \sum_i \mathbb{E}[\mathcal{S}[M_i\boldsymbol{\sigma}]^2] \\ &\leq \sum_i \|M_i\|^2 \exp(-1/2\|M_i\|^2) \\ &\leq \exp(-1/2 \max_i \|M_i\|^2) \sum_i \|M_i\|^2 \\ &= \exp(-1/2\|M\|_{2,\infty}^2) \|M\|_F^2. \end{aligned}$$

Bounding $\mathbb{E}[f]$ by $\sqrt{\mathbb{E}[f^2]}$ and plugging in to (67) completes the proof. \blacksquare

REDUCTION TO LARGE ERROR FRACTIONS

Lemma 12: Let $A \in \mathbb{R}^{m \times n}$, and suppose $(\mathbf{x}_0, \mathbf{e}_0)$ is the unique optimal solution to

$$\min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad A\mathbf{x} + \mathbf{e} = A\mathbf{x}_0 + \mathbf{e}_0. \quad (69)$$

Then if $\text{supp}(\mathbf{x}_0) \subseteq I$, then $(\mathbf{x}_0(I), \mathbf{e}_0)$ is the unique optimal solution to

$$\min \|\mathbf{x}'\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad A_{[m],I}\mathbf{x}' + \mathbf{e} = A_{[m],I}\mathbf{x}_0(I) + \mathbf{e}_0. \quad (70)$$

Proof: Any solution $(\mathbf{x}', \mathbf{e})$ to $A_{[m],I}\mathbf{x}' + \mathbf{e} = A_{[m],I}\mathbf{x}_0(I) + \mathbf{e}_0$ corresponds to a solution (\mathbf{x}, \mathbf{e}) to $A\mathbf{x} + \mathbf{e} = A_{[m],I}\mathbf{x}_0 + \mathbf{e}_0$, by setting $\mathbf{x}(I) = \mathbf{x}'$, $\mathbf{x}(I^c) = 0$. If there exists a solution to (70) with objective function $\leq \|\mathbf{x}_0(I)\|_1 + \|\mathbf{e}_0\|_1$, then either $(\mathbf{x}_0, \mathbf{e}_0)$ are not uniquely optimal for (69) or that solution is equal to $(\mathbf{x}_0(I), \mathbf{e}_0)$. ■

This lemma mimics [36]:

Lemma 13: Suppose that $(\mathbf{x}_0, \mathbf{e}_0)$ is the unique optimal solution to

$$\min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad A\mathbf{x} + \mathbf{e} = A\mathbf{x}_0 + \mathbf{e}_0. \quad (71)$$

Let $J = \text{supp}(\mathbf{e}_0)$, and suppose $J' \subset J$, and that \mathbf{e}'_0 is constructed by setting $\mathbf{e}'_0(J') = \mathbf{e}_0(J')$, $\mathbf{e}'_0(J'^c) = 0$. Then $(\mathbf{x}_0, \mathbf{e}'_0)$ is the unique optimal solution to

$$\min \|\mathbf{x}\|_1 + \|\mathbf{e}\|_1 \quad \text{subject to} \quad A\mathbf{x} + \mathbf{e} = A\mathbf{x}_0 + \mathbf{e}'_0. \quad (72)$$

Proof: Suppose on the contrary, that there exist $\tilde{\mathbf{x}}, \tilde{\mathbf{e}}$ satisfying $A\tilde{\mathbf{x}} + \tilde{\mathbf{e}} = A\mathbf{x}_0 + \mathbf{e}'_0$ and

$$\|\tilde{\mathbf{x}}\|_1 + \|\tilde{\mathbf{e}}\|_1 < \|\mathbf{x}_0\|_1 + \|\mathbf{e}'_0\|_1.$$

Set $\mathbf{h} = \mathbf{e}_0 - \mathbf{e}'_0$. Notice that $A\tilde{\mathbf{x}} + \tilde{\mathbf{e}} + \mathbf{h} = A\mathbf{x}_0 + \mathbf{e}_0$, and so $(\tilde{\mathbf{x}}, \tilde{\mathbf{e}} + \mathbf{h})$ is feasible for (71). But

$$\begin{aligned} \|\mathbf{x}_0\|_1 + \|\mathbf{e}_0\|_1 &= \|\mathbf{x}_0\|_1 + \|\mathbf{e}'_0\|_1 + \|\mathbf{h}\|_1 \\ &\geq \|\tilde{\mathbf{x}}\|_1 + \|\tilde{\mathbf{e}}\|_1 + \|\mathbf{h}\|_1 \geq \|\tilde{\mathbf{x}}\|_1 + \|\tilde{\mathbf{e}} + \mathbf{h}\|_1. \end{aligned}$$

Since the optimizer to (71) is unique, we conclude that $\tilde{\mathbf{x}} = \mathbf{x}_0$, $\tilde{\mathbf{e}} = \mathbf{e}'_0$. ■

REFERENCES

- [1] T. Figiel, J. Lindenstrauss, and V. D. Milman, "The dimension of almost spherical sections of convex bodies," *Acta Mathematica*, vol. 139, no. 1-2, pp. 53-94, 2008.
- [2] B. S. Kashin, "The widths of certain finite-dimensional sets and classes of smooth functions," *Izvestia*, vol. 41, no. 2, pp. 334-351, 2008.
- [3] M. Ledoux, *The Concentration of Measure Phenomenon*, *Mathematical Surveys and Monographs* 89. Providence, RI: American Mathematical Society, 2001.
- [4] R. Tibshirani, "Regression shrinkage and selection via the Lasso," *Journal of the Royal Statistical Society Series B*, vol. 58, no. 1, pp. 267-288, 1996.
- [5] W. Fu and K. Knight, "Asymptotics for Lasso-type estimators," *Annals of Statistics*, vol. 28, no. 5, pp. 1356-1378, 2000.
- [6] N. Meinshausen and B. Yu, "Lasso-type recovery of sparse representations for high-dimensional data," *Annals of Statistics*, vol. 37, no. 1, pp. 246-270, 2009.
- [7] R. Berinde, A. C. Gilbert, P. Indyk, H. Karloff, and M. J. Strauss, "Combining geometry and combinatorics: A unified approach to sparse signal recovery," in *Allerton Conference on Communication, Control and Computing*, 2008, pp. 798-805.
- [8] V. Guruswami, J. R. Lee, and A. Razborov, "Almost Euclidean subspaces of ℓ_1^n via expander codes," in *Symposium on Discrete Algorithms*, 2008, pp. 353-362.
- [9] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34-81, 2009.
- [10] W. Bajwa, J. Haupt, G. Raz, and R. Nowak, "Compressed channel sensing," in *Proceedings of Conference on Information Sciences and Systems*, 2008, pp. 5-10.
- [11] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Y. Ma, "Robust face recognition via sparse representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 210-227, 2009.
- [12] J. Yang, J. Wright, T. Huang, and Y. Ma, "Image super-resolution as sparse representation of raw image patches," in *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition*, 2008.
- [13] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Discriminative learned dictionaries for local image analysis," in *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition*, 2008.
- [14] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203-4215, 2005.
- [15] K. Lee, J. Ho, and D. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 5, pp. 684-698, 2005.
- [16] D. Donoho, "For most large underdetermined systems of linear equations the minimal l_1 -norm solution is also the sparsest solution," *Communications on Pure and Applied Mathematics*, vol. 59, no. 6, pp. 797-829, 2006.
- [17] R. Corless, G. Gonnet, D. Hare, D. Jeffrey, and D. Knuth, "On the Lambert W function," *Advances in Computational Mathematics*, vol. 5, no. 329, 1996.
- [18] M. Rudelson and R. Vershynin, "On sparse reconstruction from fourier and gaussian measurements," *Communications on Pure and Applied Mathematics*, vol. 61, pp. 1025-1045, 2008.
- [19] D. Donoho and J. Tanner, "Exponential bounds implying the construction of compressed sensing matrices, error-correcting codes and neighborly polytopes by random sampling," *IEEE Transactions on Information Theory*, 2010.
- [20] —, "Counting faces of randomly projected polytopes when the projection radically lowers dimension," *Journal of the American Mathematical Society*, vol. 22, no. 1, pp. 1-53, 2009.
- [21] J. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231-2242, 2004.
- [22] D. Needell and J. Tropp, "CoSAMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301-321, 2009.
- [23] Y. Eldar and M. Mishali, "Robust recovery of signals from a union of subspaces," *IEEE Transactions on Information Theory*, to be published.
- [24] M. Sipser and D. A. Spielman, "Expander codes," *IEEE Transactions on Information Theory*, vol. 42, no. 6, pp. 1710-1722, 1996.
- [25] J. Feldman, T. Malkin, R. Servedio, C. Stein, and M. Wainwright, "LP decoding corrects a constant fraction of errors," *IEEE Transactions on Information Theory*, vol. 53, no. 1, pp. 82-89, 2007.
- [26] M. R. Capalbo, O. Reingold, S. P. Vadhan, and A. Wigderson, "Randomness conductors and constant-degree lossless expanders," in *Proceedings of the 34th ACM Symposium on Theory of Computing*, 2002, pp. 659-668.
- [27] N. Kashyap, "A decomposition theorem for binary linear codes," *IEEE Transactions on Information Theory*, vol. 54, no. 7, pp. 3035-3058, 2008.
- [28] D. Donoho, "Neighborly polytopes and sparse solution of underdetermined linear equations," preprint, 2005. [Online]. Available: www-stat.stanford.edu/Reports/abstracts/05-04.pdf.
- [29] J. Fuchs, "On sparse representation in arbitrary redundant bases," *IEEE Transactions on Information Theory*, vol. 50, no. 6, pp. 1341-1344, 2004.
- [30] K. Davidson and S. Szarek, "Local operator theory, random matrices, and banach spaces," in *Handbook of the Geometry of Banach Spaces*. Amsterdam, Netherlands: Elsevier, 2001, pp. 317-366.
- [31] E. Candès and J. Romberg, " ℓ^1 -magic: Recovery of sparse signals via convex programming," 2005. [Online]. Available: <http://www.acm.caltech.edu/l1magic/>.
- [32] D. Needell and R. Vershynin, "Signal recovery from inaccurate and incomplete measurements via regularized orthogonalized matching pursuit," *Journal of Selected Topics in Signal Processing*, to be published.

- [33] D. Donoho, "For most large underdetermined systems of linear equations the minimal ℓ^1 -norm near solution approximates the sparsest solution," preprint, 2004. [Online]. Available: www-stat.stanford.edu/~donoho/Reports/2004/l1l10approx.pdf.
- [34] E. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n ," *Annals of Statistics*, vol. 35, no. 6, pp. 2313–2351, 2005.
- [35] ———, "The power of convex relaxation: Near-optimal matrix completion," *IEEE Transactions on Information Theory*, submitted for publication.
- [36] E. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *preprint*, 2009.
- [37] A. Ganesh, J. Wright, X. Li, E. Candès, and Y. Ma, "Dense error correction for low-rank matrices via principal component pursuit," in *International Symposium on Information Theory (ISIT)*, submitted for publication, 2010.
- [38] Z. Zhou, A. Wagner, H. Mobahi, and Y. Ma, "Face recognition with contiguous occlusion using markov random fields," in *Proceedings of International Conference on Computer Vision*, 2009.
- [39] M. Talagrand, "Concentration of measure and isoperimetric inequalities in product spaces," *Publications Mathématiques de l'I.H.E.S.*, vol. 81, pp. 73–205, 1995.

John Wright is a researcher with the Visual Computing Group at Microsoft Research Asia. He received his PhD in Electrical Engineering from the University of Illinois at Urbana-Champaign. His graduate work focused on developing efficient and provably correct algorithms for error correction with high-dimensional data, and on their application in automatic face recognition. His research interests encompass a number of topics in vision and signal processing, including minimum description length methods for clustering and classification, error correction and inference with non-ideal data, video analysis and tracking, as well as face and object recognition. His work has received a number of awards and honors, including a UIUC Distinguished Fellowship, Carver Fellowship, Microsoft Research Fellowship, the UIUC Martin Award for Outstanding Graduate Research, and the Lemelson-Illinois Prize for Innovation.

Yi Ma is an associate professor at the Electrical and Computer Engineering Department of the University of Illinois at Urbana-Champaign. He is also the research manager of the Visual Computing group at Microsoft Research Asia in Beijing since January 2009. His main research interest is in computer vision, high-dimensional data analysis, and systems theory. He is the first author of the popular vision textbook "An Invitation to 3-D Vision," published by Springer in 2003. Yi Ma received two Bachelors' degree in Automation and Applied Mathematics from Tsinghua University (Beijing, China) in 1995, a Master of Science degree in EECS in 1997, a Master of Arts degree in Mathematics in 2000, and a PhD degree in EECS in 2000, all from the University of California at Berkeley. Yi Ma received the David Marr Best Paper Prize at the International Conference on Computer Vision 1999, the Longuet-Higgins Best Paper Prize at the European Conference on Computer Vision 2004, and the Sang Uk Lee Best Student Paper Award with his students at the Asian Conference on Computer Vision in 2009. He also received the CAREER Award from the National Science Foundation in 2004 and the Young Investigator Award from the Office of Naval Research in 2005. He is an associate editor of IEEE Transactions on Pattern Analysis and Machine Intelligence and has served as the chief guest editor for special issues for the Proceedings of IEEE and the IEEE Signal Processing Magazine. He will also serve as Program Chair for ICCV 2013 in Sydney, Australia. He is a senior member of IEEE and a member of ACM, SIAM, and ASEE.