

Simultaneously Structured Models With Application to Sparse and Low-Rank Matrices

Samet Oymak, *Student Member, IEEE*, Amin Jalali, *Student Member, IEEE*, Maryam Fazel, Yonina C. Eldar, *Fellow, IEEE*, and Babak Hassibi, *Member, IEEE*

Abstract—Recovering structured models (e.g., sparse or group-sparse vectors, low-rank matrices) given a few linear observations have been well-studied recently. In various applications in signal processing and machine learning, the model of interest is structured in *several* ways, for example, a matrix that is simultaneously sparse and low rank. Often norms that promote the individual structures are known, and allow for recovery using an orderwise optimal number of measurements (e.g., ℓ_1 norm for sparsity, nuclear norm for matrix rank). Hence, it is reasonable to minimize a combination of such norms. We show that, surprisingly, using multiobjective optimization with these norms can do no better, orderwise, than exploiting only one of the structures, thus revealing a fundamental limitation in sample complexity. This result suggests that to fully exploit the multiple structures, we need an entirely new convex relaxation. Further, specializing our results to the case of sparse and low-rank matrices, we show that a nonconvex formulation recovers the model from very few measurements (on the order of the degrees of freedom), whereas the convex problem combining the ℓ_1 and nuclear norms requires many more measurements, illustrating a gap between the performance of the convex and nonconvex recovery problems. Our framework applies to arbitrary structure-inducing norms as well as to a wide range of measurement ensembles. This allows us to give sample complexity bounds for problems such as sparse phase retrieval and low-rank tensor completion.

Index Terms—Compressed sensing, convex relaxation, regularization, sample complexity.

Manuscript received February 11, 2013; revised July 21, 2014; accepted October 29, 2014. Date of publication February 9, 2015; date of current version April 17, 2015. A. Jalali and M. Fazel were supported by the National Science Foundation under Grant ECCS-0847077 and Grant CCF-1409836. Y. C. Eldar was supported in part by the Israel Science Foundation under Grant 170/10, in part by SRC, and in part by the Intel Collaborative Research Institute for Computational Intelligence. B. Hassibi was supported in part by the National Science Foundation under Grant CNS-0932428, Grant CCF-1018927, Grant CCF-1423663, and Grant CCF-1409204, in part by the Office of Naval Research through the MURI under Grant N00014-08-0747, in part by Qualcomm Inc., San Diego, CA, USA, and in part by King Abdulaziz University, Jeddah, Saudi Arabia.

S. Oymak is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: oymak@eecs.berkeley.edu).

A. Jalali and M. Fazel are with the Department of Electrical Engineering, University of Washington, Seattle, WA 98195 USA (e-mail: amjalali@uw.edu; mfazel@uw.edu).

Y. C. Eldar is with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel (e-mail: yonina@ee.technion.ac.il).

B. Hassibi is with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: hassibi@systems.caltech.edu).

Communicated by V. Saligrama, Associate Editor for Signal Processing.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2015.2401574

I. INTRODUCTION

RECOVERY of a structured model (signal) given a small number of linear observations has been the focus of many studies recently. Examples include recovering sparse or group-sparse vectors (which gave rise to the area of compressed sensing) [1]–[4], low-rank matrices [5], [6], and the sum of sparse and low-rank matrices [7], [8], among others. More generally, the recovery of a signal that can be expressed as the sum of a few atoms out of an appropriate atomic set has been studied in [9]. Canonical questions in this area include: How many linear measurements are enough to recover the model by any means? How many measurements are enough for a tractable approach? In the statistics literature, these questions are often posed in terms of error rates for estimators minimizing the sum of a quadratic loss function and a regularizer that reflects the desired structure [10].

There are many applications where the model of interest is known to have *several* structures at the same time (Section I-B). We then seek a signal that lies in the intersection of several sets defining the individual structures (in a sense that we will later make precise). The most common convex regularizer used to promote all structures together is a linear combination of well-known regularizers for each structure. However, there is currently no general analysis and understanding of how well such regularization performs in terms of the number of observations required for successful recovery of the desired model. This paper addresses this ubiquitous yet unexplored problem; i.e., the recovery of *simultaneously structured models*.

An example of a simultaneously structured model is a matrix that is *simultaneously sparse and low-rank*. One would like to come up with algorithms that exploit both types of structures to minimize the number of measurements required for recovery. An $n \times n$ matrix with rank $r \ll n$ can be described by $\mathcal{O}(rn)$ parameters, and can be recovered using $\mathcal{O}(rn)$ generic measurements via nuclear norm minimization [5], [11]. On the other hand, a block-sparse matrix with a $k \times k$ nonzero block where $k \ll n$ can be described by k^2 parameters and can be recovered given $\mathcal{O}(k^2 \log \frac{n}{k})$ generic measurements using ℓ_1 minimization. However, a matrix that is *both* rank r and block-sparse can be described by $\mathcal{O}(rk)$ parameters. The question is whether we can exploit this joint structure to efficiently recover such a matrix with $\mathcal{O}(rk)$ measurements.

In this paper we give a negative answer to this question in the following sense: if we use multi-objective optimization

TABLE I

SUMMARY OF RESULTS IN RECOVERY OF STRUCTURED SIGNALS. THIS PAPER SHOWS A GAP BETWEEN THE PERFORMANCE OF CONVEX AND NONCONVEX RECOVERY PROGRAMS FOR SIMULTANEOUSLY STRUCTURED MATRICES (LAST ROW)

Model	Degrees of Freedom	Nonconvex Recovery	Convex Recovery
Sparse vectors	k	$\mathcal{O}(k)$	$\mathcal{O}(k \log \frac{n}{k})$
Low rank matrices	$r(2n - r)$	$\mathcal{O}(rn)$	$\mathcal{O}(rn)$
Low rank plus sparse	$\mathcal{O}(rn + k)$	not analyzed	$\mathcal{O}((rn + k) \log^2 n)$
Low rank and sparse	$\mathcal{O}(r(k_1 + k_2))$	$\mathcal{O}(r(k_1 + k_2) \log n)$	$\Omega(rn)$

with the ℓ_1 and nuclear norms (used for sparse signals and low rank matrices, respectively), then the number of measurements required is lower bounded by $\mathcal{O}(\min\{k^2, rn\})$. In other words, we need at least this number of observations for the desired signal to be recoverable by a combination of the ℓ_1 norm and the nuclear norm. This means we can do *no better than an algorithm that exploits only one of the two structures*.

We introduce a framework to express general simultaneously structured models, and as our main result, we prove that the same phenomenon happens for a general set of structures. We analyze a wide range of measurement ensembles, including the subsampled standard basis (i.e. matrix completion), Gaussian and subgaussian measurements, and quadratic measurements. Table I summarizes known results on recovery of some common structured models, along with a result of this paper specialized to the problem of low-rank and sparse matrix recovery. The first column gives the number of parameters needed to describe the model (often referred to as its ‘degrees of freedom’), while the second and third columns show how many generic measurements are needed for successful recovery. In ‘nonconvex recovery’, we assume we are able to find the global minimum of a nonconvex problem. This is clearly intractable in general, and not a practical recovery method—we consider it as a benchmark for theoretical comparison with the (tractable) convex relaxation in order to determine how powerful the relaxation is.

The first and second rows are the results on k sparse vectors in \mathbb{R}^n and rank r matrices in $\mathbb{R}^{n \times n}$ respectively, [11], [12]. The third row considers the recovery of ‘‘low-rank plus sparse’’ matrices. Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ that can be decomposed as $\mathbf{X} = \mathbf{X}_L + \mathbf{X}_S$ where \mathbf{X}_L is a rank r matrix and \mathbf{X}_S is a matrix with only k nonzero entries. The degrees of freedom in \mathbf{X} are $\mathcal{O}(rn + k)$. Minimizing the combination of the ℓ_1 norm and nuclear norm, i.e., $f(\mathbf{X}) = \min_{\mathbf{Y}} \|\mathbf{Y}\|_* + \lambda \|\mathbf{X} - \mathbf{Y}\|_1$ subject to random Gaussian measurements on \mathbf{X} , gives a convex approach for recovering \mathbf{X} . It has been shown that under reasonable incoherence assumptions, \mathbf{X} can be recovered from $\mathcal{O}((rn + k) \log^2 n)$ measurements which is suboptimal only by a logarithmic factor [13]. Finally, the last row in Table I shows one of the results in this paper. Let $\mathbf{X} \in \mathbb{R}^{n \times n}$ be a rank r matrix whose entries are zero outside a $k_1 \times k_2$ submatrix. The degrees of freedom of \mathbf{X} are $\mathcal{O}((k_1 + k_2)r)$. We consider both convex and non-convex programs for the recovery of this type of matrix. The nonconvex method involves minimizing the number of nonzero rows, columns and rank of the matrix jointly, as discussed in Section III-B. As shown later, $\mathcal{O}((k_1 + k_2)r \log n)$ measurements suffice for this program to successfully recover

the original matrix. The convex method minimizes any convex combination of the individual structure-inducing norms, namely the nuclear norm and the $\ell_{1,2}$ norm of the matrix, which encourage low-rank and column/row-sparse solutions respectively. We show that with high probability this program cannot recover the original matrix with fewer than $\Omega(rn)$ measurements. In summary, while the nonconvex method is only slightly suboptimal, the convex method performs poorly as the number of measurements scales with n rather than $k_1 + k_2$.

A. Contributions

This paper presents an analysis for the recovery of models with more than one structure, by combining penalties corresponding to each structure. The setting considered is very broad: any number of structures can be combined, any (convex) combination of the corresponding norms can be analyzed, and the results hold for a variety of popular measurement ensembles. The framework proposed includes special cases that are of interest in their own right, e.g., sparse and low-rank matrix recovery, and low-rank tensor completion [14], [15].

More specifically, our contributions can be summarized as follows.

1) *Poor Performance of Convex Relaxations*: We consider a model with several structures and associated structure-inducing norms. For recovery, we consider a multi-objective optimization problem to minimize the individual norms simultaneously. Given the convexity of the problem, we know that minimizing a weighted sum of the norms and varying the weights traces out all points of the Pareto-optimal front (Section II). We obtain a lower bound on the number of measurements for any convex function combining the individual norms. A sketch of our main result is as follows.

Given a model \mathbf{x}_0 with τ simultaneous structures, the number of measurements required for recovery with high probability using any linear combination of the individual norms satisfies the lower bound

$$m \geq c \min_{i=1, \dots, \tau} m_i$$

where m_i is an intrinsic lower bound on the required number of measurements when minimizing the i th norm only. The term c depends on the measurement ensemble.

For the norms of interest, m_i is proportional to the degrees of freedom of the i th structure. With $\min_i m_i$ as the bottleneck, this result indicates that the combination of norms performs

no better than using only one of the norms, even though the target model has small degrees of freedom.

2) *Different Measurement Ensembles*: Our characterization of recovery failure is easy to interpret and deterministic in nature. We show that it can be used to obtain probabilistic failure results for various random measurement ensembles. In particular, our results hold for measurement matrices with i.i.d subgaussian rows, quadratic measurements and matrix completion type measurements.

3) *Understanding the Effect of Weighting*: We characterize the sample complexity of the multi-objective function as a function of the weights associated with the individual norms. Our upper and lower bounds reveal that the sample complexity of the multi-objective function is related to a certain convex combination of the sample complexities associated with the individual norms. We provide formulas for this combination as a function of the weights.

4) *Incorporating General Cone Constraints*: In addition, we can incorporate side information on \mathbf{x}_0 , expressed as convex cone constraints. This additional information helps in recovery; however, quantifying how much the cone constraints help is not trivial. Our analysis explicitly determines the role of the cone constraint: Geometric properties of the cone such as its Gaussian width determines the constant factors in the bound on the number of measurements.

5) *Illustrating a Gap for the Recovery of Sparse and Low-Rank Matrices*: As a special case, we consider the recovery of simultaneously sparse and low-rank matrices and prove that there is a significant gap between the performance of convex and non-convex recovery programs. This gap is surprising when one considers similar results on low-dimensional model recovery as shown in Table I.

B. Applications

We survey several applications where simultaneous structures arise, as well as existing results specific to these applications.

1) *Sparse Signal Recovery From Quadratic Measurements*: Sparsity has long been exploited in signal processing, applied mathematics, statistics and computer science for tasks such as compression, denoising, model selection, image processing and more. Despite the great interest in exploiting sparsity in various applications, most of the work to date has focused on recovering sparse or low rank data from linear measurements. Recently, the basic sparse recovery problem has been generalized to the case in which the measurements are given by nonlinear transforms of the unknown input, [16]. A special case of this more general setting is quadratic compressed sensing [17] in which the goal is to recover a sparse vector \mathbf{x} from quadratic measurements $b_i = \mathbf{x}^T \mathbf{A}_i \mathbf{x}$. This problem can be linearized by *lifting*, where we wish to recover a “low rank and sparse” matrix $\mathbf{X} = \mathbf{x}\mathbf{x}^T$ subject to measurements $b_i = \langle \mathbf{A}_i, \mathbf{X} \rangle$.

Sparse recovery problems from quadratic measurements arise in a variety of problems in optics. One example is sub-wavelength optical imaging [17], [18] in which the goal is to recover a sparse image from its far-field measurements,

where due to the laws of physics the relationship between the (clean) measurement and the unknown image is quadratic. In [17] the quadratic relationship is a result of using partially-incoherent light. The quadratic behavior of the measurements in [18] arises from coherent diffractive imaging in which the image is recovered from its intensity pattern. Under an appropriate experimental setup, this problem amounts to reconstruction of a sparse signal from the magnitude of its Fourier transform.

A related and notable problem involving sparse and low-rank matrices is Sparse Principal Component Analysis (SPCA), mentioned in Section IX.

2) *Sparse Phase Retrieval*: Quadratic measurements appear in phase retrieval problems, in which a signal is to be recovered from the magnitude of its measurements $b_i = |\mathbf{a}_i^T \mathbf{x}|$, where each measurement is a linear transform of the input $\mathbf{x} \in \mathbb{R}^n$ and \mathbf{a}_i 's are arbitrary, possibly complex-valued measurement vectors. An important case is when $\mathbf{a}_i^T \mathbf{x}$ is the Fourier transform and b_i^2 is the power spectral density. Phase retrieval is of great interest in many applications such as optical imaging [19], [20], crystallography [21], and more [22]–[24].

The problem becomes linear when \mathbf{x} is *lifted* and we consider the recovery of $\mathbf{X} = \mathbf{x}\mathbf{x}^T$ where each measurement takes the form $b_i^2 = \langle \mathbf{a}_i \mathbf{a}_i^T, \mathbf{X} \rangle$. In [17], an algorithm was developed to treat phase retrieval problems with sparse \mathbf{x} based on a semidefinite relaxation, and low-rank matrix recovery combined with a row-sparsity constraint on the resulting matrix. More recent works also proposed the use of semidefinite relaxation together with sparsity constraints for phase retrieval [25]–[28]. An alternative algorithm was recently designed in [29] based on a greedy search. In [27], the authors consider sparse signal recovery based on combinatorial and probabilistic approaches and provide uniqueness results under certain conditions. Stable uniqueness in phase retrieval problems is studied in [30]. The results of [31] and [32] applies to general (non-sparse) signals where in some cases *masked* versions of the signal are required.

3) *Fused Lasso*: Suppose the signal of interest \mathbf{x}_0 is sparse and its entries vary slowly, i.e., the signal can be approximated by a piecewise constant function. To encourage sparsity, one can use the ℓ_1 norm. To promote the piece-wise constant structure, discrete total variation can be used, defined as

$$\|\mathbf{x}\|_{TV} = \sum_{i=1}^{n-1} |\mathbf{x}_{i+1} - \mathbf{x}_i|,$$

where $\|\cdot\|_{TV}$ is basically the ℓ_1 norm of the gradient and is approximately sparse. The resulting optimization problem that estimates \mathbf{x}_0 from its samples $\mathbf{A}\mathbf{x}_0$ is known as fused-lasso [33], and is given as

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 + \lambda \|\mathbf{x}\|_{TV} \quad \text{s.t. } \mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}_0. \quad (1)$$

To the best of our knowledge, the sample complexity of fused lasso has not been analyzed from a compressed sensing point of view. However, there is a series of recent works [34], [35] on total variation minimization; which may lead to analysis of (1).

We remark that TV regularization is also used together with nuclear norm to encourage a low-rank and smooth (i.e., slowly varying entries) solution. This regularization finds applications in imaging and physics [36], [37].

4) *Low-Rank Tensors*: Tensors with small Tucker rank can be seen as a generalization of low-rank matrices [38]. In this setup, the signal of interest is a tensor $\mathbf{X}_0 \in \mathbb{R}^{n_1 \times \dots \times n_\tau}$, where \mathbf{X}_0 is low-rank along its unfoldings which are obtained by reshaping \mathbf{X}_0 as a matrix with size $n_i \times \frac{n}{n_i}$, with $n = \prod_{i=1}^{\tau} n_i$. Denoting the i th unfolding by $\mathcal{U}_i(\mathbf{X}_0)$, a standard approach to estimate \mathbf{X}_0 from $\mathbf{y} = \mathcal{A}(\mathbf{X}_0)$ is minimizing the weighted nuclear norms of the unfoldings,

$$\min_{\mathbf{X}} \sum_{i=1}^{\tau} \lambda_i \|\mathcal{U}_i(\mathbf{X})\|_* \quad \text{subject to} \quad \mathbf{y} = \mathcal{A}(\mathbf{X}_0). \quad (2)$$

Low-rank tensors have applications in machine learning, physics, computational finance and high dimensional PDE's [15]. The problem (2) has been investigated in several papers [14], [39]. Closer to us, [40] recently showed that the convex relaxation (2) performs poorly compared to information theoretically optimal bounds for Gaussian measurements. Our results can extend those to the more applicable tensor completion setup, where we observe the entries of the tensor.

Other applications of simultaneously structured signals include Collaborative Hierarchical Sparse Modeling [41] where sparsity is considered within the non-zero blocks in a block-sparse vector, and the recovery of hyperspectral images where we aim to recover a simultaneously block sparse and low rank matrix from compressed observations [42].

C. Outline of the Paper

The paper is structured as follows. Background and definitions are given in Section II. An overview of the main results is provided in Section III. Section IV discusses some measurement ensembles for which our results apply. Section V derives upper bounds for the convex relaxations assuming a Gaussian measurement ensemble. Proofs of the general results are presented in Section VI. The proofs for the special case of simultaneously sparse and low-rank matrices are given in Section VII, where we compare corollaries of the general results with the results on non-convex recovery approaches, and illustrate a gap. Numerical simulations in Section VIII empirically support the results on sparse and low-rank matrices. Future directions of research and a discussion on the results are presented in Section IX.

II. PROBLEM SETUP

We begin by reviewing some basic definitions. Our results will be on structure-inducing norms; examples include the ℓ_1 norm, the $\ell_{1,2}$ norm, and the nuclear norm.

A. Definitions and Signal Model

The nuclear norm of a matrix is denoted by $\|\cdot\|_*$ and is the sum of the singular values of the matrix. The $\ell_{1,2}$ norm is the sum of the ℓ_2 norms of the columns of a matrix.




Fig. 1. Depiction of the correlation between a vector \mathbf{x} and a set S . \mathbf{s}^* achieves the largest angle with \mathbf{x} , hence \mathbf{s}^* has the minimum correlation with \mathbf{x} .

Minimizing the ℓ_1 norm encourages sparse solutions, while the $\ell_{1,2}$ norm and nuclear norm encourage column-sparse and low-rank solutions respectively, [5], [6], [43]–[45]; see Section VI-D for more detailed discussion of these norms and their subdifferentials. The Euclidean norm is denoted by $\|\cdot\|_2$, i.e., the ℓ_2 norm for vectors and the Frobenius norm $\|\cdot\|_F$ for matrices.

Overlines denote normalization, i.e., for a vector \mathbf{x} and a matrix \mathbf{X} , $\bar{\mathbf{x}} = \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$ and $\bar{\mathbf{X}} = \frac{\mathbf{X}}{\|\mathbf{X}\|_F}$. The minimum and maximum singular values of a matrix \mathbf{A} are denoted by $\sigma_{\min}(\mathbf{A})$ and $\sigma_{\max}(\mathbf{A})$. The set of $n \times n$ positive semidefinite (PSD) and symmetric matrices are denoted by \mathbb{S}_+^n and \mathbb{S}^n respectively. $\text{cone}(S)$ denotes the conic hull of a given set S . $\mathcal{A}(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear measurement operator if $\mathcal{A}(\mathbf{x})$ is equivalent to the matrix multiplication $\mathbf{A}\mathbf{x}$ where $\mathbf{A} \in \mathbb{R}^{m \times n}$. If \mathbf{x} is a matrix, $\mathcal{A}(\mathbf{x})$ will be a matrix multiplication with a suitably vectorized \mathbf{x} . In some of our results, we consider Gaussian measurements, in which case \mathbf{A} has independent $\mathcal{N}(0, 1)$ entries.

For a vector $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|$ denotes a general norm and $\|\mathbf{x}\|^* = \sup_{\|\mathbf{z}\| \leq 1} \langle \mathbf{x}, \mathbf{z} \rangle$ is the corresponding dual norm. A subgradient of the norm $\|\cdot\|$ at \mathbf{x} is a vector \mathbf{g} for which $\|\mathbf{z}\| \geq \|\mathbf{x}\| + \langle \mathbf{g}, \mathbf{z} - \mathbf{x} \rangle$ holds for any \mathbf{z} . The set of all subgradients is called the subdifferential and is denoted by $\partial\|\mathbf{x}\|$. The Lipschitz constant of the norm is defined as

$$L = \sup_{\mathbf{z}_1 \neq \mathbf{z}_2 \in \mathbb{R}^n} \frac{\|\mathbf{z}_1\| - \|\mathbf{z}_2\|}{\|\mathbf{z}_1 - \mathbf{z}_2\|_2}.$$

Definition 1 (Correlation): Given a nonzero vector \mathbf{x} and a set S , $\rho(\mathbf{x}, S)$ is defined as

$$\rho(\mathbf{x}, S) := \inf_{\mathbf{s} \in S} \frac{|\mathbf{x}^T \mathbf{s}|}{\|\mathbf{x}\|_2 \|\mathbf{s}\|_2}.$$

$\rho(\mathbf{x}, S)$ corresponds to the minimum absolute-valued correlation between the vector \mathbf{x} and elements of S . Let $\bar{\mathbf{x}} = \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$. The correlation between \mathbf{x} and the associated subdifferential has a simple form:

$$\rho(\mathbf{x}, \partial\|\mathbf{x}\|) = \inf_{\mathbf{g} \in \partial\|\mathbf{x}\|} \frac{\bar{\mathbf{x}}^T \mathbf{g}}{\|\mathbf{g}\|_2} = \frac{\|\bar{\mathbf{x}}\|}{\sup_{\mathbf{g} \in \partial\|\mathbf{x}\|} \|\mathbf{g}\|_2}.$$

Here, we used the fact that, for norms, subgradients $\mathbf{g} \in \partial\|\mathbf{x}\|$ satisfy $\mathbf{x}^T \mathbf{g} = \|\mathbf{x}\|$, [46]. The denominator of the right hand side is the local Lipschitz constant of $\|\cdot\|$ at \mathbf{x} and is upper bounded by L . Consequently, $\rho(\mathbf{x}, \partial\|\mathbf{x}\|) \geq \frac{\|\bar{\mathbf{x}}\|}{L}$. We will denote $\frac{\|\bar{\mathbf{x}}\|}{L}$ by κ . Recently, this quantity has been studied




Fig. 2. For a scaled norm ball passing through \mathbf{x}_0 , $\kappa = \frac{\|\mathbf{p}\|_2}{\|\mathbf{x}_0\|_2}$, where \mathbf{p} is any of the closest points on the scaled norm ball to the origin.

by Mu et al. to analyze the simultaneously structured signals in a similar spirit to us for Gaussian measurements [40].¹ Similar calculations as above gives an alternative interpretation for κ which is illustrated in Figure 2.

κ is a measure of alignment between the vector \mathbf{x} and the subdifferential. For the norms of interest, it is associated with the model complexity. For instance, for a k -sparse vector \mathbf{x} , $\|\bar{\mathbf{x}}\|_1$ lies between 1 and \sqrt{k} depending on how spiky nonzero entries are. Also the Lipschitz constant for the ℓ_1 norm in \mathbb{R}^n is $L = \sqrt{n}$. Hence, when the nonzero entries are ± 1 , we find $\kappa^2 = \frac{k}{n}$. Similarly, given a $d \times d$, rank r matrix \mathbf{X} , $\|\bar{\mathbf{X}}\|_*$ lies between 1 and \sqrt{r} . If the singular values are spread (i.e. ± 1), we find $\kappa^2 = \frac{r}{d} = \frac{rd}{d^2}$. In these cases, κ^2 is proportional to the model complexity normalized by the ambient dimension.

Simultaneously Structured Models: We consider a signal \mathbf{x}_0 which has several low-dimensional structures S_1, S_2, \dots, S_τ (e.g., sparsity, group sparsity, low-rank). Suppose each structure i corresponds to a norm denoted by $\|\cdot\|_{(i)}$ which promotes that structure (e.g., $\ell_1, \ell_{1,2}$, nuclear norm). We refer to such an \mathbf{x}_0 as a *simultaneously structured model*.

B. Convex Recovery Program

We investigate the recovery of the simultaneously structured \mathbf{x}_0 from its linear measurements $\mathcal{A}(\mathbf{x}_0)$. To recover \mathbf{x}_0 , we would like to simultaneously minimize the norms $\|\cdot\|_{(i)}$, $i = 1, \dots, \tau$, which leads to a multi-objective (vector-valued) optimization problem. For all feasible points \mathbf{x} satisfying $\mathcal{A}(\mathbf{x}) = \mathcal{A}(\mathbf{x}_0)$ and side information $\mathbf{x} \in \mathcal{C}$, consider the set of achievable norms $\{\|\mathbf{x}\|_{(i)}\}_{i=1}^\tau$ denoted as points in \mathbb{R}^τ . The minimal points of this set with respect to the positive orthant \mathbb{R}_+^τ form the *Pareto-optimal* front, as illustrated in Figure 3. Since the problem is convex, one can alternatively consider the set

$$\{\mathbf{v} \in \mathbb{R}^\tau : \exists \mathbf{x} \in \mathbb{R}^n \text{ s.t. } \mathbf{x} \in \mathcal{C}, \mathcal{A}(\mathbf{x}) = \mathcal{A}(\mathbf{x}_0), \\ v_i \geq \|\mathbf{x}\|_{(i)}, \text{ for } i = 1, \dots, \tau\},$$

which is convex and has the same Pareto optimal points as the original set (see [47, Ch. 4]).

¹The work [40] was submitted after our initial manuscript; in which we projected the subdifferential onto a carefully chosen subspace to obtain bounds on the sample complexity (see Proposition 6). Inspired by [40], projection onto \mathbf{x}_0 and the use of κ led to the simplification of the notation and improvement of the results in the current manuscript, in particular, Section IV.




Fig. 3. Consider a point \mathbf{x}_0 represented in the figure by the dot. We need at least m measurements for \mathbf{x}_0 to be recoverable since for any $\underline{m} < m$ this point is not on the Pareto optimal front.

Definition 2 (Recoverability): We call \mathbf{x}_0 recoverable if it is a Pareto optimal point; i.e., there does not exist a feasible $\mathbf{x}' \neq \mathbf{x}$ satisfying $\mathcal{A}(\mathbf{x}') = \mathcal{A}(\mathbf{x}_0)$ and $\mathbf{x}' \in \mathcal{C}$, with $\|\mathbf{x}'\|_{(i)} \leq \|\mathbf{x}_0\|_{(i)}$ for $i = 1, \dots, \tau$.

The vector-valued convex recovery program can be turned into a scalar optimization problem as

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{C}}{\text{minimize}} && f(\mathbf{x}) = h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)}) \\ & \text{subject to} && \mathcal{A}(\mathbf{x}) = \mathcal{A}(\mathbf{x}_0), \end{aligned} \quad (3)$$

where $h : \mathbb{R}_+^\tau \rightarrow \mathbb{R}_+$ is convex and non-decreasing in each argument (i.e., non-decreasing and strictly increasing in at least one coordinate). For convex problems with strong duality, it is known that we can recover all of the Pareto optimal points by optimizing weighted sums $f(\mathbf{x}) = \sum_{i=1}^\tau \lambda_i \|\mathbf{x}\|_{(i)}$, with positive weights λ_i , among all possible functions $f(\mathbf{x}) = h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)})$. For each \mathbf{x}_0 on the Pareto, the coefficients of such a recovering function are given by the hyperplane supporting the Pareto at \mathbf{x}_0 [47, Ch. 4].

In Figure 3, consider the smallest m that makes \mathbf{x}_0 recoverable. Then one can choose a function h and recover \mathbf{x}_0 by (3) using the m measurements. If the number of measurements is any less, then *no* function can recover \mathbf{x}_0 . Our goal is to provide lower bounds on m .

In [9], Chandrasekaran et al. propose a general theory for constructing a suitable penalty, called an *atomic norm*, given a single set of atoms that describes the structure of the target object. In the case of simultaneous structures, this construction requires defining new atoms, and then ensuring the resulting atomic norm can be minimized in a computationally tractable way, which is nontrivial and often intractable. We briefly discuss such constructions as a future research direction in Section IX.

III. MAIN RESULTS

In this section, we state our main theorems that aim to characterize the number of measurements needed to recover a simultaneously structured signal by convex or nonconvex programs. We first present our general results, followed by results for simultaneously sparse and low-rank matrices as a specific but important instance of the general case. The proofs

are given in Sections VI and VII. All of our statements will implicitly assume $\mathbf{x}_0 \neq 0$. This will ensure that \mathbf{x}_0 is not a trivial minimizer and 0 is not in the subdifferentials.

A. General Simultaneously Structured Signals

Consider the recovery of a signal \mathbf{x}_0 that is simultaneously structured with S_1, S_2, \dots, S_τ as described in Section II-A. We provide a lower bound on the required number of measurements, using the geometric properties of the individual norms.

Theorem 1 (Deterministic Failure): Suppose $\mathcal{C} = \mathbb{R}^n$ and,

$$\rho(\mathbf{x}_0, \partial f(\mathbf{x}_0)) := \inf_{\mathbf{g} \in \partial f(\mathbf{x}_0)} |\bar{\mathbf{g}}^T \bar{\mathbf{x}}_0| > \frac{\|\mathbf{A} \bar{\mathbf{x}}_0\|_2}{\sigma_{\min}(\mathbf{A}^T)}. \quad (4)$$

Then, \mathbf{x}_0 is not a minimizer of (3).

Theorem 1 is deterministic in nature. However, it can be easily specialized to specific random measurement ensembles. The left hand side of (4) depends only on the vector \mathbf{x}_0 and the subdifferential $\partial f(\mathbf{x}_0)$, hence it is independent of the measurement matrix \mathbf{A} . For simultaneously structured models, we will argue that, the left hand side cannot be made too small, as the subgradients are *aligned* with the signal. On the other hand, the right hand side depends only on \mathbf{A} and \mathbf{x}_0 and is independent of the subdifferential. In linear inverse problems, \mathbf{A} is often assumed to be random. For large class of random matrices, we will argue that, the right hand side is approximately $\sim \sqrt{\frac{m}{n}}$ which will yield a lower bound on the number of required measurements.

Typical measurement ensembles include the following:

- **Sampling Entries:** In low-rank matrix and tensor completion problems, we observe the entries of \mathbf{x}_0 uniformly at random. In this case, rows of \mathbf{A} are chosen from the standard basis in \mathbb{R}^n . We should remark that, instead of the standard basis, one can consider other orthonormal bases such as the Fourier basis.
- **Matrices With i.i.d. Rows:** \mathbf{A} has independent and identically distributed rows with certain moment conditions. This is a widely used setup in compressed sensing as each measurement we make is associated with the corresponding row of \mathbf{A} [48].
- **Quadratic Measurements:** Arises in the phase retrieval problem as discussed in Section I-B.

In Section IV, we find upper bounds on the right hand side of (4) for these ensembles. As discussed in Section IV, we can modify the rows of \mathbf{A} to get better bounds as long as this does not affect its null space. For instance, one can discard the identical rows to improve conditioning. However, as m increases and \mathbf{A} has more linearly independent rows, $\sigma_{\min}(\mathbf{A}^T)$ will naturally decrease and (4) will no longer hold after a certain point. In particular, (4) cannot hold beyond $m \geq n$ as $\sigma_{\min}(\mathbf{A}^T) = 0$. This is indeed natural as the system becomes overdetermined.

The following proposition lower bounds the left hand side of (4) in an interpretable manner. In particular, the correlation $\rho(\mathbf{x}_0, \partial f(\mathbf{x}_0))$ can be lower bounded by the smallest individual correlation.

Proposition 1: Let L_i be the Lipschitz constant of the i th norm and $\kappa_i = \frac{\|\bar{\mathbf{x}}_0\|_{(i)}}{L_i}$ for $1 \leq i \leq \tau$. Set $\kappa_{\min} = \min\{\kappa_i : i = 1, \dots, \tau\}$. Then:

- All functions $f(\cdot)$ in (3) satisfy, $\rho(\mathbf{x}_0, \partial f(\mathbf{x}_0)) \geq \kappa_{\min}$.²
- Suppose $f(\cdot)$ is a weighted linear combination $f(\mathbf{x}) = \sum_{i=1}^{\tau} \lambda_i \|\mathbf{x}\|_{(i)}$ for nonnegative $\{\lambda_i\}_{i=1}^{\tau}$. Let $\bar{\lambda}_i = \frac{\lambda_i L_i}{\sum_{i=1}^{\tau} \lambda_i L_i}$ for $1 \leq i \leq \tau$. Then, $\rho(\mathbf{x}_0, \partial f(\mathbf{x}_0)) \geq \sum_{i=1}^{\tau} \bar{\lambda}_i \kappa_i$.

Proof: From Lemma 3, any subgradient of $f(\cdot)$ can be written as, $\mathbf{g} = \sum_{i=1}^{\tau} w_i \mathbf{g}_i$ for some nonnegative w_i 's. On the other hand, from [46], $\langle \bar{\mathbf{x}}_0, \mathbf{g}_i \rangle = \|\bar{\mathbf{x}}_0\|_{(i)}$. Combining these results,

$$\mathbf{g}^T \bar{\mathbf{x}}_0 = \sum_{i=1}^{\tau} w_i \|\bar{\mathbf{x}}_0\|_{(i)}.$$

From the triangle inequality, $\|\mathbf{g}\|_2 \leq \sum_{i=1}^{\tau} w_i L_i$. Therefore,

$$\frac{\sum_{i=1}^{\tau} w_i \|\bar{\mathbf{x}}_0\|_{(i)}}{\sum_{i=1}^{\tau} w_i L_i} \geq \min_{1 \leq i \leq \tau} \frac{w_i \|\bar{\mathbf{x}}_0\|_{(i)}}{w_i L_i} = \kappa_{\min}. \quad (5)$$

To prove the second part, we use the fact that for the weighted sums of norms, $w_i = \lambda_i$ and the subgradients have the form $\mathbf{g} = \sum_{i=1}^{\tau} \lambda_i \mathbf{g}_i$, [47]. Then, substitute $\bar{\lambda}_i$ for λ_i in the left hand side of (5). ■

Before stating the next result, let us give a relevant definition regarding the average distance between a set and a random vector.

Definition 3 (Gaussian Distance): Let \mathcal{M} be a closed convex set in \mathbb{R}^n and let $\mathbf{h} \in \mathbb{R}^n$ be a vector with independent standard normal entries. Then, the Gaussian distance of \mathcal{M} is defined as

$$\mathbf{D}(\mathcal{M}) = \mathbb{E}[\inf_{\mathbf{v} \in \mathcal{M}} \|\mathbf{h} - \mathbf{v}\|_2]$$

When \mathcal{M} is a cone, we have $0 \leq \mathbf{D}(\mathcal{M}) \leq \sqrt{n}$. Similar definitions have been used extensively in the literature, such as Gaussian width [9], statistical dimension [49] and mean width [50]. For notational simplicity, let the normalized distance be $\bar{\mathbf{D}}(\mathcal{M}) = \frac{\mathbf{D}(\mathcal{M})}{\sqrt{n}}$.

We will now state our result for Gaussian measurements; which can additionally include cone constraints for the lower bound. Other ensembles are considered in Section IV.

Theorem 2 (Gaussian Lower Bound): Suppose \mathbf{A} has independent $\mathcal{N}(0, 1)$ entries. Whenever $m \leq m_{low}$, \mathbf{x}_0 will not be a minimizer of any of the recovery programs in (3) with probability at least $1 - 10 \exp(-\frac{1}{16} \min\{m_{low}, (1 - \bar{\mathbf{D}}(\mathcal{C}))^2 n\})$, where

$$m_{low} \triangleq \frac{(1 - \bar{\mathbf{D}}(\mathcal{C})) n \kappa_{\min}^2}{100}.$$

Remark: When $\mathcal{C} = \mathbb{R}^n$, $\bar{\mathbf{D}}(\mathcal{C}) = 0$ hence, the lower bound simplifies to $m_{low} = \frac{n \kappa_{\min}^2}{100}$.

Note that $\bar{\mathbf{D}}(\mathcal{C})$ depends only on \mathcal{C} and can be viewed as a constant. For instance, for the positive semidefinite cone, we show that $\bar{\mathbf{D}}(\mathbb{S}_+^n) < \frac{\sqrt{3}}{2}$. Observe that for a smaller cone \mathcal{C} , it is reasonable to expect a smaller lower bound on the

²The lower bound κ_{\min} is directly comparable to [40, Th. 5]. Indeed, our lower bounds on the sample complexity will have the form $\mathcal{O}(\kappa_{\min}^2 n)$.

TABLE II

SUMMARY OF THE PARAMETERS THAT ARE DISCUSSED IN THIS SECTION. THE LAST THREE LINES ARE FOR A $d \times d$ S&L (k, k, r) MATRIX WHERE $n = d^2$. IN THE FOURTH COLUMN, THE CORRESPONDING ENTRY FOR S&L IS $\kappa_{\min} = \min\{\kappa_{\ell_1}, \kappa_{\star}\}$

Model	$f(\cdot)$	L	$\ \bar{\mathbf{x}}_0\ \leq$	$n\kappa^2 \leq$
k sparse vector	$\ \cdot\ _1$	\sqrt{n}	\sqrt{k}	k
k column-sparse matrix	$\ \cdot\ _{1,2}$	\sqrt{d}	\sqrt{k}	kd
Rank r matrix	$\ \cdot\ _{\star}$	\sqrt{d}	\sqrt{r}	rd
S&L (k, k, r) matrix	$h(\ \cdot\ _{\star}, \ \cdot\ _1)$	–	–	$\min\{k^2, rd\}$

required number of measurements. Indeed, as \mathcal{C} gets smaller, $\bar{\mathbf{D}}(\mathcal{C})$ increases.

As discussed before, there are various options for the scalarizing function in (3), with one choice being the weighted sum of norms. In fact, for a recoverable point \mathbf{x}_0 there always exists a weighted sum of norms which recovers it. This function is also often the choice in applications, where the space of positive weights is searched for a good combination. Thus, we can state the following corollary as a general result.

Corollary 1 (Weighted Lower Bound): Suppose \mathbf{A} has i.i.d $\mathcal{N}(0, 1)$ entries and $f(\mathbf{x}) = \sum_{i=1}^{\tau} \lambda_i \|\mathbf{x}\|_{(i)}$ for nonnegative weights $\{\lambda_i\}_{i=1}^{\tau}$. Whenever $m \leq m'_{low}$, \mathbf{x}_0 will not be a minimizer of the recovery program (3) with probability at least $1 - 10 \exp(-\frac{1}{16} \min\{m'_{low}, (1 - \bar{\mathbf{D}}(\mathcal{C}))^2 n\})$, where

$$m'_{low} \triangleq \frac{n(1 - \bar{\mathbf{D}}(\mathcal{C}))(\sum_{i=1}^{\tau} \bar{\lambda}_i \kappa_i)^2}{100},$$

$$\text{and } \bar{\lambda}_i = \frac{\lambda_i L_i}{\sum_{i=1}^{\tau} \lambda_i L_i}.$$

Observe that Theorem 2 is stronger than stating “a particular function $h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)})$ will not work”. Instead, our result states that with high probability none of the programs in the class (3) can return \mathbf{x}_0 as optimal unless the number of measurements is sufficiently large.

To understand the result better, note that the required number of measurements is proportional to $\kappa_{\min}^2 n$ which is often proportional to the sample complexity of the best individual norm. As we have argued in Section II-A, $\kappa_i^2 n$ corresponds to how structured the signal is. For sparse signals it is equal to the sparsity, and for a rank r matrix, it is equal to the degrees of freedom of the set of rank r matrices. Consequently, Theorem 2 suggests that even if the signal satisfies multiple structures, the required number of measurements is effectively determined by only one dominant structure.

Intuitively, the degrees of freedom of a simultaneously structured signal can be much lower, which is provable for simultaneously sparse and low-rank (S&L) matrices. Hence, there is a considerable gap between the expected measurements based on model complexity and the number of measurements needed for recovery via (3) ($\kappa_{\min}^2 n$).

B. Simultaneously Sparse and Low-Rank Matrices

We now focus on a special case, namely simultaneously sparse and low-rank (S&L) matrices. We consider matrices with nonzero entries contained in a small submatrix where the submatrix itself is low rank. Here, norms of interest are $\|\cdot\|_{1,2}$, $\|\cdot\|_1$ and $\|\cdot\|_{\star}$ and the cone of interest is the PSD cone. We also consider nonconvex approaches and contrast the

results with convex approaches. For the nonconvex problem, we replace the norms $\|\cdot\|_1, \|\cdot\|_{1,2}, \|\cdot\|_{\star}$ with the functions $\|\cdot\|_0, \|\cdot\|_{0,2}, \text{rank}(\cdot)$ which give the number of nonzero entries, the number of nonzero columns and rank of a matrix respectively and use the same cone constraint as the convex method. We show that convex methods perform poorly as predicted by the general result in Theorem 2, while nonconvex methods require optimal number of measurements (up to a logarithmic factor). Proofs are given in Section VII.

Definition 4: We say $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times d_2}$ is an S&L matrix with (k_1, k_2, r) if the smallest submatrix that contains nonzero entries of \mathbf{X}_0 has size $k_1 \times k_2$ and $\text{rank}(\mathbf{X}_0) = r$. When \mathbf{X}_0 is symmetric, let $d = d_1 = d_2$ and $k = k_1 = k_2$. We consider the following cases.

(a) *General:* $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times d_2}$ is S&L with (k_1, k_2, r) .

(b) *PSD model:* $\mathbf{X}_0 \in \mathbb{R}^{n \times n}$ is PSD and S&L with (k, k, r) .

We are interested in S&L matrices with $k_1 \ll d_1, k_2 \ll d_2$ so that the matrix is sparse, and $r \ll \min\{k_1, k_2\}$ so that the submatrix containing the nonzero entries is low rank. Recall from Section II-B that our goal is to recover \mathbf{X}_0 from linear observations $\mathcal{A}(\mathbf{X}_0)$ via convex or nonconvex optimization programs. The measurements can be equivalently written as $\mathbf{A} \text{vec}(\mathbf{X}_0)$, where $\mathbf{A} \in \mathbb{R}^{m \times d_1 d_2}$ and $\text{vec}(\mathbf{X}_0) \in \mathbb{R}^{d_1 d_2}$ denotes the vector obtained by stacking the columns of \mathbf{X}_0 .

Based on the results in Section III-A, we obtain lower bounds on the number of measurements for convex recovery. We additionally show that significantly fewer measurements are sufficient for non-convex programs to uniquely recover \mathbf{X}_0 ; thus proving a performance gap between convex and nonconvex approaches. The following theorem summarizes the results.

Theorem 3 (Performance of S&L Matrix Recovery): Suppose $\mathcal{A}(\cdot)$ is an i.i.d Gaussian map and consider recovering $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times d_2}$ via

$$\underset{\mathbf{X} \in \mathcal{C}}{\text{minimize}} f(\mathbf{X}) \quad \text{subject to } \mathcal{A}(\mathbf{X}) = \mathcal{A}(\mathbf{X}_0). \quad (6)$$

For the cases given in Definition 4, the following convex and nonconvex recovery results hold for some positive constants c_1, c_2 .

(a) *General Model:*

(a1) Let $f(\mathbf{X}) = \|\mathbf{X}\|_{1,2} + \lambda_1 \|\mathbf{X}^T\|_{1,2} + \lambda_2 \|\mathbf{X}\|_{\star}$ where $\lambda_1, \lambda_2 \geq 0$ and $\mathcal{C} = \mathbb{R}^{d_1 \times d_2}$.

Then, (6) will fail to recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m_0)$ whenever $m \leq c_2 m_0$ where $m_0 = \min\{d_1 k_2, d_2 k_1, (d_1 + d_2)r\}$.

(a2) Let $f(\mathbf{X}) = \frac{1}{k_2} \|\mathbf{X}\|_{0,2} + \frac{1}{k_1} \|\mathbf{X}^T\|_{0,2} + \frac{1}{r} \text{rank}(\mathbf{X})$ and $\mathcal{C} = \mathbb{R}^{d_1 \times d_2}$. Then, (6) will uniquely

TABLE III

SUMMARY OF RECOVERY RESULTS FOR MODELS IN DEFINITION 4, ASSUMING $d_1 = d_2 = d$ AND $k_1 = k_2 = k$. FOR THE PSD WITH ℓ_1 CASE, WE ASSUME $\frac{\|\bar{\mathbf{X}}_0\|_1}{k}$ AND $\frac{\|\bar{\mathbf{X}}_0\|_*}{\sqrt{r}}$ TO BE APPROXIMATELY CONSTANT FOR THE SAKE OF SIMPLICITY. NONCONVEX APPROACHES ARE OPTIMAL UP TO A LOGARITHMIC FACTOR, WHILE CONVEX APPROACHES PERFORM POORLY

Setting	Nonconvex Sufficient m	Convex Required m
General model	$\mathcal{O}\left(\max\{rk, k \log \frac{d}{k}\}\right)$	$\Omega(rd)$
PSD with $\ell_{1,2}$	$\mathcal{O}\left(\max\{rk, k \log \frac{d}{k}\}\right)$	$\Omega(rd)$
PSD with ℓ_1	$\mathcal{O}\left(k \log \frac{d}{k}\right)$	$\Omega(\min\{k^2, rd\})$

recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m)$ whenever $m \geq c_2 \max\{(k_1 + k_2)r, k_1 \log \frac{d_1}{k_1}, k_2 \log \frac{d_2}{k_2}\}$.

(b) PSD with $\ell_{1,2}$:

- (b1) Let $f(\mathbf{X}) = \|\mathbf{X}\|_{1,2} + \lambda \|\mathbf{X}\|_*$ where $\lambda \geq 0$ and $\mathcal{C} = \mathbb{S}_+^d$. Then, (6) will fail to recover \mathbf{X}_0 with probability $1 - \exp(-c_1 rd)$ whenever $m \leq c_2 rd$.
- (b2) Let $f(\mathbf{X}) = \frac{2}{k} \|\mathbf{X}\|_{0,2} + \frac{1}{r} \text{rank}(\mathbf{X})$ and $\mathcal{C} = \mathbb{S}^d$. Then, (6) will uniquely recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m)$ whenever $m \geq c_2 \max\{rk, k \log \frac{d}{k}\}$.

(c) PSD with ℓ_1 :

- (c1) Let $f(\mathbf{X}) = \|\mathbf{X}\|_1 + \lambda \|\mathbf{X}\|_*$ and $\mathcal{C} = \mathbb{S}_+^d$. Then, (6) will fail to recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m_0)$ for all possible $\lambda \geq 0$ whenever $m \leq c_2 m_0$ where $m_0 = \min\{\|\bar{\mathbf{X}}_0\|_1^2, \|\bar{\mathbf{X}}_0\|_*^2 d\}$.
- (c2) Suppose $\text{rank}(\mathbf{X}_0) = 1$. Let $f(\mathbf{X}) = \frac{1}{k^2} \|\mathbf{X}\|_0 + \text{rank}(\mathbf{X})$ and $\mathcal{C} = \mathbb{S}^d$. Then, (6) will uniquely recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m)$ whenever $m \geq c_2 k \log \frac{d}{k}$.

Remark on ‘‘PSD With ℓ_1 ’’: In the special case, $\mathbf{X}_0 = \mathbf{a}\mathbf{a}^T$ for a k -sparse vector \mathbf{a} , we have $m_0 = \min\{\|\bar{\mathbf{a}}\|_1^4, d\}$. When nonzero entries of \mathbf{a} are ± 1 , we have $m_0 = \min\{k^2, d\}$.

The nonconvex programs require almost the same number of measurements as the degrees of freedom (or number of parameters) of the underlying model. For instance, it is known that the degrees of freedom of a rank r matrix of size $k_1 \times k_2$ is simply $r(k_1 + k_2 - r)$ which is $\mathcal{O}((k_1 + k_2)r)$. Hence, the nonconvex results are optimal up to a logarithmic factor. On the other hand, our results on the convex programs that follow from Theorem 2 show that the required number of measurements are significantly larger. Table III provides a quick comparison of the results on S&L.

For the S&L (k,k,r) model, from standard results one can easily deduce that [3], [5], [43],

- ℓ_1 penalty only: requires at least k^2 measurements,
- $\ell_{1,2}$ penalty only: requires at least kd measurements,
- Nuclear norm penalty only: requires at least rd measurements.

These follow from the model complexity of the sparse, column-sparse and low-rank matrices. Theorem 2 shows that, combination of norms require at least as much as the best individual norm. For instance, combination of ℓ_1 and the nuclear

norm penalization yields the lower bound $\mathcal{O}(\min\{k^2, rd\})$ for S&L matrices whose singular values and nonzero entries are spread. This is indeed what we would expect from the interpretation that $\kappa^2 n$ is often proportional to the sample complexity of the corresponding norm and, the lower bound $\kappa_{\min}^2 n$ is proportional to that of the best individual norm.

As we saw in Section III-A, adding a cone constraint to the recovery program does not help in reducing the lower bound by more than a constant factor. In particular, we discuss the positive semidefiniteness assumption that is beneficial in the sparse phase retrieval problem and show that the number of measurements remain high even when we include this extra information. On the other hand, the nonconvex recovery programs performs well even without the PSD constraint.

We remark that, we could have stated Theorem 3 for more general measurements given in Section IV without the cone constraint. For instance, the following result holds for the weighted linear combination of individual norms and for the subgaussian ensemble.

Corollary 2: Suppose $\mathbf{X}_0 \in \mathbb{R}^{d \times d}$ obeys the general model with $k_1 = k_2 = k$ and \mathcal{A} is a linear subgaussian map as described in Proposition 2. Choose $f(\mathbf{X}) = \lambda_{\ell_1} \|\mathbf{X}\|_1 + \lambda_* \|\mathbf{X}\|_*$, where $\lambda_{\ell_1} = \beta$, $\lambda_* = (1 - \beta)\sqrt{d}$ and $0 \leq \beta \leq 1$. Then, whenever, $m \leq \min\{m_{low}, c_1 n\}$, where,

$$m_{low} = \frac{(\beta \|\bar{\mathbf{X}}_0\|_1 + (1 - \beta) \|\bar{\mathbf{X}}_0\|_* \sqrt{d})^2}{2},$$

(6) fails with probability $1 - 4 \exp(-c_2 m_{low})$. Here $c_1, c_2 > 0$ are constants as described in Proposition 2.

Remark: Choosing $\mathbf{X}_0 = \mathbf{a}\mathbf{a}^T$ where nonzero entries of \mathbf{a} are ± 1 yields $\frac{1}{2}(\beta k + (1 - \beta)\sqrt{d})^2$ on the right hand side. An explicit construction of an S&L matrix with maximal $\|\bar{\mathbf{X}}\|_1, \|\bar{\mathbf{X}}\|_*$ is provided in Section VII-C.

This corollary compares well with the upper bound obtained in Corollary 5 of Section V. In particular, both the bounds and the penalty parameters match up to logarithmic factors. Hence, together, they sandwich the sample complexity of the combined cost $f(\mathbf{X})$.

IV. MEASUREMENT ENSEMBLES

This section will make use of standard results on sub-gaussian random variables and random matrix theory to obtain probabilistic statements. We will explain how one can analyze the right hand side of (4) for,

- Matrices with sub-gaussian rows,
- Subsampled standard basis (in matrix completion),
- Quadratic measurements arising in phase retrieval.

A. Sub-Gaussian Measurements

We first consider the measurement maps with sub-gaussian entries. The following definitions are borrowed from [51].

Definition 5 (Sub-Gaussian Random Variable): A random variable x is sub-gaussian if there exists a constant $K > 0$ such that for all $p \geq 1$,

$$(\mathbb{E} |x|^p)^{1/p} \leq K \sqrt{p}.$$

The smallest such K is called the sub-gaussian norm of x and is denoted by $\|x\|_{\Psi_2}$. A sub-exponential random

variable y is one for which there exists a constant K' such that, $\mathbb{P}(|y| > t) \leq \exp(1 - \frac{t}{K'})$. The variable x is sub-gaussian if and only if x^2 is sub-exponential.

Definition 6 (Isotropic Sub-Gaussian Vector): A random vector $\mathbf{x} \in \mathbb{R}^n$ is sub-gaussian if the one dimensional marginals $\mathbf{x}^T \mathbf{v}$ are sub-gaussian random variables for all $\mathbf{v} \in \mathbb{R}^n$. The sub-gaussian norm of \mathbf{x} is defined as,

$$\|\mathbf{x}\|_{\Psi_2} = \sup_{\|\mathbf{v}\|=1} \|\mathbf{x}^T \mathbf{v}\|_{\Psi_2}.$$

The vector \mathbf{x} is also isotropic if its covariance is equal to the identity, i.e. $\mathbb{E} \mathbf{x} \mathbf{x}^T = \mathbf{I}_n$.

Proposition 2 (Sub-Gaussian Measurements): Suppose \mathbf{A} has i.i.d rows in either of the following forms,

- a copy of a zero-mean isotropic sub-gaussian vector $\mathbf{a} \in \mathbb{R}^n$, where $\|\mathbf{a}\|_2 = \sqrt{n}$ almost surely.
- the rows consist of i.i.d zero-mean and unit-variance sub-gaussian entries.

Then, there exists constants c_1, c_2 depending only on the sub-gaussian norm of the rows, such that, whenever $m \leq c_1 n$, with probability $1 - 4 \exp(-c_2 m)$, we have,

$$\frac{\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2}{\sigma_{\min}^2(\mathbf{A}^T)} \leq \frac{2m}{n}.$$

Proof: Using [51, Th. 5.58], there exists constants c, C depending only on the sub-gaussian norm of \mathbf{a} such that for any $t \geq 0$, with probability $1 - 2 \exp(-ct^2)$

$$\sigma_{\min}(\mathbf{A}^T) \geq \sqrt{n} - C\sqrt{m} - t.$$

Choosing $t = C\sqrt{m}$ and $m \leq \frac{n}{100C^2}$ ensures that $\sigma_{\min}(\mathbf{A}^T) \geq \frac{4\sqrt{n}}{5}$.

Next, we shall estimate $\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2$, which is a sum of i.i.d. sub-exponential random variables identical to $|\mathbf{a}^T \bar{\mathbf{x}}_0|^2$. Note that $\mathbb{E}[|\mathbf{a}^T \bar{\mathbf{x}}_0|^2] = 1$. Hence, [51, Proposition 5.16] gives,

$$\mathbb{P}(\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2 \geq m + t) \leq 2 \exp(-c' \min\{\frac{t^2}{m}, t\}).$$

Choosing $t = \frac{7m}{25}$, we find that $\mathbb{P}(\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2 \geq \frac{32m}{25}) \leq 2 \exp(-c''m)$. Combining the two, we obtain,

$$\mathbb{P}\left(\frac{\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2}{\sigma_{\min}^2(\mathbf{A}^T)} \leq \frac{2m}{n}\right) \geq 1 - 4 \exp(-c'''m).$$

The second statement can be proved in the exact same manner by using [51, Th. 5.39] instead of Theorem 5.58. ■

Remark: While Proposition 2 assumes \mathbf{a} has fixed ℓ_2 norm, this can be ensured by properly normalizing rows of \mathbf{A} (assuming they stay sub-gaussian). For instance, if the ℓ_2 norm of the rows are larger than $c\sqrt{n}$ for a positive constant c , normalization will not affect sub-gaussianity. Note that, scaling rows of a matrix does not change its null space.

B. Randomly Sampling Entries

We now consider the scenario where each row of \mathbf{A} is chosen from the standard basis uniformly at random. Note that, when m is comparable to n , there is a nonnegligible probability that \mathbf{A} will have duplicate rows. Theorem 1 does not take this situation into account which would make $\sigma_{\min}(\mathbf{A}^T) = 0$. In this case, one can discard the copies as they

don't affect the recoverability of \mathbf{x}_0 . This would get rid of the ill-conditioning, since the new matrix is well-conditioned with the exact same null space as the original. This corresponds to a ‘‘sampling without replacement’’ scheme where we ensure each row is different.

Similar to achievability results in matrix completion [6], the following failure result requires true signal to be incoherent with the standard basis, where incoherence is characterized by $\|\bar{\mathbf{x}}_0\|_{\infty}$, which lies between $\frac{1}{\sqrt{n}}$ and 1.

Proposition 3 (Sampling Entries): Let $\{\mathbf{e}_i\}_{i=1}^n$ be the standard basis in \mathbb{R}^n and suppose each row of \mathbf{A} is chosen from $\{\mathbf{e}_i\}_{i=1}^n$ uniformly at random. Let $\hat{\mathbf{A}}$ be the matrix obtained by removing duplicate rows in \mathbf{A} . Then, with probability $1 - \exp(-\frac{m}{4n\|\bar{\mathbf{x}}_0\|_{\infty}^2})$, we have,

$$\frac{\|\hat{\mathbf{A}} \bar{\mathbf{x}}_0\|_2^2}{\sigma_{\min}^2(\hat{\mathbf{A}})} \leq \frac{2m}{n}.$$

Proof: Let $\hat{\mathbf{A}}$ be the matrix obtained by discarding the rows of \mathbf{A} that occur multiple times except one of them. Clearly $\text{Null}(\hat{\mathbf{A}}) = \text{Null}(\mathbf{A})$ hence they are equivalent for the purpose of recovering \mathbf{x}_0 . Furthermore, $\sigma_{\min}(\hat{\mathbf{A}}) = 1$. Therefore, we are interested in upper bounding $\|\hat{\mathbf{A}} \bar{\mathbf{x}}_0\|_2$.

Clearly $\|\hat{\mathbf{A}} \bar{\mathbf{x}}_0\|_2 \leq \|\mathbf{A} \bar{\mathbf{x}}_0\|_2$. Hence, we will bound $\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2$ probabilistically. Let \mathbf{a} be the first row of \mathbf{A} . $|\mathbf{a}^T \bar{\mathbf{x}}_0|^2$ is a random variable, with mean $\frac{1}{n}$ and is upper bounded by $\|\bar{\mathbf{x}}_0\|_{\infty}^2$. Applying the Chernoff Bound yields

$$\mathbb{P}(\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2 \geq \frac{m}{n}(1 + \delta)) \leq \exp(-\frac{m\delta^2}{2(1 + \delta)n\|\bar{\mathbf{x}}_0\|_{\infty}^2}).$$

Setting $\delta = 1$, we find that, with probability $1 - \exp(-\frac{m}{4n\|\bar{\mathbf{x}}_0\|_{\infty}^2})$, we have,

$$\frac{\|\hat{\mathbf{A}} \bar{\mathbf{x}}_0\|_2^2}{\sigma_{\min}(\hat{\mathbf{A}})^2} \leq \frac{\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2}{\sigma_{\min}(\hat{\mathbf{A}})^2} \leq \frac{2m}{n},$$

completing the proof. ■

A significant application of this result would be for the low-rank tensor completion problem, where we randomly observe some entries of a low-rank tensor and try to reconstruct it. A promising approach for this problem is using the weighted linear combinations of nuclear norms of the unfoldings of the tensor to induce the low-rank tensor structure described in (2), [14], [15]. Related work [40] shows the poor performance of (2) for the special case of Gaussian measurements. Combination of Theorem 1 and Proposition 3 will immediately extend the results of [40] to the more applicable tensor completion setup (under proper incoherence conditions that bound $\|\bar{\mathbf{x}}_0\|_{\infty}$).

Remark: In Propositions 2 and 3, we can make the upper bound for the ratio $\frac{\|\mathbf{A} \bar{\mathbf{x}}_0\|_2^2}{\sigma_{\min}(\mathbf{A})^2}$ arbitrarily close to $\frac{m}{n}$ by changing the proof parameters. Combined with Proposition 1, this would suggest that, failure happens, when $m < n\kappa_{\min}$.

C. Quadratic Measurements

As mentioned in the phase retrieval problem, quadratic measurements $|\mathbf{v}^T \mathbf{a}|^2$ of the vector $\mathbf{a} \in \mathbb{R}^d$ can be linearized by the change of variable $\mathbf{a} \rightarrow \mathbf{X}_0 = \mathbf{a} \mathbf{a}^T$ and using $\mathbf{V} = \mathbf{v} \mathbf{v}^T$.

The following proposition can be used to obtain a lower bound for such ensembles when combined with Theorem 1.

Proposition 4: Suppose we observe quadratic measurements $\mathcal{A}(\mathbf{X}_0) \in \mathbb{R}^m$ of a matrix $\mathbf{X}_0 = \mathbf{a}\mathbf{a}^T \in \mathbb{R}^{d \times d}$. Here, assume that i th entry of $\mathcal{A}(\mathbf{X}_0)$ is equal to $|\mathbf{v}_i^T \mathbf{a}|^2$ where $\{\mathbf{v}_i\}_{i=1}^m$ are independent vectors, either with $\mathcal{N}(0, 1)$ entries or are uniformly distributed over the sphere with radius \sqrt{d} . Then, there exists absolute constants $c_1, c_2 > 0$ such that whenever $m < \frac{c_1 d}{\log d}$, with probability $1 - 2ed^{-2}$,

$$\frac{\|\mathcal{A}(\bar{\mathbf{X}}_0)\|_2}{\sigma_{\min}(\mathbf{A}^T)} \leq \frac{c_2 \sqrt{m} \log d}{d}.$$

Proof: Let $\mathbf{V}_i = \mathbf{v}_i \mathbf{v}_i^T$. Without loss of generality, assume \mathbf{v}_i 's are uniformly distributed over the sphere with radius \sqrt{d} . To lower bound $\sigma_{\min}(\mathbf{A}^T)$, we estimate the coherence of its columns, defined by,

$$\mu(\mathbf{A}^T) = \max_{i \neq j} \frac{|\langle \mathbf{V}_i, \mathbf{V}_j \rangle|}{\|\mathbf{V}_i\|_F \|\mathbf{V}_j\|_F} = \frac{(\mathbf{v}_i^T \mathbf{v}_j)^2}{d^2}.$$

[51, Sec. 5.2.5] states that the sub-gaussian norm of \mathbf{v}_i is bounded by an absolute constant. Hence, conditioned on \mathbf{v}_j (which satisfies $\|\mathbf{v}_j\|_2 = \sqrt{d}$), $\frac{(\mathbf{v}_i^T \mathbf{v}_j)^2}{d}$ is a subexponential random variable with mean 1. Using Definition 5, there exists a constant $c > 0$ such that,

$$\mathbb{P}\left(\frac{(\mathbf{v}_i^T \mathbf{v}_j)^2}{d} > c \log d\right) \leq ed^{-4}.$$

Union bounding over all i, j pairs ensures that with probability ed^{-2} we have $\mu(\mathbf{A}^T) \leq c \frac{\log d}{d}$. Next, we use the standard result that for a matrix with columns of equal length, $\sigma_{\min}(\mathbf{A}^T) \geq d(1 - (m-1)\mu)$. The reader is referred to [52, Proposition 1]. Hence, $m \leq \frac{d}{2c \log d}$, gives $\sigma_{\min}(\mathbf{A}^T) \geq \frac{d}{2}$.

It remains to upper bound $\|\mathcal{A}(\bar{\mathbf{X}}_0)\|_2$. The i th entry of $\mathcal{A}(\bar{\mathbf{X}}_0)$ is equal to $|\mathbf{v}_i^T \bar{\mathbf{a}}|^2$, hence it is subexponential. Consequently, there exists a constant c' so that each entry is upper bounded by $\frac{c'}{2} \log d$ with probability $1 - ed^{-3}$. Union bounding, and using $m \leq d$, we find that $\|\mathcal{A}(\bar{\mathbf{X}}_0)\|_2 \leq \frac{c'}{2} \sqrt{m} \log d$ with probability $1 - ed^{-2}$. Combining with the estimate of $\sigma_{\min}(\mathbf{A}^T)$ completes the proof. ■

Comparison to Existing Literature: Proposition 4 is useful to estimate the performance of the sparse phase retrieval problem, in which \mathbf{a} is a k sparse vector, and we minimize a combination of the ℓ_1 norm and the nuclear norm to recover \mathbf{X}_0 . Combined with Theorem 1, Proposition 4 gives that, whenever $m \leq \frac{c_1 d}{\log d}$ and $\frac{c_2 \sqrt{m} \log d}{d} \leq \min\{\frac{\|\bar{\mathbf{X}}_0\|_1}{d}, \frac{\|\bar{\mathbf{X}}_0\|_*}{\sqrt{d}}\}$, recovery fails with high probability. Since $\|\bar{\mathbf{X}}_0\|_* = 1$ and $\|\bar{\mathbf{X}}_0\|_1 = \|\bar{\mathbf{a}}\|_1^2$, the failure condition reduces to,

$$m \leq \frac{c}{\log^2 d} \min\{\|\bar{\mathbf{a}}\|_1^4, d\}.$$

When $\bar{\mathbf{a}}$ is a k -sparse vector with ± 1 entries, in a similar flavor to Theorem 3, the right hand side has the form $\frac{c}{\log^2 d} \min\{k^2, d\}$.

We emphasize that the lower bound provided in [26] is directly comparable to our results. The authors in [26]

consider the same problem and give two results: first, if $m \geq \mathcal{O}(\|\bar{\mathbf{a}}\|_1^2 k \log d)$ then minimizing $\|\mathbf{X}\|_1 + \lambda \text{tr}(\mathbf{X})$ for suitable value of λ over the set of PSD matrices will exactly recover \mathbf{X}_0 with high probability. Secondly, their Theorem 1.3 provides a necessary condition (lower bound) on the number of measurements, under which the recovery program fails to recover \mathbf{X}_0 with high probability. In particular, their failure condition is $m \leq \min\{m_0, \frac{d}{40 \log d}\}$ where $m_0 = \frac{\max(\|\bar{\mathbf{a}}\|_1^2 - k/2, 0)^2}{500 \log^2 d}$.

Observe that both results require $m \leq \mathcal{O}\left(\frac{d}{\log d}\right)$. Focusing on the sparsity requirements, when the nonzero entries are sufficiently diffused (i.e. $\|\bar{\mathbf{a}}\|_1^2 \approx k$) both results yield $\mathcal{O}\left(\frac{\|\bar{\mathbf{a}}\|_1^4}{\log^2 d}\right)$ as a lower bound. On the other hand, if $\|\bar{\mathbf{a}}\|_1 \leq \sqrt{k/2}$, their lower bound becomes trivial while our lower bound still requires $\mathcal{O}\left(\frac{\|\bar{\mathbf{a}}\|_1^4}{\log^2 d}\right)$ measurements. $\|\bar{\mathbf{a}}\|_1 \leq \sqrt{k/2}$ can happen as soon as the nonzero entries are spiky, i.e. some of the entries are much larger than the rest. In this sense, our bounds are tighter. On the other hand, their lower bound includes the PSD constraint unlike ours.

D. Asymptotic Regime

While we discussed two cases in the nonasymptotic setup, we believe significantly more general results can be stated asymptotically ($m, n \rightarrow \infty$). For instance, under finite fourth moment constraint, thanks to the Bai-Yin law [53], asymptotically, the smallest singular value of a matrix with i.i.d. unit variance entries concentrates around $\sqrt{n} - \sqrt{m}$. Similarly, $\|\mathbf{A}\bar{\mathbf{x}}_0\|_2^2$ is the sum of independent variables; hence thanks to the law of large numbers, we will have $\frac{\|\mathbf{A}\bar{\mathbf{x}}_0\|_2^2}{m} \rightarrow 1$. Together, these yield $\frac{\|\mathbf{A}\bar{\mathbf{x}}_0\|_2}{\sigma_{\min}(\mathbf{A}^T)} \rightarrow \frac{\sqrt{m}}{\sqrt{n} - \sqrt{m}}$.

V. UPPER BOUNDS

We now state an upper bound on the simultaneous optimization for Gaussian measurement ensemble. Our upper bound will be in terms of distance to the dilated subdifferentials.

To accomplish this, we make use of the recent theory on the sample complexity of linear inverse problems. It has been recently shown that (3) exhibits a phase transition from failure with high probability to success with high probability when the number of Gaussian measurements are around the quantity $m_{PT} = \mathbf{D}(\text{cone}(\partial f(\mathbf{x}_0)))^2$ [9], [49]. This phenomenon was first observed by Donoho and Tanner, who calculated the phase transitions for ℓ_1 minimization and showed that $\mathbf{D}(\text{cone}(\partial \|\mathbf{x}_0\|_1))^2 \leq 2k \log \frac{en}{k}$ for a k -sparse vector in \mathbb{R}^n [54]. All of these works focus on signals with a single structure and do not study properties of a penalty that is a combination of norms. The next theorem relates the phase transition point of the joint optimization (3) to the individual subdifferentials.

Theorem 4: Suppose \mathbf{A} has i.i.d. $\mathcal{N}(0, 1)$ entries and let $f(\mathbf{x}) = \sum_{i=1}^{\tau} \lambda_i \|\mathbf{x}\|_{(i)}$. For positive scalars $\{\alpha_i\}_{i=1}^{\tau}$, let $\bar{\lambda}_i = \frac{\lambda_i \alpha_i}{\sum_{i=1}^{\tau} \lambda_i \alpha_i^{-1}}$ and define,

$$m_{up}(\{\alpha_i\}_{i=1}^{\tau}) := \left(\sum_i \bar{\lambda}_i \mathbf{D}(\alpha_i \partial \|\mathbf{x}_0\|_{(i)}) \right)^2.$$

If $m \geq (\sqrt{m_{up}} + t)^2 + 1$, then program (3) will succeed with probability $1 - 2 \exp(-\frac{t^2}{2})$.

Proof: Fix \mathbf{h} as an i.i.d. standard normal vector. Let \mathbf{g}_i be such that $\alpha_i \mathbf{g}_i$ is closest to \mathbf{h} over $\alpha_i \partial \|\mathbf{x}_0\|_{(i)}$. Let $\gamma = (\sum_i \frac{\lambda_i}{\alpha_i})^{-1}$. Then, we may write,

$$\begin{aligned} \inf_{\mathbf{g}' \in \text{cone}(\partial f(\mathbf{x}_0))} \|\mathbf{h} - \mathbf{g}'\|_2 &\leq \inf_{\mathbf{g} \in \partial f(\mathbf{x}_0)} \|\mathbf{h} - \gamma \mathbf{g}\|_2 \\ &\leq \|\mathbf{h} - \gamma \sum_i \lambda_i \mathbf{g}_i\|_2 \\ &= \|\mathbf{h} - \gamma \sum_i \frac{\lambda_i}{\alpha_i} \alpha_i \mathbf{g}_i\|_2 \\ &= \|\mathbf{h} - \sum_i \bar{\lambda}_i \alpha_i \mathbf{g}_i\|_2 \\ &\leq \sum_i \bar{\lambda}_i \|\mathbf{h} - \alpha_i \mathbf{g}_i\|_2 \\ &= \sum_i \bar{\lambda}_i \inf_{\mathbf{g}'_i \in \partial \|\mathbf{x}_0\|_{(i)}} \|\mathbf{h} - \alpha_i \mathbf{g}'_i\|_2. \end{aligned}$$

Taking the expectations of both sides and using the definition of $\mathbf{D}(\cdot)$, we find that

$$\mathbf{D}(\text{cone}(\partial f(\mathbf{x}_0))) \leq \sum_i \bar{\lambda}_i \mathbf{D}(\alpha_i \partial \|\mathbf{x}_0\|_{(i)}),$$

and $m_{up} \geq \mathbf{D}(\text{cone}(\partial f(\mathbf{x}_0)))^2$. The result then follows from the fact that, when $m \geq (\mathbf{D}(\text{cone}(\partial f(\mathbf{x}_0))) + t)^2 + 1$, recovery succeeds with probability $1 - 2 \exp(-\frac{t^2}{2})$. To see this, first, as discussed in [9, Proposition 3.6], $\mathbf{D}(\text{cone}(\partial f(\mathbf{x}_0)))$ is equal to the Gaussian width of the ‘‘tangent cone intersected with the unit ball’’ (see Theorem 6 for a definition of Gaussian width). Then, [9, Corollary 3.3] yields the probabilistic statement. ■

For Theorem 4 to be useful, choices of α_i should be made wisely. An obvious choice is letting,

$$\alpha_i^* = \arg \min_{\alpha_i \geq 0} \mathbf{D}(\alpha_i \partial \|\mathbf{x}_0\|_{(i)}). \quad (7)$$

With this choice, our upper bounds can be related to the individual sample complexities, which is equal to $\mathbf{D}(\text{cone}(\partial \|\mathbf{x}_0\|_{(i)}))^2$. [55, Proposition 1] shows that, if $\|\cdot\|_{(i)}$ is a decomposable norm, then,

$$0 \leq \mathbf{D}(\alpha_i^* \partial \|\mathbf{x}_0\|_{(i)}) - \mathbf{D}(\text{cone}(\partial \|\mathbf{x}_0\|_{(i)})) \leq 6.$$

Decomposability is defined and discussed in detail in Section VI-D. In particular, $\ell_1, \ell_{1,2}$ and the nuclear norm are decomposable. With this assumption, our upper bound suggests that the sample complexity of the simultaneous optimization is smaller than a certain convex combination of individual sample complexities.

Corollary 3: Suppose \mathbf{A} has i.i.d $\mathcal{N}(0, 1)$ entries and let $f(\mathbf{x}) = \sum_{i=1}^T \lambda_i \|\mathbf{x}\|_{(i)}$ for decomposable norms $\{\|\cdot\|_{(i)}\}_{i=1}^T$. Let $\{\alpha_i^*\}_{i=1}^T$ be as in (7) and assume they are strictly positive. Let $\bar{\lambda}_i^* = \frac{\lambda_i (\alpha_i^*)^{-1}}{\sum_{i=1}^T \lambda_i (\alpha_i^*)^{-1}}$ and define

$$\sqrt{m_{up}(\{\alpha_i^*\}_{i=1}^T)} := \sum_i \bar{\lambda}_i^* \mathbf{D}(\text{cone}(\partial \|\mathbf{x}_0\|_{(i)})) + 6.$$

If $m \geq (\sqrt{m_{up}} + t)^2 + 1$, then program (3) will succeed with probability $1 - 2 \exp(-\frac{t^2}{2})$.

Here, we used the fact that $\sum_i \bar{\lambda}_i^* = 1$ to take 6 out of the sum over i . We note that Corollaries 1 and 3 can be related in the case of sparse and low-rank matrices. For norms of interest, roughly speaking,

- $n\kappa_i^2$ is proportional to the sample complexity $\mathbf{D}(\text{cone}(\partial \|\mathbf{x}_0\|_{(i)}))^2$.
- L_i is proportional to $\frac{\sqrt{n}}{\alpha_i^*}$.

Consequently, the sample complexity of (3) will be upper and lower bounded by similar convex combinations.

A. Upper Bounds for the S&L Model

We now apply the bound of Theorem 4 to S&L matrices. To obtain simple and closed form bounds, we make use of existing results in the literature.

- [55, Table II]: If $\mathbf{x}_0 \in \mathbb{R}^n$ is a k sparse vector, choosing $\alpha_{\ell_1} = \sqrt{2 \log \frac{n}{k}}$, $\mathbf{D}(\alpha_{\ell_1} \partial \|\mathbf{x}_0\|_1)^2 \leq 2k \log \frac{en}{k}$.
- [56, Table III]: If $\mathbf{X}_0 \in \mathbb{R}^{d \times d}$ is a rank r matrix, choosing $\alpha_* = 2\sqrt{d}$, $\mathbf{D}(\alpha_* \partial \|\mathbf{X}_0\|_*)^2 \leq 6dr + 2d$.

Proposition 5: Suppose \mathbf{A} has i.i.d $\mathcal{N}(0, 1)$ entries and $\mathbf{X}_0 \in \mathbb{R}^{d \times d}$ is a rank r matrix whose nonzero entries lie on a $k \times k$ submatrix. For $0 \leq \beta \leq 1$, let $f(\mathbf{X}) = \lambda_{\ell_1} \|\mathbf{X}\|_1 + \lambda_* \|\mathbf{X}\|_*$ where $\lambda_{\ell_1} = \beta \sqrt{\log \frac{d}{k}}$ and $\lambda_* = (1 - \beta)\sqrt{d}$. Then, whenever,

$$m \geq \left(2\beta k \sqrt{\log \frac{ed}{k}} + (1 - \beta)\sqrt{6dr + 2d} + t \right)^2 + 1,$$

\mathbf{X}_0 can be recovered via (3) with probability $1 - 2 \exp(-\frac{t^2}{2})$.

Proof: To apply Theorem 4, we choose $\alpha_{\ell_1} = \sqrt{4 \log \frac{d}{k}}$ and $\alpha_* = 2\sqrt{d}$. \mathbf{X}_0 is effectively an (at most) k^2 sparse vector of size d^2 . Hence, $\alpha_{\ell_1} = \sqrt{2 \log \frac{d^2}{k^2}}$ and $\mathbf{D}(\alpha_{\ell_1} \|\mathbf{X}_0\|_1)^2 \leq 4k^2 \log \frac{ed}{k}$.

Now, for the choice of α_* , we have, $\mathbf{D}(\alpha_* \|\mathbf{X}_0\|_*)^2 \leq 6dr + 2d$. Observe that $\alpha_{\ell_1}^{-1} \lambda_{\ell_1} = \frac{\beta}{2}$, $\alpha_*^{-1} \lambda_* = \frac{1-\beta}{2}$ and apply Theorem 4 to conclude. ■

VI. GENERAL SIMULTANEOUSLY STRUCTURED MODEL RECOVERY

Recall the setup from Section II where we consider a vector $\mathbf{x}_0 \in \mathbb{R}^n$ whose structures are associated with a family of norms $\{\|\cdot\|_{(i)}\}_{i=1}^T$ and \mathbf{x}_0 satisfies the cone constraint $\mathbf{x}_0 \in \mathcal{C}$. This section is dedicated to the proofs of theorems in Section III-A and additional side results where the goal is to find lower bounds on the required number of measurements to recover \mathbf{x}_0 .

The following definitions will be helpful for the rest of our discussion. For a subspace M , denote its orthogonal complement by M^\perp . For a convex set M and a point \mathbf{x} , we define the projection operator as

$$\mathcal{P}_M(\mathbf{x}) = \arg \min_{\mathbf{u} \in M} \|\mathbf{x} - \mathbf{u}\|_2.$$

Given a cone \mathcal{C} , denote its dual cone by \mathcal{C}^* and polar cone by $\mathcal{C}^\circ = -\mathcal{C}^*$, where \mathcal{C}^* is defined as

$$\mathcal{C}^* = \{\mathbf{z} \mid \langle \mathbf{z}, \mathbf{v} \rangle \geq 0 \text{ for all } \mathbf{v} \in \mathcal{C}\}.$$

A. Preliminary Lemmas

We first show that the objective function $\max_{1 \leq i \leq \tau} \frac{\|\mathbf{x}\|_{(i)}}{\|\mathbf{x}_0\|_{(i)}}$ can be viewed as the ‘best’ among the functions mentioned in (3) for recovery of \mathbf{x}_0 .

Lemma 1: Consider the class of recovery programs in (3). If the program

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{C}}{\text{minimize}} \quad f_{\text{best}}(\mathbf{x}) \triangleq \max_{i=1, \dots, \tau} \frac{\|\mathbf{x}\|_{(i)}}{\|\mathbf{x}_0\|_{(i)}} \\ & \text{subject to} \quad \mathcal{A}(\mathbf{x}) = \mathcal{A}(\mathbf{x}_0) \end{aligned} \quad (8)$$

fails to recover \mathbf{x}_0 , then any member of this class will also fail to recover \mathbf{x}_0 .

Proof: Suppose (8) does not have \mathbf{x}_0 as an optimal solution and there exists \mathbf{x}' such that $f_{\text{best}}(\mathbf{x}') \leq f_{\text{best}}(\mathbf{x}_0)$. Then

$$\frac{1}{\|\mathbf{x}_0\|_{(i)}} \|\mathbf{x}'\|_{(i)} \leq f_{\text{best}}(\mathbf{x}') \leq f_{\text{best}}(\mathbf{x}_0) = 1,$$

for $i = 1, \dots, \tau$. This implies that

$$\|\mathbf{x}'\|_{(i)} \leq \|\mathbf{x}_0\|_{(i)}, \quad \text{for all } i = 1, \dots, \tau. \quad (9)$$

Conversely, given (9), we have $f_{\text{best}}(\mathbf{x}') \leq f_{\text{best}}(\mathbf{x}_0)$ from the definition of f_{best} .

Furthermore, since we assume $h(\cdot)$ in (3) is non-decreasing in its arguments and increasing in at least one of them, (9) implies $f(\mathbf{x}') \leq f(\mathbf{x}_0)$ for any such function $f(\cdot)$. Thus, failure of $f_{\text{best}}(\cdot)$ in recovery of \mathbf{x}_0 implies failure of any other function in (3). ■

The following lemma provides necessary conditions for \mathbf{x}_0 to be a minimizer of the problem (3).

Lemma 2: If \mathbf{x}_0 is a minimizer of the program (3), then there exist $\mathbf{v} \in \mathcal{C}^*$, \mathbf{z} , and $\mathbf{g} \in \partial f(\mathbf{x}_0)$ such that

$$\mathbf{g} - \mathbf{v} - \mathbf{A}^T \mathbf{z} = 0 \quad \text{and} \quad \langle \mathbf{x}_0, \mathbf{v} \rangle = 0.$$

The proof of Lemma 2 follows from the KKT conditions for (3) to have \mathbf{x}_0 as an optimal solution [57, Sec. 4.7].

The next lemma describes the subdifferential of any general function $f(\mathbf{x}) = h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)})$ as discussed in Section II-B.

Lemma 3: For any subgradient of the function $f(\mathbf{x}) = h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)})$ at $\mathbf{x} \neq 0$ defined by convex function $h(\cdot)$, there exists non-negative constants w_i , $i = 1, \dots, \tau$ such that

$$\mathbf{g} = \sum_{i=1}^{\tau} w_i \mathbf{g}_i$$

where $\mathbf{g}_i \in \partial \|\mathbf{x}_0\|_{(i)}$.

Proof: Consider the function $N(\mathbf{x}) = [\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)}]^T$ by which we have $f(\mathbf{x}) = h(N(\mathbf{x}))$. By [58, Th. 10.49] we have

$$\partial f(\mathbf{x}) = \bigcup \{ \partial(\mathbf{y}^T N(\mathbf{x})) : \mathbf{y} \in \partial h(N(\mathbf{x})) \}$$

where we used the convexity of f and h . Now notice that any $\mathbf{y} \in \partial h(N(\mathbf{x}))$ is a non-negative vector because of the monotonicity assumption on $h(\cdot)$. This implies that any subgradient $\mathbf{g} \in \partial f(\mathbf{x})$ is in the form of $\partial(\mathbf{w}^T N(\mathbf{x}))$ for

some nonnegative vector \mathbf{w} . The desired result simply follows because subgradients of conic combinations of norms are conic combinations of their subgradients (see [59]). ■

Using Lemmas 2 and 3, we now provide the proofs of Theorems 1 and 2.

B. Proof of Theorem 1

We prove the more general version of Theorem 1, which can take care of the cone constraint and alignment of subgradients over arbitrary subspaces. This will require us to extend the definition of correlation to handle subspaces. For a linear subspace $\mathcal{R} \in \mathbb{R}^n$ and a set $S \in \mathbb{R}^n$, we define,

$$\rho(\mathcal{R}, S) = \inf_{0 \neq \mathbf{s} \in S} \frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{s})\|_2}{\|\mathbf{s}\|_2}.$$

Proposition 6: Let

$$\sigma_{\mathcal{C}}(\mathbf{A}^T) = \inf_{\|\mathbf{z}\|_2=1} \frac{\|\mathcal{P}_{\mathcal{C}}(\mathbf{A}^T \mathbf{z})\|_2}{\|\mathbf{A}^T \mathbf{z}\|_2}.$$

Let \mathcal{R} be an arbitrary linear subspace orthogonal to the following cone,

$$\{ \mathbf{y} \in \mathbb{R}^n \mid \mathbf{x}_0^T \mathbf{y} = 0, \mathbf{y} \in \mathcal{C}^* \}. \quad (10)$$

Suppose,

$$\begin{aligned} \rho(\mathcal{R}, \partial f(\mathbf{x}_0)) &:= \inf_{\mathbf{g} \in \partial f(\mathbf{x}_0)} \frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{\|\mathbf{g}\|_2} \\ &> \frac{\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T))}{\sigma_{\mathcal{C}}(\mathbf{A}^T) \sigma_{\min}(\mathbf{A}^T)}. \end{aligned} \quad (11)$$

Then, \mathbf{x}_0 is not a minimizer of (3).

Proof: Suppose \mathbf{x}_0 is a minimizer of (3). From Lemma 2, there exist a $\mathbf{g} \in \partial f(\mathbf{x}_0)$, $\mathbf{z} \in \mathbb{R}^m$ and $\mathbf{v} \in \mathcal{C}^*$ such that

$$\mathbf{g} = \mathbf{A}^T \mathbf{z} + \mathbf{v} \quad (12)$$

and $\langle \mathbf{x}_0, \mathbf{v} \rangle = 0$. We first eliminate the contribution of \mathbf{v} in equation (12). Projecting both sides of (12) onto the subspace \mathcal{R} gives,

$$\mathcal{P}_{\mathcal{R}}(\mathbf{g}) = \mathcal{P}_{\mathcal{R}}(\mathbf{A}^T \mathbf{z}) = \mathcal{P}_{\mathcal{R}}(\mathbf{A}^T) \mathbf{z}. \quad (13)$$

Taking the ℓ_2 norms,

$$\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2 = \|\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T) \mathbf{z}\|_2 \leq \sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T)) \|\mathbf{z}\|_2. \quad (14)$$

Since $\mathbf{v} \in \mathcal{C}^*$, from Lemma 5 we have $\mathcal{P}_{\mathcal{C}}(-\mathbf{v}) = \mathcal{P}_{\mathcal{C}}(\mathbf{A}^T \mathbf{z} - \mathbf{g}) = 0$. Using Corollary 4,

$$\|\mathbf{g}\|_2 \geq \|\mathcal{P}_{\mathcal{C}}(\mathbf{A}^T \mathbf{z})\|_2. \quad (15)$$

From the initial assumption, for any $\mathbf{z} \in \mathbb{R}^m$, we have,

$$\sigma_{\mathcal{C}}(\mathbf{A}^T) \|\mathbf{A}^T \mathbf{z}\|_2 \leq \|\mathcal{P}_{\mathcal{C}}(\mathbf{A}^T \mathbf{z})\|_2 \quad (16)$$

Combining (15) and (16) yields $\|\mathbf{g}\|_2 \geq \sigma_{\mathcal{C}}(\mathbf{A}^T) \|\mathbf{A}^T \mathbf{z}\|_2$. Further incorporating (14), we find,

$$\begin{aligned} \frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T))} &\leq \|\mathbf{z}\|_2 \leq \frac{\|\mathbf{A}^T \mathbf{z}\|_2}{\sigma_{\min}(\mathbf{A}^T)} \\ &\leq \frac{\|\mathbf{g}\|_2}{\sigma_{\mathcal{C}}(\mathbf{A}^T) \sigma_{\min}(\mathbf{A}^T)}. \end{aligned}$$

Hence, if \mathbf{x}_0 is recoverable, then there exists $\mathbf{g} \in \partial f(\mathbf{x}_0)$ satisfying

$$\frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{\|\mathbf{g}\|_2} \leq \frac{\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T))}{\sigma_{\mathcal{C}}(\mathbf{A}^T)\sigma_{\min}(\mathbf{A}^T)},$$

completing the proof. \blacksquare

To obtain Theorem 1, choose $\mathcal{R} = \text{span}(\{\mathbf{x}_0\})$ and $\mathcal{C} = \mathbb{R}^n$. This choice of \mathcal{R} yields $\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T)) = \|\bar{\mathbf{x}}_0 \bar{\mathbf{x}}_0^T \mathbf{A}^T\|_2 = \|\mathbf{A} \bar{\mathbf{x}}_0\|_2$ and $\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2 = |\bar{\mathbf{x}}_0^T \mathbf{g}|$. Choosing $\mathcal{C} = \mathbb{R}^n$ yields $\sigma_{\mathcal{C}}(\mathbf{A}) = 1$. Also note that, for any choice of \mathcal{C} , \mathbf{x}_0 is orthogonal to (10) by definition.

C. Proof of Theorem 2

Rotational invariance of Gaussian measurements allows us to make full use of Proposition 6. The following is a generalization of Theorem 2.

Proposition 7: Consider the setup in Proposition 6 where \mathbf{A} has i.i.d $\mathcal{N}(0, 1)$ entries. Let,

$$m_{low} = \frac{n(1 - \bar{\mathbf{D}}(\mathcal{C}))\rho(\mathcal{R}, \partial f(\mathbf{x}_0))^2}{100},$$

and suppose that $\dim(\mathcal{R}) \leq m_{low}$. Then, whenever $m \leq m_{low}$, with probability $1 - 10 \exp(-\frac{1}{16} \min\{m_{low}, (1 - \bar{\mathbf{D}}(\mathcal{C}))^2 n\})$, (3) will fail for all functions $f(\cdot)$.

Proof: More measurements can only increase the chance of success. Hence, without losing generality, assume $m = m_{low}$ and $\dim(\mathcal{R}) \leq m$. The result will follow from Proposition 6. Recall that $m \leq \frac{(1 - \bar{\mathbf{D}}(\mathcal{C}))n}{100}$.

- $\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T)$ is statistically identical to a $\dim(\mathcal{R}) \times m$ matrix with i.i.d. $\mathcal{N}(0, 1)$ entries under proper unitary rotation. Hence, using [51, Corollary 5.35], with probability $1 - 2 \exp(-\frac{m}{8})$, $\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T)) \leq 1.5\sqrt{m} + \sqrt{\dim(\mathcal{R})} \leq 2.5\sqrt{m}$. With the same probability, $\sigma_{\min}(\mathbf{A}^T) \geq \sqrt{n} - 1.5\sqrt{m}$.
- From Theorem 7, using $m \leq \frac{(1 - \bar{\mathbf{D}}(\mathcal{C}))n}{100}$, with probability $1 - 6 \exp(-\frac{(1 - \bar{\mathbf{D}}(\mathcal{C}))^2 n}{16})$, $\sigma_{\mathcal{C}}^2(\mathbf{A}^T) \geq \frac{1 - \bar{\mathbf{D}}(\mathcal{C})}{4(1 + \bar{\mathbf{D}}(\mathcal{C}))} \geq \frac{1 - \bar{\mathbf{D}}(\mathcal{C})}{8}$.

Since $\frac{m}{n} \leq \frac{1}{30}$, combining these, with the desired probability,

$$\begin{aligned} \frac{\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T))}{\sigma_{\mathcal{C}}(\mathbf{A}^T)\sigma_{\min}(\mathbf{A}^T)} &\leq \sqrt{\frac{8}{1 - \bar{\mathbf{D}}(\mathcal{C})}} \frac{2.5\sqrt{m}}{\sqrt{n} - 1.5\sqrt{m}} \\ &< \frac{10\sqrt{m}}{\sqrt{(1 - \bar{\mathbf{D}}(\mathcal{C}))n}}. \end{aligned}$$

Finally, using Proposition 6 and the fact that $m \leq \frac{n(1 - \bar{\mathbf{D}}(\mathcal{C}))}{100} \rho(\mathcal{R}, \partial f(\mathbf{x}_0))^2$, with the same probability (3) fails. \blacksquare

To prove Theorem 2, choose $\mathcal{R} = \text{span}(\{\mathbf{x}_0\})$ and use the first statement of Proposition 1.

To show Corollary 1, choose $\mathcal{R} = \text{span}(\{\mathbf{x}_0\})$ and use the second statement of Proposition 1.

D. Enhanced Lower Bounds

From our initial results, it may look like our lower bounds are suboptimal. For instance, considering only the ℓ_1 norm,

$\kappa = \frac{\|\bar{\mathbf{x}}_0\|_1}{\sqrt{n}}$ lies between $\frac{1}{\sqrt{n}}$ and $\sqrt{\frac{k}{n}}$ for a k sparse signal. Combined with Theorem 2, this gives a lower bound of $\|\bar{\mathbf{x}}_0\|_1^2$ measurements. On the other hand, clearly, we need at least $\mathcal{O}(k)$ measurements to estimate a k sparse vector.

Indeed, Proposition 6 gives such a bound with a better choice of \mathcal{R} . In particular, let us choose $\mathcal{R} = \text{span}(\{\text{sign}(\mathbf{x}_0)\})$. For any $\mathbf{g} \in \partial \|\mathbf{x}_0\|_1$, we have that,

$$\frac{\left\langle \mathbf{g}, \frac{\text{sign}(\mathbf{x}_0)}{\sqrt{k}} \right\rangle}{L} = \sqrt{\frac{k}{n}} \implies \rho(\text{sign}(\mathbf{x}_0), \partial \|\mathbf{x}_0\|_1) = \sqrt{\frac{k}{n}}.$$

Hence, we immediately have $m \geq \mathcal{O}(k)$ as a lower bound. The idea of choosing such sign vectors can be generalized to the so-called decomposable norms.

Definition 7 (Decomposable Norm): A norm $\|\cdot\|$ is decomposable at $\mathbf{x} \in \mathbb{R}^n$ if there exist a subspace $T \subset \mathbb{R}^n$ and a vector $\mathbf{e} \in T$ such that the subdifferential at \mathbf{x} has the form

$$\partial \|\mathbf{x}\| = \{\mathbf{z} \in \mathbb{R}^n : \mathcal{P}_T(\mathbf{z}) = \mathbf{e}, \|\mathcal{P}_{T^\perp}(\mathbf{z})\|^* \leq 1\}.$$

We refer to T as the *support* and \mathbf{e} as the *sign vector* of \mathbf{x} with respect to $\|\cdot\|$.

Similar definitions are used in [13] and [60]. Our definition is simpler and less strict compared to these works. Note that L is a global property of the norm while \mathbf{e} and T depend on both the norm and the point under consideration (decomposability is a local property in this sense).

To give some intuition for Definition 7, we review examples of norms that arise when considering simultaneously sparse and low rank matrices. For a matrix $\mathbf{X} \in \mathbb{R}^{d_1 \times d_2}$, let $\mathbf{X}_{i,j}$, $\mathbf{X}_{i,\cdot}$ and $\mathbf{X}_{\cdot,j}$ denote its (i, j) entry, i th row and j th column respectively.

Lemma 4 (See [60]): The ℓ_1 norm, the $\ell_{1,2}$ norm and the nuclear norm are decomposable as follows.

- **ℓ_1 norm** is decomposable at every $\mathbf{x} \in \mathbb{R}^n$, with sign $\mathbf{e} = \text{sgn}(\mathbf{x})$, and support

$$T = \text{supp}(\mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^n : \mathbf{x}_i = 0 \Rightarrow \mathbf{y}_i = 0 \ \forall i \leq n\}.$$

- **$\ell_{1,2}$ norm** is decomposable at every $\mathbf{X} \in \mathbb{R}^{d_1 \times d_2}$. The support is

$$T = \{\mathbf{Y} \in \mathbb{R}^{d_1 \times d_2} : \mathbf{X}_{\cdot,i} = \mathbf{0} \Rightarrow \mathbf{Y}_{\cdot,i} = \mathbf{0} \ \forall i \leq d_2\},$$

and the sign vector $\mathbf{e} \in \mathbb{R}^{d_1 \times d_2}$ is obtained by normalizing the columns of \mathbf{X} present in the support, $\mathbf{e}_{\cdot,j} = \frac{\mathbf{X}_{\cdot,j}}{\|\mathbf{X}_{\cdot,j}\|_2}$ if $\|\mathbf{X}_{\cdot,j}\|_2 \neq 0$, and setting the rest of the columns to zero.

- **Nuclear norm** is decomposable at every $\mathbf{X} \in \mathbb{R}^{d_1 \times d_2}$. For a matrix \mathbf{X} with rank r and compact singular value decomposition $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^T$ where $\Sigma \in \mathbb{R}^{r \times r}$, we have $\mathbf{e} = \mathbf{U}\mathbf{V}^T$ and

$$\begin{aligned} T &= \{\mathbf{Y} \in \mathbb{R}^{d_1 \times d_2} : (\mathbf{I} - \mathbf{U}\mathbf{U}^T)\mathbf{Y}(\mathbf{I} - \mathbf{V}\mathbf{V}^T) = \mathbf{0}\} \\ &= \{\mathbf{Z}_1 \mathbf{V}^T + \mathbf{U}\mathbf{Z}_2^T \mid \mathbf{Z}_1 \in \mathbb{R}^{d_1 \times r}, \mathbf{Z}_2 \in \mathbb{R}^{d_2 \times r}\}. \end{aligned}$$

The next lemma shows that the sign vector \mathbf{e} will yield the largest correlation with the subdifferential and the best lower bound for such norms.




Fig. 4. An example of a decomposable norm: ℓ_1 norm is decomposable at $\mathbf{x}_0 = (1, 0)$. The sign vector \mathbf{e} , the support T , and shifted subspace T^\perp are illustrated. A subgradient \mathbf{g} at \mathbf{x}_0 and its projection onto T^\perp are also shown.

Lemma 5: Let $\|\cdot\|$ be a decomposable norm with support T and sign vector \mathbf{e} . For any $\mathbf{v} \neq 0$, we have that,

$$\rho(\mathbf{v}, \partial\|\mathbf{x}_0\|) \leq \rho(\mathbf{e}, \partial\|\mathbf{x}_0\|) \quad (17)$$

Also $\rho(\mathbf{e}, \partial\|\mathbf{x}_0\|) \geq \frac{\|\mathbf{e}\|_2}{L}$.

Proof: Let \mathbf{v} be a unit vector. Without losing generality, assume $\mathbf{v}^T \mathbf{e} \geq 0$. Pick a vector $\mathbf{z} \in T^\perp$ with $\|\mathbf{z}\|^* = 1$ such that $\mathbf{z}^T \mathbf{v} \leq 0$ (otherwise pick $-\mathbf{z}$). Now, consider the class of subgradients $\mathbf{g}(\alpha) = \mathbf{e} + \alpha \mathbf{z}$ for $1 \geq \alpha \geq -1$. Then,

$$\begin{aligned} \inf_{-1 \leq \alpha \leq 1} \frac{|\mathbf{v}^T \mathbf{g}(\alpha)|}{\|\mathbf{g}(\alpha)\|_2} &= \inf_{0 \leq \alpha \leq 1} \frac{|\mathbf{v}^T \mathbf{g}(\alpha)|}{\|\mathbf{g}(\alpha)\|_2} \\ &= \inf_{0 \leq \alpha \leq 1} \frac{|\mathbf{e}^T \mathbf{v} - \alpha |\mathbf{z}^T \mathbf{v}||}{(\|\mathbf{e}\|_2^2 + \alpha^2 \|\mathbf{z}\|_2^2)^{1/2}}. \end{aligned}$$

If $|\mathbf{z}^T \mathbf{v}| \geq \mathbf{e}^T \mathbf{v}$, then, the numerator can be made 0 and $\rho(\mathbf{v}, \partial\|\mathbf{x}_0\|) = 0$. Otherwise, the right hand side is decreasing function of α , hence the minimum is achieved at $\alpha = 1$, which gives,

$$\begin{aligned} \inf_{-1 \leq \alpha \leq 1} \frac{|\mathbf{v}^T \mathbf{g}(\alpha)|}{\|\mathbf{g}(\alpha)\|_2} &= \frac{|\mathbf{e}^T \mathbf{v} - |\mathbf{z}^T \mathbf{v}||}{(\|\mathbf{e}\|_2^2 + \|\mathbf{z}\|_2^2)^{1/2}} \\ &\leq \frac{|\mathbf{e}^T \mathbf{v}|}{(\|\mathbf{e}\|_2^2 + \|\mathbf{z}\|_2^2)^{1/2}} \\ &\leq \frac{\|\mathbf{e}\|_2}{(\|\mathbf{e}\|_2^2 + \|\mathbf{z}\|_2^2)^{1/2}} \\ &= \inf_{-1 \leq \alpha \leq 1} \frac{|\mathbf{e}^T \mathbf{g}(\alpha)|}{\|\mathbf{g}(\alpha)\|_2} \end{aligned}$$

where we used $\mathbf{e}^T \mathbf{g}(\alpha) = \mathbf{e}^T \mathbf{e} = \|\mathbf{e}\|_2^2$. Hence, along any direction \mathbf{z} , \mathbf{e} yields a higher minimum correlation than \mathbf{v} . To obtain (17), further take the infimum over all $\mathbf{z} \in T^\perp$, $\|\mathbf{z}\|^* \leq 1$ which yields the infimum over $\partial\|\mathbf{x}_0\|$. Finally, use $\|\mathbf{g}(\alpha)\|_2 \leq L$ to lower bound $\rho(\mathbf{e}, \partial\|\mathbf{x}_0\|)$. ■

Based on Lemma 5, the individual lower bounds are $\mathcal{O}(\frac{\|\mathbf{e}\|_2^2}{L^2})n$. Calculating $\frac{\|\mathbf{e}\|_2^2}{L^2}n$ for the norms in Lemma 4, reveals that, this quantity is k for a k sparse vector, cd_1 for a c -column sparse matrix and $r \max\{d_1, d_2\}$ for a rank r matrix. Compared to bounds obtained by using $\bar{\mathbf{x}}_0$, these new quantities are directly proportional to the true model complexities. Finally, we remark that, these new bounds correspond to choosing \mathbf{x}_0 that maximizes the value of

$\|\bar{\mathbf{x}}_0\|_1, \|\bar{\mathbf{x}}_0\|_*$ or $\|\bar{\mathbf{x}}_0\|_{1,2}$ while keeping sparsity, rank or column sparsity fixed. In particular, in these examples, \mathbf{e} has the same sparsity, rank, column sparsity as \mathbf{x}_0 .

The next proposition derives a correlation bound for the combination of decomposable norms as well as a simple lower bound on the sample complexity.

Proposition 8: Given decomposable norms $\|\cdot\|_{(i)}$ with supports T_i and sign vectors \mathbf{e}_i . Let $T_\cap = \bigcap_{1 \leq i \leq \tau} T_i$. Choose the subspace \mathcal{R} to be a subset of T_\cap .

- Assume $\langle \mathcal{P}_{\mathcal{R}}(\mathbf{e}_i), \mathcal{P}_{\mathcal{R}}(\mathbf{e}_j) \rangle \geq 0$ for all i, j and $\min_{1 \leq i \leq \tau} \frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{e}_i)\|_2}{\|\mathbf{e}_i\|_2} \geq v$. Then,

$$\rho(\mathcal{R}, \partial f(\mathbf{x}_0)) \geq \frac{v}{\sqrt{\tau}} \min_{1 \leq i \leq \tau} \rho(\mathbf{e}_i, \partial\|\mathbf{x}_0\|_{(i)}).$$

- Consider Proposition 6 with Gaussian measurements and suppose \mathcal{R} is orthogonal to the set (10). Let $f(\mathbf{x}) = \sum_{i=1}^{\tau} \lambda_i \|\mathbf{x}\|_{(i)}$ for some nonnegative $\{\lambda_i\}$'s. Then, if $m < \dim(\mathcal{R})$, (3) fails with probability 1.

Proof: Let $\mathbf{g} = \sum_{i=1}^{\tau} w_i \mathbf{g}_i$ for some $\mathbf{g}_i \in \partial\|\mathbf{x}_0\|_{(i)}$. First, $\|\mathbf{g}\|_2 \leq \sum_{i=1}^{\tau} w_i \|\mathbf{g}_i\|_2$. Next,

$$\begin{aligned} \|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2^2 &= \left\| \sum_{i=1}^{\tau} w_i \mathcal{P}_{\mathcal{R}}(\mathbf{e}_i) \right\|_2^2 \geq \sum_{i=1}^{\tau} w_i^2 \|\mathcal{P}_{\mathcal{R}}(\mathbf{e}_i)\|_2^2 \\ &\geq v^2 \sum_{i=1}^{\tau} w_i^2 \|\mathbf{e}_i\|_2^2 \geq \frac{v^2}{\tau} \left(\sum_{i=1}^{\tau} w_i \|\mathbf{e}_i\|_2 \right)^2. \end{aligned}$$

To see the second statement, consider the line (13) from the proof of Proposition 6. $\mathcal{P}_{\mathcal{R}}(\mathbf{g}) = \sum_{i=1}^{\tau} \lambda_i \mathcal{P}_{\mathcal{R}}(\mathbf{e}_i)$. On the other hand, the column space of $\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T)$ is an m -dimensional random subspace of \mathcal{R} . If $m < \dim(\mathcal{R})$, $\mathcal{P}_{\mathcal{R}}(\mathbf{g})$ is linearly independent of $\mathcal{P}_{\mathcal{R}}(\mathbf{A}^T)$ with probability 1 and (13) will not hold. ■

In the next section, we show how better choices of \mathcal{R} (based on the decomposability assumption) can improve the lower bounds for S&L recovery.

VII. PROOFS FOR SECTION III-B

Using the general framework of Section III-A, we now prove Theorem 3, which states various convex and nonconvex recovery results for the S&L models. We start with proofs of the convex recovery.

A. Convex Recovery Results for S&L

In this section, we prove the statements of Theorem 3 regarding convex approaches, using Theorem 2 and Proposition 7. We make use of the decomposable norms to obtain better lower bounds. Hence, we first state a result on the sign vectors and the supports of the S&L model following Lemma 4. The proof is provided in Appendix B.

Lemma 6: Denote the norm $\|\mathbf{X}^T\|_{1,2}$ by $\|\cdot\|^T_{1,2}$. Given a matrix $\mathbf{X}_0 \in \mathbb{R}^{d_1 \times d_2}$, let $\mathbf{E}_*, \mathbf{E}_c, \mathbf{E}_r$ and T_*, T_c, T_r be the sign vectors and supports for the norms $\|\cdot\|_*$, $\|\cdot\|_{1,2}$, $\|\cdot\|^T_{1,2}$ respectively. Then,

- $\mathbf{E}_*, \mathbf{E}_r, \mathbf{E}_c \in T_* \cap T_c \cap T_r$,
- $\langle \mathbf{E}_*, \mathbf{E}_r \rangle \geq 0$, $\langle \mathbf{E}_*, \mathbf{E}_c \rangle \geq 0$, and $\langle \mathbf{E}_c, \mathbf{E}_r \rangle \geq 0$.

1) *Proof of Theorem 3 Convex Cases:*

Proof of (a1): We use the functions $\|\cdot\|_{1,2}, \|\cdot\|^T_{1,2}$ and $\|\cdot\|_*$ without the cone constraint, i.e., $\mathcal{C} = \mathbb{R}^{d_1 \times d_2}$. We will apply Proposition 7 with $\mathcal{R} = T_* \cap T_c \cap T_r$. From Lemma 6 all the sign vectors lie on \mathcal{R} and they have pairwise nonnegative inner products. Consequently, applying Proposition 8

$$\rho(\mathcal{R}, \partial f(\mathbf{X}_0))^2 \geq \frac{1}{3} \min\left\{\frac{k_1}{d_1}, \frac{k_2}{d_2}, \frac{r}{\min\{d_1, d_2\}}\right\}.$$

If $m < \dim(\mathcal{R})$, then we have failure with probability 1. Hence, assume $m \geq \dim(\mathcal{R})$. Now, apply Proposition 7 with the given m_{low} .

Proof of (b1): In this case, we apply Lemma 12. We choose $\mathcal{R} = T_* \cap T_c \cap T_r \cap \mathbb{S}^n$, the norms are the same as in the general model, and $v \geq \frac{1}{\sqrt{2}}$. Also, pairwise inner products are positive, hence, using Proposition 8, $\rho(\mathcal{R}, \partial f(\mathbf{X}_0))^2 \geq \frac{1}{4} \min\{\frac{k}{d}, \frac{r}{d}\}$. Again, we may assume $m \geq \dim(\mathcal{R})$. Finally, based on Corollary 10, for the PSD cone we have $\bar{\mathbf{D}}(\mathcal{C}) \geq \frac{\sqrt{3}}{2}$. The result follows from Proposition 7 with the given m_{low} .

Proof of (c1): For the PSD cone, $\bar{\mathbf{D}}(\mathcal{C}) \geq \frac{\sqrt{3}}{2}$ and we simply use Theorem 2 to obtain the result by using $\kappa_{\ell_1}^2 = \frac{\|\bar{\mathbf{X}}_0\|_1^2}{d^2}$ and $\kappa_*^2 = \frac{\|\bar{\mathbf{X}}_0\|_*^2}{d}$.

2) *Proof of Corollary 2:* To show this, we use Theorem 1 and substitute κ 's corresponding to ℓ_1 and the nuclear norm. $\kappa_* = \frac{\|\bar{\mathbf{X}}_0\|_*}{\sqrt{d}}$ and $\kappa_{\ell_1} = \frac{\|\bar{\mathbf{X}}_0\|_{\ell_1}}{d}$. Also observe that, $\lambda_{\ell_1} L_{\ell_1} = \beta d$ and $\lambda_* L_* = (1 - \beta)d$. Hence, $\sum_{i=1}^2 \bar{\lambda}_i \kappa_i = \alpha \|\bar{\mathbf{X}}_0\|_1 + (1 - \alpha) \|\bar{\mathbf{X}}_0\|_* \sqrt{d}$. We then use Proposition 2 to conclude with sufficiently small $c_1, c_2 > 0$.

B. Nonconvex Recovery Results for S&L

While Theorem 3 states the result for Gaussian measurements, we prove the nonconvex recovery for the more general sub-gaussian measurements. We first state a lemma that will be useful in proving the nonconvex results. The proof is provided in Appendix C and uses standard arguments.

Lemma 7: Consider the set of matrices M in $\mathbb{R}^{d_1 \times d_2}$ that are supported over an $s_1 \times s_2$ submatrix with rank at most q . There exists a constant $c > 0$ such that whenever $m \geq c \min\{(s_1 + s_2)q, s_1 \log \frac{d_1}{s_1}, s_2 \log \frac{d_2}{s_2}\}$, with probability $1 - 2 \exp(-cm)$, $\mathcal{A}(\cdot) : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ with i.i.d. zero-mean and isotropic sub-gaussian rows satisfies

$$\mathcal{A}(\mathbf{X}) \neq 0, \quad \text{for all } \mathbf{X} \in M. \quad (18)$$

1) *Proof of Theorem 3 Nonconvex Cases:* Denote the sphere in $\mathbb{R}^{d_1 \times d_2}$ with unit Frobenius norm by $\mathcal{S}^{d_1 \times d_2}$.

Proof of (a2): Observe that the function $f(\mathbf{X}) = \frac{\|\mathbf{X}\|_{0,2}}{\|\mathbf{X}_0\|_{0,2}} + \frac{\|\mathbf{X}^T\|_{0,2}}{\|\mathbf{X}_0^T\|_{0,2}} + \frac{\text{rank}(\mathbf{X})}{\text{rank}(\mathbf{X}_0)}$ satisfies the triangle inequality and we have $f(\mathbf{X}_0) = 3$. Hence, if all null space elements $\mathbf{W} \in \text{Null}(\mathcal{A})$ satisfy $f(\mathbf{W}) > 6$, we have

$$f(\mathbf{X}) \geq f(\mathbf{X} - \mathbf{X}_0) - f(-\mathbf{X}_0) > 3,$$

for all feasible \mathbf{X} which implies \mathbf{X}_0 being the unique minimizer.

Consider the set M of matrices, which are supported over a $6k_1 \times 6k_2$ submatrix with rank at most $6r$. Observe that any \mathbf{Z} satisfying $f(\mathbf{Z}) \leq 6$ belongs to M . Hence ensuring

$\text{Null}(\mathcal{A}) \cap M = \{0\}$ would ensure $f(\mathbf{W}) > 6$ for all $\mathbf{W} \in \text{Null}(\mathcal{A})$. Since M is a cone, this is equivalent to $\text{Null}(\mathcal{A}) \cap (M \cap \mathcal{S}^{d_1 \times d_2}) = \emptyset$. Now, applying Lemma 7 with set M and $s_1 = 6k_1, s_2 = 6k_2, q = 6r$ we find the desired result.

Proof of (b2): Observe that due to the symmetry constraint,

$$f(\mathbf{X}) = \frac{\|\mathbf{X}\|_{0,2}}{\|\mathbf{X}_0\|_{0,2}} + \frac{\|\mathbf{X}^T\|_{0,2}}{\|\mathbf{X}_0^T\|_{0,2}} + \frac{\text{rank}(\mathbf{X})}{\text{rank}(\mathbf{X}_0)}.$$

Hence, the minimization is the same as (a2), the matrix is rank r contained in a $k \times k$ submatrix and we additionally have the positive semidefinite constraint which can only reduce the amount of required measurements compared to (a2). Consequently, the result follows by applying Lemma 7, similar to (a2).

Proof of (c2): Let $C = \{\mathbf{X} \neq 0 \mid f(\mathbf{X}) \leq f(\mathbf{X}_0)\}$. Since $\text{rank}(\mathbf{X}_0) = 1$, if $f(\mathbf{X}) \leq f(\mathbf{X}_0) = 2$, then $\text{rank}(\mathbf{X}) = 1$. With the symmetry constraint, this means $\mathbf{X} = \pm \mathbf{x}\mathbf{x}^T$ for some l -sparse \mathbf{x} . Observe that $\mathbf{X} - \mathbf{X}_0$ has rank at most 2 and is contained in a $2k \times 2k$ submatrix as $l \leq k$. Let M be the set of matrices that are symmetric and whose support lies in a $2k \times 2k$ submatrix. Using Lemma 7 with $q = 2, s_1 = s_2 = 2k$, whenever $m \geq ck \log \frac{n}{k}$, with desired probability all nonzero $\mathbf{W} \in M$ will satisfy $\mathcal{A}(\mathbf{W}) \neq 0$. Consequently, any $\mathbf{X} \in C$ has $\mathcal{A}(\mathbf{X}) \neq \mathcal{A}(\mathbf{X}_0)$, so that \mathbf{X}_0 is the unique minimizer.

C. Existence of a Matrix With Large κ 's

We now argue that, there exists an S&L matrix that has large $\kappa_{\ell_1}, \kappa_{\ell_1,2}$ and κ_* simultaneously. We provide have a deterministic construction that is close to optimal. Our construction is based on Hadamard matrices. $\mathbf{H}_n \in \mathbb{R}^{n \times n}$ is called a Hadamard matrix if it has ± 1 entries and orthogonal rows. Hadamard matrices exist for n that is an integer power of 2.

Using \mathbf{H}_n , our aim is to construct a $d_1 \times d_2$ S&L (k_1, k_2, r) matrix \mathbf{X}_0 that satisfies $\|\bar{\mathbf{X}}_0\|_1^2 \approx k_1 k_2$, $\|\bar{\mathbf{X}}_0\|_*^2 \approx r$, $\|\bar{\mathbf{X}}_0\|_{1,2}^2 \approx k_2$ and $\|\bar{\mathbf{X}}_0^T\|_{1,2}^2 \approx k_1$. To do this, we construct a $k_1 \times k_2$ matrix and then embed it into a larger $d_1 \times d_2$ matrix. The following lemma summarizes the construction.

Lemma 8: Without loss of generality, assume $k_2 \geq k_1 \geq r$. Let $\mathbf{H} := \mathbf{H}_{\lfloor \log_2 k_2 \rfloor}$. Let $\mathbf{X} \in \mathbb{R}^{k_1 \times k_2}$ be such that its i th row is equal to $[i - 1 \pmod{r}] + 1$ 'th row of \mathbf{H} followed by 0's for $1 \leq i \leq k_1$. Then,

$$\|\bar{\mathbf{X}}_0\|_1^2 \geq \frac{k_1 k_2}{2}, \quad \|\bar{\mathbf{X}}_0\|_*^2 \geq \frac{r}{2},$$

$$\|\bar{\mathbf{X}}_0\|_{1,2}^2 \geq \frac{k_2}{2}, \quad \|\bar{\mathbf{X}}_0^T\|_{1,2}^2 = k_1.$$

In particular, if $k_1 \equiv 0 \pmod{r}$ and k_2 is an integer power of 2, then,

$$\|\bar{\mathbf{X}}_0\|_1^2 = k_1 k_2, \quad \|\bar{\mathbf{X}}_0\|_*^2 = r,$$

$$\|\bar{\mathbf{X}}_0\|_{1,2}^2 = k_2, \quad \|\bar{\mathbf{X}}_0^T\|_{1,2}^2 = k_1.$$

Proof: The left $k_1 \times 2^{\lfloor \log_2 k_2 \rfloor}$ entries of \mathbf{X} are ± 1 , and the remaining entries are 0. This makes the calculation of ℓ_1 and $\ell_{1,2}$ and Frobenius norms trivial.




Fig. 5. Performance of the recovery program minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_{1,2}}{\|\mathbf{X}_0\|_{1,2}}\}$ with a PSD constraint. The dark region corresponds to the experimental region of failure due to insufficient measurements. As predicted by Theorem 3, the number of required measurements increases linearly with rd .

In particular, $\|\mathbf{X}_0\|_F^2 = \|\mathbf{X}_0\|_1 = k_1 2^{\lceil \log_2 k_2 \rceil}$, $\|\mathbf{X}_0\|_{1,2} = \sqrt{k_1} 2^{\lceil \log_2 k_2 \rceil}$ and $\|\mathbf{X}_0^T\|_{1,2} = k_1 2^{\frac{\lceil \log_2 k_2 \rceil}{2}}$. Substituting these yield the results for these norms.

To lower bound the nuclear norm, observe that, each of the first r rows of \mathbf{H} are repeated at least $\lfloor \frac{k_1}{r} \rfloor$ times in \mathbf{X} . Combined with orthogonality, this ensures that each singular value of \mathbf{X} that is associated with the j th row of \mathbf{H} is at least $\sqrt{2^{\lceil \log_2 k_2 \rceil} \lfloor \frac{k_1}{r} \rfloor}$ for all $1 \leq j \leq r$. Consequently

$$\|\mathbf{X}\|_* \geq r \sqrt{2^{\lceil \log_2 k_2 \rceil} \lfloor \frac{k_1}{r} \rfloor}.$$

Hence,

$$\|\bar{\mathbf{X}}\|_* \geq \frac{r \sqrt{2^{\lceil \log_2 k_2 \rceil} \lfloor \frac{k_1}{r} \rfloor}}{\sqrt{k_1} 2^{\lceil \log_2 k_2 \rceil}} = r \sqrt{\frac{1}{k_1} \lfloor \frac{k_1}{r} \rfloor}.$$

Finally, we use the fact that $\lfloor \frac{k_1}{r} \rfloor \geq \frac{k_1}{2r}$ as $k_1 \geq r$. ■

If we are allowed to use complex numbers, then one can apply the same idea with the Discrete Fourier Transform (DFT) matrix. Similar to \mathbf{H}_n , DFT has orthogonal rows and its entries have the same absolute value. However, it exists for any $n \geq 1$ which makes the argument more concise.

VIII. NUMERICAL EXPERIMENTS

In this section, we numerically verify our theoretical bounds on the number of measurements for the sparse and low-rank recovery problem. We demonstrate the empirical performance of the weighted maximum of the norms f_{best} (see Lemma 1), as well as the weighted sum of norms.

The experimental setup is as follows. Our goal is to explore how the number of required measurements m scales with the size of the matrix d . We consider a grid of (m, d) values, and generate at least 100 test instances for each grid point (in the boundary areas, we increase the number of instances to at least 200).

We generate the target matrix \mathbf{X}_0 by generating a $k \times r$ i.i.d. Gaussian matrix \mathbf{G} , and inserting the $k \times k$ matrix $\mathbf{G}\mathbf{G}^T$




Fig. 6. Performance of the recovery program minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ with a PSD constraint. $r = 1, k = 8$ and d is allowed to vary. The plot shows m versus d to illustrate the lower bound $\Omega(\min\{k^2, dr\})$ predicted by Theorem 3.

in an $d \times d$ matrix of zeros. We take $r = 1$ and $k = 8$ in all of the following experiments; even with these small values, we can observe the scaling predicted by our bounds. In each test, we measure the normalized recovery error $\frac{\|\mathbf{X} - \mathbf{X}_0\|_F}{\|\mathbf{X}_0\|_F}$ and declare successful recovery when this error is less than 10^{-4} . The optimization programs are solved using the CVX package [61], which calls the SDP solver SeDuMi [62].

We first test our bound in part (b) of Theorem 3, $\Omega(rd)$, on the number of measurements for recovery in the case of minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_{1,2}}{\|\mathbf{X}_0\|_{1,2}}\}$ over the set of PSD matrices. Figure 5 shows the results, which demonstrates m scaling linearly with d (note that $r = 1$).

Next, we replace the $\ell_{1,2}$ norm with the ℓ_1 norm and consider a recovery program that emphasizes entry-wise sparsity rather than block sparsity. Figure 6 demonstrates the lower bound $\Omega(\min\{k^2, d\})$ in Part (c) of Theorem 3 where we attempt to recover a rank-1 PSD matrix \mathbf{X}_0 by minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ subject to the measurements and a PSD constraint. The dashed green curve in the figure shows the empirical 95% failure boundary, depicting the region of failure with high probability that our results have predicted. It starts off growing linearly with d , when the term rd dominates the term k^2 , and then saturates as d grows and the k^2 term (which is a constant in our experiments) becomes dominant.

The penalty function $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ depends on the norm of \mathbf{X}_0 . In practice the norm of the solution is not known beforehand; a weighted sum of norms is used instead. In Figure 7 we examine the performance of the weighted sum of norms penalty in recovery of a rank-1 PSD matrix, for different weights. We pick $\lambda = 0.20$ and $\lambda = 0.35$ for a randomly generated matrix \mathbf{X}_0 . It can be seen that we get a reasonable result which is comparable to the performance of $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$.

In addition, we consider the *amount of error* in the recovery when the program fails. Figure 8 shows two curves below which we get a 90% failure, where for the upper (green) curve the normalized error threshold for declaring failure is 10^{-4} ,




Fig. 7. Performance of the recovery program minimizing $\text{tr}(\mathbf{X}) + \lambda\|\mathbf{X}\|_1$ with a PSD constraint, for $\lambda = 0.2$ (left) and $\lambda = 0.35$ (right).




Fig. 8. 90% frequency of failure where the threshold of recovery is 10^{-4} for the green (upper) and 0.05 for the red (lower) curve. $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ is minimized subject to the PSD constraint and the measurements.

and for the lower (red) curve it is a larger value of 0.05. We minimize $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ as the objective. We observe that when the recovery program has an error, it is very likely that this error is large, as the curves for 10^{-4} and 0.05 almost overlap. Thus, when the program fails, it fails badly. This observation agrees with intuition from similar problems in compressed sensing where sharp phase transition is observed.

As a final comment, observe that, in Figures 6, 7 and 8 the required amount of measurements slowly increases even when d is large and $k^2 = 64$ is the dominant constant term. While this is consistent with our lower bound of $\Omega(k^2, d)$,

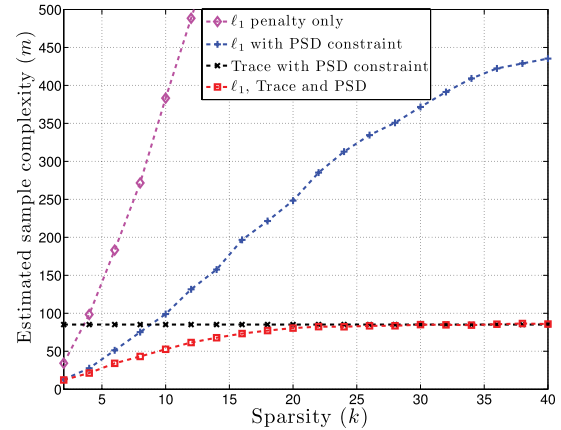


Fig. 9. We compare sample complexities of different approaches for a rank 1, 40×40 matrix as function of sparsity. The sample complexities were estimated by a search over m , where we chose the m with success rate closest to 50% (over 100 iterations).

the slow increase for constant k , can be explained by the fact that, as d gets larger, sparsity becomes the dominant structure and ℓ_1 minimization by itself requires $\mathcal{O}(k^2 \log \frac{d}{k})$ measurements rather than $\mathcal{O}(k^2)$. Hence for large d , the number of measurements can be expected to grow logarithmically in d .

In Figure 9, we compare the estimated phase transition points for different approaches for varying sparsity levels.

The algorithms we compare are,

- Minimize ℓ_1 norm,
- Minimize ℓ_1 norm subject to the positive-semidefinite constraint,
- Minimize trace norm subject to the positive-semidefinite constraint,
- Minimize $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ subject to the positive-semidefinite constraint.

Not surprisingly, the last option outperforms the rest in all cases. On the other hand, its performance is highly comparable to the minimum of the second and third approaches. For all regimes of sparsity, we observe that, measurements required by the last method is at least half as much as the minimum of the second and third methods.

IX. DISCUSSION

We have considered the problem of recovery of a simultaneously structured object from limited measurements. It is common in practice to combine known norm penalties corresponding to the individual structures (also known as regularizers in statistics and machine learning applications), and minimize this combined objective in order to recover the object of interest. The common use of this approach motivated us to analyze its performance, in terms of the smallest number of generic measurements needed for correct recovery. We showed that, under a certain assumption on the norms involved, the combined penalty requires more generic measurements than one would expect based on the degrees of freedom of the desired object. Our lower bounds on the required number of measurements implies that the combined norm penalty cannot perform significantly better than the best individual norm.

These results raise several interesting questions, and lead to directions for future work. We briefly outline some of these directions, as well as connections to some related problems.

A. Quantifying Recovery Failure via Error Bounds

We observe from the recovery error plots shown in Figure 8 that whenever our recovery program fails, it fails with a significant recovery error. The figure shows two curves under which recovery fails with high probability, where failure is defined by the normalized error $\|\mathbf{X} - \mathbf{X}_0\|_F / \|\mathbf{X}_0\|_F$ being above 10^{-4} and 0.05. The two curves almost coincide. This observation leads to the question of whether we can characterize how large the error is with a high probability over the random measurements. A lower bound on the recovery error as a function of the number of problem parameters will be very insightful.

B. Defining New Atoms for Simultaneously Structured Models

Our results show that combinations of individual norms do not exhibit a strong recovery performance. On the other hand, the seminal paper [9] proposes a remarkably general construction for an appropriate penalty given a set of atoms. Can we revisit a simultaneously structured recovery problem, and define new atoms that capture all structures at the same time? And can we obtain a new norm penalty induced by the convex hull of the atoms? Abstractly, the answer is yes, but such convex hulls may be hard to characterize, and the corresponding penalty may not be efficiently computable. It is interesting to find special cases where this construction can be carried out and results in a tractable problem. Recent developments in this direction include the “square norm” proposed by [40] for the low-rank tensor recovery which provably outperforms (2) for Gaussian measurements and the (k, q) -trace norm introduced by Richard *et al.* to estimate S&L matrices [63].

C. Algorithms for Minimizing Combination of Norms

Despite the limitation in their theoretical performance, in practice one may still need to solve convex relaxations that combine the different norms, *i.e.*, problem (3). Consider the special case of sparse and low-rank matrix recovery. All corresponding optimization problems mentioned in Theorem 3 can be expressed as a semidefinite program and solved by standard solvers; for example, for the numerical experiments in Section VIII we used the interior-point solver SeDuMi [62] via the modeling environment CVX [61]. However, interior point methods do not scale for problems with tens of thousands of matrix entries, which are common in machine learning applications. One future research direction is to explore first-order methods, which have been successful in solving ℓ_1 problems with a single structure (for example ℓ_1 or nuclear norm regularization alone). In particular, Alternating Directions Methods of Multipliers (ADMM) appears to be a promising candidate.

D. Characterizing the Tightness of the Lower Bounds

The results provided in this paper are negative in nature, as we characterize the lower bounds on the required amount of measurements for mixed convex recovery problems. However, it would be interesting to see how much we can gain by making use of multiple norms and how tight are these lower bounds. In [64], the authors investigate a specific simultaneous model where a signal $\mathbf{x} \in \mathbb{R}^n$ is sparse in both time and frequency domains, *i.e.*, \mathbf{x} and $\mathbf{D}\mathbf{x}$ are k_1, k_2 sparse respectively where \mathbf{D} is the Discrete Fourier Transform matrix. For recovery, the authors consider minimizing $\|\mathbf{x}\|_1 + \lambda \|\mathbf{D}\mathbf{x}\|_1$ subject to measurements. Intuitively, results of this paper would suggest the necessity of $\Omega(\min\{k_1, k_2\})$ measurements for successful recovery. On the other hand, best of the individual functions (ℓ_1 norms) will require $\Omega(\min\{k_1 \log \frac{n}{k_1}, k_2 \log \frac{n}{k_2}\})$ measurements. In [64], it is shown that the mixed approach will require as little as $\max\{k_1, k_2\} \log \log n$ under mild assumptions.

This shows that the mixed approach can result in a logarithmic improvement over the individual functions when $k_1 \approx k_2$ and the lower bound given by this paper might be achievable up to a small factor.

E. Connection to Sparse PCA

The sparse PCA problem (see [65]–[67]) seeks sparse principal components given a (possibly noisy) data matrix. Several formulations for this problem exist, and many algorithms have been proposed. In particular, a popular algorithm is the SDP relaxation proposed in [67], which is based on the following formulation.

For the first principal component to be sparse, we seek an $\mathbf{x} \in \mathbb{R}^n$ that maximizes $\mathbf{x}^T \mathbf{A} \mathbf{x}$ for a given data matrix \mathbf{A} , and minimizes $\|\mathbf{x}\|_0$. Similar to the sparse phase retrieval problem, this problem can be reformulated in terms of a rank-1, PSD matrix $\mathbf{X} = \mathbf{x}\mathbf{x}^T$ which is also row- and column-sparse. Thus we seek a simultaneously low-rank and sparse \mathbf{X} . This problem is different from the recovery problem studied in this paper, since we do not have m random measurements of \mathbf{X} . Yet, it will be interesting to connect this paper’s results to the sparse PCA problem to potentially provide new insights for sparse PCA.

APPENDIX A PROPERTIES OF CONES

In this appendix, we state some results regarding cones which are used in the proof of general recovery. Recall the definitions of polar and dual cones from Section II.

Theorem 5 (Moreau’s Decomposition Theorem, [68]): Let \mathcal{C} be a closed and convex cone in \mathbb{R}^n . Then, for any $\mathbf{x} \in \mathbb{R}^n$, we have

- $\mathbf{x} = \mathcal{P}_{\mathcal{C}}(\mathbf{x}) + \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x})$.
- $\langle \mathcal{P}_{\mathcal{C}}(\mathbf{x}), \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x}) \rangle = 0$.

Lemma 9 (Projection is Nonexpansive): Let $\mathcal{C} \in \mathbb{R}^n$ be a closed and convex set and $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ be vectors. Then,

$$\|\mathcal{P}_{\mathcal{C}}(\mathbf{a}) - \mathcal{P}_{\mathcal{C}}(\mathbf{b})\|_2 \leq \|\mathbf{a} - \mathbf{b}\|_2.$$

Corollary 4: Let \mathcal{C} be a closed convex cone and \mathbf{a}, \mathbf{b} be vectors satisfying $\mathcal{P}_{\mathcal{C}}(\mathbf{a} - \mathbf{b}) = 0$. Then

$$\|\mathbf{b}\|_2 \geq \|\mathcal{P}_{\mathcal{C}}(\mathbf{a})\|_2.$$

Proof: Using Lemma 9, we have $\|\mathcal{P}_{\mathcal{C}}(\mathbf{a})\|_2 = \|\mathcal{P}_{\mathcal{C}}(\mathbf{a}) - \mathcal{P}_{\mathcal{C}}(\mathbf{a} - \mathbf{b})\|_2 \leq \|\mathbf{b}\|_2$. ■

The unit sphere in \mathbb{R}^n will be denoted by \mathcal{S}^{n-1} for the following theorems.

Theorem 6 (Escape Through a Mesh, [69]): For a given set $\mathcal{D} \in \mathcal{S}^{n-1}$, define the *Gaussian width* as

$$\omega(\mathcal{D}) = \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{D}} \langle \mathbf{x}, \mathbf{g} \rangle \right],$$

in which $\mathbf{g} \in \mathbb{R}^n$ has i.i.d. standard Gaussian entries. Given m , let $d = \sqrt{n-m} - \frac{1}{4\sqrt{n-m}}$. Provided that $\omega(\mathcal{D}) \leq d$ a random m -dimensional subspace which is uniformly drawn w.r.t. Haar measure will have no intersection with \mathcal{D} with probability at least

$$1 - 3.5 \exp(-(d - \omega(\mathcal{D}))^2). \quad (\text{A.1})$$

Theorem 7: Consider a random Gaussian map $\mathcal{G}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ with i.i.d. entries and the corresponding adjoint operator \mathcal{G}^* . Let \mathcal{C} be a closed and convex cone and recalling Definition 3, let

$$\zeta(\mathcal{C}) := 1 - \bar{\mathbf{D}}(\mathcal{C}), \quad \gamma(\mathcal{C}) := 2\sqrt{\frac{1 + \bar{\mathbf{D}}(\mathcal{C})}{1 - \bar{\mathbf{D}}(\mathcal{C})}}.$$

where $\bar{\mathbf{D}}(\mathcal{C}) = \frac{\mathbf{D}(\mathcal{C})}{\sqrt{n}}$. Then, if $m \leq \frac{7\zeta(\mathcal{C})}{16}n$, with probability at least $1 - 6 \exp(-(\frac{\zeta(\mathcal{C})}{4})^2 n)$, for all $\mathbf{z} \in \mathbb{R}^n$ we have

$$\|\mathcal{G}^*(\mathbf{z})\|_2 \leq \gamma(\mathcal{C}) \|\mathcal{P}_{\mathcal{C}}(\mathcal{G}^*(\mathbf{z}))\|_2. \quad (\text{A.2})$$

Proof: For notational simplicity, let $\zeta = \zeta(\mathcal{C})$ and $\gamma = \gamma(\mathcal{C})$. Consider the set

$$\mathcal{D} = \{\mathbf{x} \in \mathcal{S}^{n-1} : \|\mathbf{x}\|_2 \geq \gamma \|\mathcal{P}_{\mathcal{C}}(\mathbf{x})\|_2\}$$

and we are going to show that with high probability, the range of \mathcal{G}^* misses \mathcal{D} . Using Theorem 5, for any $\mathbf{x} \in \mathcal{D}$, we may write

$$\begin{aligned} \langle \mathbf{x}, \mathbf{g} \rangle &= \langle \mathcal{P}_{\mathcal{C}}(\mathbf{x}) + \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x}), \mathcal{P}_{\mathcal{C}}(\mathbf{g}) + \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g}) \rangle \\ &\leq \langle \mathcal{P}_{\mathcal{C}}(\mathbf{x}), \mathcal{P}_{\mathcal{C}}(\mathbf{g}) \rangle + \langle \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x}), \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g}) \rangle \\ &\leq \|\mathcal{P}_{\mathcal{C}}(\mathbf{x})\|_2 \|\mathcal{P}_{\mathcal{C}}(\mathbf{g})\|_2 + \|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x})\|_2 \|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g})\|_2 \\ &\leq \gamma^{-1} \|\mathcal{P}_{\mathcal{C}}(\mathbf{g})\|_2 + \|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g})\|_2 \end{aligned} \quad (\text{A.3})$$

where in (A.3) we used the fact that elements of \mathcal{C} and \mathcal{C}° have nonpositive inner products and $\|\mathcal{P}_{\mathcal{C}}(\mathbf{x})\|_2 \leq \|\mathbf{x}\|_2$ is by Lemma 9. Hence, from the definition of Gaussian width,

$$\begin{aligned} \omega(\mathcal{D}) &= \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{D}} \langle \mathbf{x}, \mathbf{g} \rangle \right] \\ &\leq \gamma^{-1} \mathbb{E} [\|\mathcal{P}_{\mathcal{C}}(\mathbf{g})\|_2] + \mathbb{E} [\|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g})\|_2] \\ &\leq \sqrt{n}(\gamma^{-1} \bar{\mathbf{D}}(\mathcal{C}^\circ) + \bar{\mathbf{D}}(\mathcal{C})) \leq \frac{2-\zeta}{2} \sqrt{n}. \end{aligned}$$

Where we used the fact that $\gamma \geq \frac{2\bar{\mathbf{D}}(\mathcal{C}^\circ)}{1-\bar{\mathbf{D}}(\mathcal{C})}$; which follows from $\bar{\mathbf{D}}(\mathcal{C})^2 + \bar{\mathbf{D}}(\mathcal{C}^\circ)^2 \leq 1$ (see Theorem 5 above). Hence, whenever,

$$m \leq \frac{7\zeta}{16}n \leq (1 - (\frac{4-\zeta}{4})^2)n = m',$$

using the upper bound on $\omega(\mathcal{D})$, we have,

$$\begin{aligned} (\sqrt{n-m} - \omega(\mathcal{D}) - \frac{1}{4\sqrt{n-m}})^2 \\ \geq (\sqrt{n-m} - \omega(\mathcal{D}))^2 - \frac{1}{2} \geq (\frac{\zeta}{4})^2 n - \frac{1}{2}. \end{aligned}$$

Now, using Theorem 6, the range space of \mathcal{G}^* will miss the undesired set \mathcal{D} with probability at least $1 - 3.5 \exp(-(\frac{\zeta}{4})^2 n + \frac{1}{2}) \geq 1 - 6 \exp(-(\frac{\zeta}{4})^2 n)$. ■

Lemma 10: Consider the cones \mathbb{S}^d and \mathbb{S}_+^d in the space $\mathbb{R}^{d \times d}$. Then, $\bar{\mathbf{D}}(\mathbb{S}^d) < \frac{1}{\sqrt{2}}$ and $\bar{\mathbf{D}}(\mathbb{S}_+^d) < \frac{\sqrt{3}}{2}$.

Proof: Let \mathbf{G} be a $d \times d$ matrix with i.i.d. standard normal entries. Set of symmetric matrices \mathbb{S}^d is an $\frac{d(d+1)}{2}$ dimensional subspace of $\mathbb{R}^{d \times d}$. Hence, $\mathbb{E} \|\mathcal{P}_{\mathbb{S}^d}(\mathbf{G})\|_F^2 = \frac{d(d+1)}{2}$ and $\mathbb{E} \|\mathcal{P}_{(\mathbb{S}^d)^\circ}(\mathbf{G})\|_F^2 = \frac{d(d-1)}{2}$. Hence,

$$\bar{\mathbf{D}}(\mathbb{S}^d) = \sqrt{\frac{d(d-1)}{2d^2}} < \frac{1}{\sqrt{2}}.$$

To prove the second statement, observe that projection of a matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ onto \mathbb{S}_+^d is obtained by first projecting \mathbf{A} onto \mathbb{S}^d and then taking the matrix induced by the positive eigenvalues of $\mathcal{P}_{\mathbb{S}^d}(\mathbf{A})$. Since, \mathbf{G} and $-\mathbf{G}$ are identically distributed and \mathbb{S}_+^d is a self dual cone, $\mathcal{P}_{\mathbb{S}_+^d}(\mathbf{G})$ is identically distributed as $-\mathcal{P}_{\mathbb{S}_+^d}(\mathbf{G})$ where $\mathbb{S}_-^d = (\mathbb{S}_+^d)^\circ$ stands for negative semidefinite matrices. Hence,

$$\mathbb{E} \|\mathcal{P}_{\mathbb{S}_+^d}(\mathbf{G})\|_F^2 = \frac{\mathbb{E} \|\mathcal{P}_{\mathbb{S}^d}(\mathbf{G})\|_F^2}{2} = \frac{d(d+1)}{4},$$

$$\mathbb{E} \|\mathcal{P}_{(\mathbb{S}_+^d)^\circ}(\mathbf{G})\|_F^2 = \frac{d(3d-1)}{4}.$$

Consequently, $\bar{\mathbf{D}}(\mathbb{S}_+^d) = \sqrt{\frac{3}{4} - \frac{1}{4d}} < \sqrt{\frac{3}{4}}$. ■

APPENDIX B

NORMS IN SPARSE AND LOW-RANK MODEL

A. Relevant Notation for the Proofs

Let $[k]$ denote the set $\{1, 2, \dots, k\}$. Let S_c, S_r denote the indexes of the nonzero columns and rows of \mathbf{X}_0 so that nonzero entries of \mathbf{X}_0 lies on $S_r \times S_c$ submatrix. $\mathcal{S}_c, \mathcal{S}_r$ denotes the k_1, k_2 dimensional subspaces of vectors whose nonzero entries lie on S_c and S_r respectively.

Let \mathbf{X}_0 have singular value decomposition $\mathbf{U}\Sigma\mathbf{V}^T$ such that $\Sigma \in \mathbb{R}^{r \times r}$ and columns of \mathbf{U}, \mathbf{V} lies on $\mathcal{S}_c, \mathcal{S}_r$ respectively.

B. Proof of Lemma 6

Proof: Observe that $T_c = \mathbb{R}^d \times \mathcal{S}_c$ and $T_r = \mathcal{S}_r \times \mathbb{R}^d$ hence $T_c \cap T_r$ is the set of matrices that lie on $\mathcal{S}_r \times \mathcal{S}_c$. Hence, $\mathbf{E}_* = \mathbf{U}\mathbf{V}^T \in T_c \cap T_r$. Similarly, \mathbf{E}_c and \mathbf{E}_r are the matrices obtained by scaling columns and rows of \mathbf{X}_0 to have unit size. As a result, they also lie on $\mathcal{S}_r \times \mathcal{S}_c$ and $T_c \cap T_r$. $\mathbf{E}_* \in T_*$ by definition.

Next, we may write $\mathbf{E}_c = \mathbf{X}_0 \mathbf{D}_c$ where \mathbf{D}_c is the scaling nonnegative diagonal matrix. Consequently, \mathbf{E}_c lies on the range space of \mathbf{X}_0 and belongs to T_* . This follows from definition of T_* in Lemma 4 and the fact that $(\mathbf{I} - \mathbf{U}\mathbf{U}^T)\mathbf{E}_c = 0$.

In the exact same way, $\mathbf{E}_r = \mathbf{D}_r \mathbf{X}_0$ for some nonnegative diagonal \mathbf{D}_r and lies on the range space of \mathbf{X}^T and hence lies on T_* . Consequently, $\mathbf{E}_*, \mathbf{E}_c, \mathbf{E}_r$ lies on $T_c \cap T_r \cap T_*$.

Now, consider

$$\begin{aligned} \langle \mathbf{E}_c, \mathbf{E}_* \rangle &= \langle \mathbf{X}_0 \mathbf{D}_c, \mathbf{U} \mathbf{V}^T \rangle = \text{tr}(\mathbf{V} \mathbf{U}^T \mathbf{U} \Sigma \mathbf{V}^T \mathbf{D}_c) \\ &= \text{tr}(\mathbf{V} \Sigma \mathbf{V}^T \mathbf{D}_c) \geq 0 \end{aligned}$$

since both $\mathbf{V} \Sigma \mathbf{V}^T$ and \mathbf{D}_c are positive semidefinite matrices. In the exact same way, we have $\langle \mathbf{E}_c, \mathbf{E}_* \rangle \geq 0$. Finally,

$$\langle \mathbf{E}_c, \mathbf{E}_r \rangle = \langle \mathbf{X}_0 \mathbf{D}_c, \mathbf{D}_r \mathbf{X}_0 \rangle = \text{tr}(\mathbf{D}_c \mathbf{X}_0^T \mathbf{D}_r \mathbf{X}_0) \geq 0,$$

since both \mathbf{D}_c and $\mathbf{X}_0^T \mathbf{D}_r \mathbf{X}_0$ are PSD matrices. Overall, the pairwise inner products of $\mathbf{E}_r, \mathbf{E}_c, \mathbf{E}_*$ are nonnegative. ■

C. Results on the Positive Semidefinite Constraint

Lemma 11: Assume $\mathbf{X}, \mathbf{Y} \in \mathbb{S}_+^d$ have eigenvalue decompositions $\mathbf{X} = \sum_{i=1}^{\text{rank}(\mathbf{X})} \sigma_i \mathbf{u}_i \mathbf{u}_i^T$ and $\mathbf{Y} = \sum_{i=1}^{\text{rank}(\mathbf{Y})} c_i \mathbf{v}_i \mathbf{v}_i^T$. Further, assume $\langle \mathbf{Y}, \mathbf{X} \rangle = 0$. Then, $\mathbf{U}^T \mathbf{Y} = \mathbf{0}$ where $\mathbf{U} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_{\text{rank}(\mathbf{X})}]$.

Proof: Observe that,

$$\langle \mathbf{Y}, \mathbf{X} \rangle = \sum_{i=1}^{\text{rank}(\mathbf{X})} \sum_{j=1}^{\text{rank}(\mathbf{Y})} \sigma_i c_j |\mathbf{u}_i^T \mathbf{v}_j|^2.$$

Since $\sigma_i, c_j > 0$, right hand side is 0 if and only if $\mathbf{u}_i^T \mathbf{v}_j = 0$ for all i, j . Hence, the result follows. ■

Lemma 12: Assume $\mathbf{X}_0 \in \mathbb{S}_+^d$ so that in Section IX-A, $S_c = S_r, T_c = T_r, k_1 = k_2 = k$ and $\mathbf{U} = \mathbf{V}$. Let $\mathcal{R} = T_c \cap T_r \cap T_* \cap \mathbb{S}^d, S_* = T_* \cap \mathbb{S}^d$, and,

$$\mathcal{Y} = \{\mathbf{Y} \mid \mathbf{Y} \in (\mathbb{S}_+^d)^*, \langle \mathbf{Y}, \mathbf{X}_0 \rangle = 0\},$$

Then, the following statements hold.

- $S_* \subseteq \text{span}(\mathcal{Y})^\perp$. Hence, $\mathcal{R} \subseteq S_*$ and is orthogonal to \mathcal{Y} .
- $\mathbf{E}_* \in \mathcal{R}, \frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{E}_c)\|_F}{\|\mathbf{E}_c\|_F} = \frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{E}_r)\|_F}{\|\mathbf{E}_r\|_F} \geq \frac{1}{\sqrt{2}}$.

Proof: The dual of \mathbb{S}_+^d with respect to $\mathbb{R}^{d \times d}$ is the set sum of \mathbb{S}_+^d and Skew^d where Skew^d is the set of skew-symmetric matrices. Now, assume, $\mathbf{Y} \in \mathcal{Y}$ and $\mathbf{X} \in S_*$. Then, $\langle \mathbf{Y}, \mathbf{X} \rangle = \langle \frac{\mathbf{Z}}{2}, \mathbf{X} \rangle$ where $\mathbf{Z} = \mathbf{Y} + \mathbf{Y}^T \in \mathbb{S}_+^d$ and $\langle \mathbf{Z}, \mathbf{X}_0 \rangle = 0$. Since \mathbf{X}_0, \mathbf{Z} are both PSD, applying Lemma 11, we have $\mathbf{U}^T \mathbf{Z} = \mathbf{0}$ hence $(\mathbf{I} - \mathbf{U} \mathbf{U}^T) \mathbf{Z} (\mathbf{I} - \mathbf{U} \mathbf{U}^T) = \mathbf{Z}$ which means $\mathbf{Z} \in T_*^\perp$. Hence, $\langle \mathbf{Z}, \mathbf{X} \rangle = \langle \mathbf{Y}, \mathbf{X} \rangle = 0$ as $\mathbf{X} \in S_* \subset T_*$. Hence, $\text{span}(\mathcal{Y}) \subseteq S_*^\perp$.

For the second statement, let $T_\cap = T_* \cap T_c \cap T_r$. Recalling Lemma 6, observe that $\mathbf{E}_* \in T_\cap$. Since \mathbf{E}_* is also symmetric, $\mathbf{E}_* \in \mathcal{R}$. Similarly, $\mathbf{E}_c, \mathbf{E}_r \in T_\cap, \langle \mathbf{E}_c, \mathbf{E}_r \rangle \geq 0$ and $\|\mathcal{P}_{\mathcal{R}}(\mathbf{E}_c)\|_F = \frac{\|\mathbf{E}_c + \mathbf{E}_r\|_F}{2} \geq \frac{\|\mathbf{E}_c\|_F}{\sqrt{2}}$. Similar result is true for \mathbf{E}_r . ■

APPENDIX C

RESULTS ON NON-CONVEX RECOVERY

Next two lemmas are standard results on sub-gaussian measurement operators.

Lemma 13 (Properties of Sub-Gaussian Mappings): Assume \mathbf{X} is an arbitrary matrix with unit Frobenius norm. A measurement operator $\mathcal{A}(\cdot)$ with i.i.d zero-mean

isotropic subgaussian rows (see Section IV) satisfies the following:

- $\mathbb{E}[\|\mathcal{A}(\mathbf{X})\|_2^2] = m$.
- There exists an absolute constant $c > 0$ such that, for all $1 \geq \varepsilon \geq 0$, we have

$$\mathbb{P}(\|\mathcal{A}(\mathbf{X})\|_2^2 - m \geq \varepsilon m) \leq 2 \exp(-c\varepsilon^2 m).$$

Proof: Observe that, when $\|\mathbf{X}\|_F = 1$, entries of $\mathcal{A}(\mathbf{X})$ are zero-mean with unit variance. Hence, the first statement follows directly. For the second statement, we use the fact that square of a sub-gaussian random variable is sub-exponential and view $\|\mathcal{A}(\mathbf{X})\|_2^2$ as a sum of m i.i.d. subexponentials with unit mean. Then, result follows from [51, Corollary 5.17]. ■

For the consequent lemmas, $\mathcal{S}^{d_1 \times d_2}$ denotes the unit Frobenius norm sphere in $\mathbb{R}^{d_1 \times d_2}$.

Lemma 14: Let $\mathcal{D} \in \mathbb{R}^{d_1 \times d_2}$ be an arbitrary cone and $\mathcal{A}(\cdot) : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ be a measurement operator with i.i.d zero-mean and isotropic sub-gaussian rows. Assume that the set $\bar{\mathcal{D}} = \mathcal{S}^{d_1 \times d_2} \cap \mathcal{D}$ has ε -covering number bounded above by $\eta(\varepsilon)$. Then, there exists constants $c_1, c_2 > 0$ such that whenever $m \geq c_1 \log \eta(1/4)$, with probability $1 - 2 \exp(-c_2 m)$, we have

$$\mathcal{D} \cap \text{Null}(\mathcal{A}) = \{0\}.$$

Proof: Let $\eta = \eta(\frac{1}{4})$, and $\{\mathbf{X}_i\}_{i=1}^\eta$ be a $\frac{1}{4}$ -covering of $\bar{\mathcal{D}}$. With probability at least $1 - 2\eta \exp(-c\varepsilon^2 m)$, for all i , we have

$$(1 - \varepsilon)m \leq \|\mathcal{A}(\mathbf{X}_i)\|_2^2 \leq (1 + \varepsilon)m.$$

Now, let $\mathbf{X}_{\text{sup}} = \arg \sup_{\mathbf{X} \in \bar{\mathcal{D}}} \|\mathcal{A}(\mathbf{X})\|_2$. Choose $1 \leq a \leq \eta$ such that $\|\mathbf{X}_a - \mathbf{X}_{\text{sup}}\|_2 \leq 1/4$. Then:

$$\begin{aligned} \|\mathcal{A}(\mathbf{X}_{\text{sup}})\|_2 &\leq \|\mathcal{A}(\mathbf{X}_a)\|_2 + \|\mathcal{A}(\mathbf{X}_{\text{sup}} - \mathbf{X}_a)\|_2 \\ &\leq (1 + \varepsilon)m + \frac{1}{4} \|\mathcal{A}(\mathbf{X}_{\text{sup}})\|_2. \end{aligned}$$

Hence, $\|\mathcal{A}(\mathbf{X}_{\text{sup}})\|_2 \leq \frac{4}{3}(1 + \varepsilon)m$. Similarly, let $\mathbf{X}_{\text{inf}} = \arg \inf_{\mathbf{X} \in \bar{\mathcal{D}}} \|\mathcal{A}(\mathbf{X})\|_2$. Choose $1 \leq b \leq \eta$ satisfying $\|\mathbf{X}_b - \mathbf{X}_{\text{inf}}\|_2 \leq 1/4$. Then,

$$\begin{aligned} \|\mathcal{A}(\mathbf{X}_{\text{inf}})\|_2 &\geq \|\mathcal{A}(\mathbf{X}_b)\|_2 - \|\mathcal{A}(\mathbf{X}_{\text{inf}} - \mathbf{X}_b)\|_2 \\ &\geq (1 - \varepsilon)m - \frac{1}{3}(1 + \varepsilon)m. \end{aligned}$$

This yields $\|\mathcal{A}(\mathbf{X}_{\text{inf}})\|_2 \geq \frac{2-4\varepsilon}{3}m$. Choosing $\varepsilon = 1/4$ whenever $m \geq \frac{32}{c} \log(\eta)$ with the desired probability, $\|\mathcal{A}(\mathbf{X}_{\text{inf}})\|_2 > 0$. Equivalently, $\bar{\mathcal{D}} \cap \text{Null}(\mathcal{A}) = \emptyset$. Since $\mathcal{A}(\cdot)$ is linear and \mathcal{D} is a cone, the claim is proved. ■

The following lemma gives a covering number of the set of low rank matrices.

Lemma 15 (Candes and Plan, [11]): Let M be the set of matrices in $\mathbb{R}^{d_1 \times d_2}$ with rank at most r . Then, for any $\varepsilon > 0$, there exists a covering of $\mathcal{S}^{d_1 \times d_2} \cap M$ with size at most $(\frac{c_3}{\varepsilon})^{(d_1+d_2)r}$ where c_3 is an absolute constant. In particular, $\log(\eta(1/4))$ is upper bounded by $C^{(d_1+d_2)r}$ for some constant $C > 0$.

Now, we use Lemma 15 to find the covering number of the set of simultaneously low rank and sparse matrices.

A. Proof of Lemma 7

Proof: Assume M has $\frac{1}{4}$ -covering number N . Then, using Lemma 14, whenever $m \geq c_1 \log N$, (18) will hold. What remains is to find N . To do this, we cover each individual $s_1 \times s_2$ submatrix and then take the union of the covers. For a fixed submatrix, using Lemma 15, $\frac{1}{4}$ -covering number is given by $C^{(s_1+s_2)q}$. In total there are $\binom{d_1}{s_1} \times \binom{d_2}{s_2}$ distinct submatrices. Consequently, by using $\log \binom{d}{s} \approx s \log \frac{d}{s} + s$, we find

$$\begin{aligned} \log N &\leq \log \left(\binom{d_1}{s_1} \times \binom{d_2}{s_2} C^{(s_1+s_2)q} \right) \\ &\leq s_1 \log \frac{d_1}{s_1} + s_1 + s_2 \log \frac{d_2}{s_2} + s_2 + (s_1 + s_2)q \log C, \end{aligned}$$

and obtain the desired result. ■

REFERENCES

- [1] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.
- [2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [3] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [4] Y. C. Eldar and G. Kutyniok, Eds., *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2012.
- [5] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010.
- [6] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Found. Comput. Math.*, vol. 9, no. 6, pp. 717–772, 2009.
- [7] V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky, "Latent variable graphical model selection via convex optimization," in *Proc. 48th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Sep./Oct. 2010, pp. 1610–1613.
- [8] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *J. ACM*, vol. 58, no. 3, 2011, Art. ID 11.
- [9] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Found. Comput. Math.*, vol. 12, no. 6, pp. 805–849, 2012.
- [10] S. Negahban, B. Yu, M. J. Wainwright, and P. K. Ravikumar, "A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers," in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 1348–1356.
- [11] E. J. Candès and Y. Plan, "Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements," *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 2342–2359, Apr. 2011.
- [12] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [13] J. Wright, A. Ganesh, K. Min, and Y. Ma, "Compressive principal component pursuit," *Inf. Inference*, vol. 2, no. 1, pp. 32–68, 2013.
- [14] S. Gandy, B. Recht, and I. Yamada, "Tensor completion and low-n-rank tensor recovery via convex optimization," *Inverse Problems*, vol. 27, no. 2, p. 025010, 2011.
- [15] L. Grasedyck, D. Kressner, and C. Tobler, "A literature survey of low-rank tensor approximation techniques," *GAMM-Mitteilungen*, vol. 36, no. 1, pp. 53–78, 2013.
- [16] A. Beck and Y. C. Eldar, "Sparsity constrained nonlinear optimization: Optimality conditions and algorithms," *SIAM J. Optim.*, vol. 23, no. 3, pp. 1480–1509, 2013.
- [17] Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," *Opt. Exp.*, vol. 19, no. 16, pp. 14807–14822, 2011.
- [18] A. Szameit *et al.*, "Sparsity-based single-shot subwavelength coherent diffractive imaging," *Nature Mater.*, vol. 11, no. 5, pp. 455–459, 2012.
- [19] A. Walther, "The question of phase retrieval in optics," *J. Modern Opt.*, vol. 10, no. 1, pp. 41–49, 1963.
- [20] R. P. Millane, "Phase retrieval in crystallography and optics," *J. Opt. Soc. Amer. A*, vol. 7, no. 3, pp. 394–411, 1990.
- [21] R. W. Harrison, "Phase problem in crystallography," *J. Opt. Soc. Amer. A*, vol. 10, no. 5, pp. 1046–1055, 1993.
- [22] N. E. Hurt, *Phase Retrieval and Zero Crossings: Mathematical Methods in Image Reconstruction*, vol. 52. New York, NY, USA: Springer-Verlag, 2001.
- [23] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of the phase from image and diffraction plane pictures," *Optik*, vol. 35, no. 2, pp. 237–246, 1972.
- [24] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, 1982.
- [25] H. Ohlsson, A. Y. Yang, R. Dong, and S. S. Sastry. (2011). "Compressive phase retrieval from squared output measurements via semidefinite programming." [Online]. Available: <http://arxiv.org/abs/1111.6323>
- [26] X. Li and V. Voroninski, "Sparse signal recovery from quadratic measurements via convex programming," *SIAM J. Math. Anal.*, vol. 45, no. 5, pp. 3019–3033, 2013.
- [27] K. Jaganathan, S. Oymak, and B. Hassibi, "Recovery of sparse 1-D signals from the magnitudes of their Fourier transform," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2012, pp. 1473–1477.
- [28] Y. M. Lu and M. Vetterli, "Sparse spectral factorization: Unicity and reconstruction algorithms," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2011, pp. 5976–5979.
- [29] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPARG: Efficient phase retrieval of sparse signals," *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 928–938, Feb. 2014.
- [30] Y. C. Eldar and S. Mendelson, "Phase retrieval: Stability and recovery guarantees," *Appl. Comput. Harmon. Anal.*, vol. 36, no. 3, pp. 473–494, 2014.
- [31] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM J. Imag. Sci.*, vol. 6, no. 1, pp. 199–225, 2013.
- [32] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [33] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight, "Sparsity and smoothness via the fused lasso," *J. Roy. Statist. Soc., Ser. B (Statistical Methodology)*, vol. 67, no. 1, pp. 91–108, 2005.
- [34] D. Needell and R. Ward, "Stable image reconstruction using total variation minimization," *SIAM J. Imag. Sci.*, vol. 6, no. 2, pp. 1035–1058, 2013.
- [35] D. Needell and R. Ward, "Near-optimal compressed sensing guarantees for total variation minimization," *IEEE Trans. Image Process.*, vol. 22, no. 10, pp. 3941–3949, Oct. 2013.
- [36] O. Semerci, N. Hao, M. E. Kilmer, and E. L. Miller. (2013). "Tensor-based formulation and nuclear norm regularization for multi-energy computed tomography." [Online]. Available: <http://arxiv.org/abs/1307.5348>
- [37] M. Golbabaee and P. Vandergheynst, "Joint trace/TV norm minimization: A new efficient approach for spectral compressive imaging," in *Proc. 19th IEEE Int. Conf. Image Process. (ICIP)*, Sep./Oct. 2012, pp. 933–936.
- [38] L. R. Tucker, "Some mathematical notes on three-mode factor analysis," *Psychometrika*, vol. 31, no. 3, pp. 279–311, 1966.
- [39] J. Liu, P. Musialski, P. Wonka, and J. Ye, "Tensor completion for estimating missing values in visual data," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 208–220, Jan. 2013.
- [40] C. Mu, B. Huang, J. Wright, and D. Goldfarb. (2013). "Square deal: Lower bounds and improved relaxations for tensor recovery." [Online]. Available: <http://arxiv.org/abs/1307.5870>
- [41] P. Sprechmann, I. Ramirez, G. Sapiro, and Y. C. Eldar, "C-HiLasso: A collaborative hierarchical sparse modeling framework," *IEEE Trans. Signal Process.*, vol. 59, no. 9, pp. 4183–4198, Sep. 2011.
- [42] M. Golbabaee and P. Vandergheynst, "Hyperspectral image compressed sensing via low-rank and joint-sparse matrix recovery," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2012, pp. 2741–2744.
- [43] M. Stojnic, F. Parvaresh, and B. Hassibi, "On the reconstruction of block-sparse signals with an optimal number of measurements," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 3075–3085, Aug. 2009.
- [44] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. Roy. Statist. Soc., Ser. B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [45] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 5302–5316, Nov. 2009.
- [46] G. A. Watson, "Characterization of the subdifferential of some matrix norms," *Linear Algebra Appl.*, vol. 170, pp. 33–45, Jun. 1992.

- [47] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [48] E. J. Candès and Y. Plan, "A probabilistic and RIPless theory of compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7235–7254, Nov. 2011.
- [49] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp, "Living on the edge: Phase transitions in convex programs with random data," *Inf. Inference*, 2014, doi: 10.1093/imaia/iau005.
- [50] R. Vershynin. (2014). "Estimation in high dimensions: A geometric perspective." [Online]. Available: <http://arxiv.org/abs/1405.5103>
- [51] R. Vershynin. (2010). "Introduction to the non-asymptotic analysis of random matrices." [Online]. Available: <http://arxiv.org/abs/1011.3027>
- [52] J. A. Tropp, "On the conditioning of random subictionaries," *Appl. Comput. Harmon. Anal.*, vol. 25, no. 1, pp. 1–24, 2008.
- [53] Z. D. Bai and Y. Q. Yin, "Limit of the smallest eigenvalue of a large dimensional sample covariance matrix," *Ann. Probab.*, vol. 21, no. 3, pp. 1275–1294, 1993.
- [54] D. L. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," *Proc. Nat. Acad. Sci. United States Amer.*, vol. 102, no. 27, pp. 9446–9451, 2005.
- [55] R. Foygel and L. Mackey, "Corrupted sensing: Novel guarantees for separating structured signals," *IEEE Trans. Inf. Theory*, vol. 60, no. 2, pp. 1223–1247, Feb. 2014.
- [56] S. Oymak, C. Thrampoulidis, and B. Hassibi. (2013). "The squared-error of generalized LASSO: A precise analysis." [Online]. Available: <http://arxiv.org/abs/1311.0830>
- [57] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar, *Convex Analysis and Optimization*. Belmont, MA, USA: Athena Scientific, 2003.
- [58] R. T. Rockafellar, R. J.-B. Wets, and M. Wets, *Variational Analysis*, vol. 317. New York, NY, USA: Springer-Verlag, 1998.
- [59] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ, USA: Princeton Univ. Press, 1997.
- [60] E. Candès and B. Recht, "Simple bounds for recovering low-complexity models," *Math. Program.*, vol. 141, nos. 1–2, pp. 577–589, 2013.
- [61] M. Grant, S. Boyd, and Y. Ye. (2008). "CVX: MATLAB software for disciplined convex programming." [Online]. Available: <http://cvxr.com/cvx/download/>
- [62] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optim. Methods Softw.*, vol. 11, nos. 1–4, pp. 625–653, 1999.
- [63] E. Richard, G. Obozinski, and J.-P. Vert. (2014). "Tight convex relaxations for sparse matrix factorization." [Online]. Available: <http://arxiv.org/abs/1407.5158>
- [64] S. Oymak and B. Hassibi. (2012). "Recovering jointly sparse signals via joint basis pursuit." [Online]. Available: <http://arxiv.org/abs/1202.3531>
- [65] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre, "Generalized power method for sparse principal component analysis," *J. Mach. Learn. Res.*, vol. 11, pp. 517–553, Mar. 2010.
- [66] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis," *J. Comput. Graph. Statist.*, vol. 15, no. 2, pp. 265–286, 2006.
- [67] A. d'Aspremont, F. Bach, and L. El Ghaoui, "Optimal solutions for sparse principal component analysis," *J. Mach. Learn. Res.*, vol. 9, pp. 1269–1294, Jun. 2008.
- [68] J.-J. Moreau, "Décomposition orthogonale d'un espace hilbertien selon deux cônes mutuellement polaires," *CR Acad. Sci. Paris*, vol. 255, pp. 238–240, 1962.
- [69] Y. Gordon, *On Milman's Inequality and Random Subspaces Which Escape Through a Mesh in \mathbb{R}^n* . New York, NY, USA: Springer-Verlag, 1988.

Samet Oymak (S'11) is a postdoctoral scholar in AMPLab at UC Berkeley. He received his B.Sc. degree from Bilkent University in 2009, M.Sc. and Ph.D. degrees from Caltech in 2011 and 2014 respectively, all in electrical engineering. At Caltech, he was awarded Atwood Fellowship in 2009 and he will be a Simons Fellow in UC Berkeley during Spring 2015. His current research interests are high dimensional estimation, statistical machine learning, convex optimization, and random matrix theory.

Amin Jalali (S'11) was born in Mashhad, Iran, and received the B.Sc. degrees in electrical engineering and in pure mathematics both from Sharif University of Technology, Tehran, Iran, in 2010. He is currently working toward the Ph.D. degree in Electrical Engineering at University of Washington, Seattle, WA, USA. His research interests include machine learning, convex optimization, and statistical signal processing, with a current emphasis on high-dimensional statistics and structure learning. He has been awarded University of Washington Graduate School Top Scholar Fellowship and Paul C. Leach Fellowship. Mr. Jalali is a student member of the Association for Computing Machinery and the American Statistical Association.

Maryam Fazel is an Associate Professor of Electrical Engineering at the University of Washington, with adjunct appointments in the departments of Computer Science and Engineering, Mathematics, and Statistics. Maryam received her MS and PhD in EE from Stanford University, her BS in EE from Sharif University of Technology in Iran, and was a postdoctoral scholar at the Control and Dynamical Systems Department at Caltech, prior to joining UW in 2008. She is a recipient of the NSF Career Award (2009), the UW EE Outstanding Teaching Award (2009), and the Best Student Paper award at the Uncertainty in Artificial Intelligence (UAI) conference (2014) with her student and collaborator. She coauthored a paper selected as a Fast-Breaking paper by ScienceWatch (2011), and was a plenary speaker at the Intl. Linear Algebra Society conference (ILAS, 2013).

Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering both from Tel-Aviv University (TAU), Tel-Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2002. From January 2002 to July 2002, she was a Postdoctoral Fellow at the Digital Signal Processing Group at MIT. She is currently a Professor in the Department of Electrical Engineering at the Technion - Israel Institute of Technology, Haifa and holds the The Edwards Chair in Engineering. She is also a Research Affiliate with the Research Laboratory of Electronics at MIT and a Visiting Professor at Stanford University, Stanford, CA. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, optimization methods, and their applications to biology and optics.

Dr. Eldar was a Horev Fellow of the Leaders in Science and Technology program at the Technion from 2002–2005 and an Alon Fellow. In 2004, she was awarded the Wolf Foundation Krill Prize for Excellence in Scientific Research, in 2005 the Andre and Bella Meyer Lectureship, in 2007 the Henry Taub Prize for Excellence in Research, in 2008 the Hershel Rich Innovation Award, the Award for Women with Distinguished Contributions, the Muriel & David Jacknow Award for Excellence in Teaching, and the Technion Outstanding Lecture Award, in 2009 the Technions Award for Excellence in Teaching, in 2010 the Michael Bruno Memorial Award from the Rothschild Foundation, and in 2011 the Weizmann Prize for Exact Sciences. In 2012 she was elected to the Young Israel Academy of Science and to the Israel Committee for Higher Education. In 2013 she received the Technions Award for Excellence in Teaching, the Hershel Rich Innovation Award, and the IEEE Signal Processing Technical Achievement Award, and in 2014 the IEEE/AESS Fred Nathanson Memorial Radar Award. She received several best paper awards together with her research students and colleagues. She is the Editor in Chief of *Foundations and Trends in Signal Processing* and a member of the IEEE Sensor Array and Multichannel Technical Committee. In the past, she was a Signal Processing Society Distinguished Lecturer, member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing technical committees, and served as an associate editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*.

Babak Hassibi (M'08) was born in Tehran, Iran, in 1967. He received the B.S. degree from the University of Tehran in 1989, and the M.S. and Ph.D. degrees from Stanford University in 1993 and 1996, respectively, all in electrical engineering.

He has been with the California Institute of Technology since January 2001, where he is currently Gordon M Binder/Amgen Professor and Executive Officer of Electrical Engineering, and Associate Director of Information Science and Technology. From October 1996 to October 1998 he was a research associate at the Information Systems Laboratory, Stanford University, and from November 1998 to December 2000 he was a Member of the Technical Staff in the Mathematical Sciences Research Center at Bell Laboratories, Murray Hill, NJ. He has also held short-term appointments at Ricoh California Research Center, the Indian Institute of Science, and Linköping University, Sweden. His research interests include wireless communications and networks, robust estimation and control, adaptive signal processing and linear algebra. He is the coauthor of the books (both with A.H. Sayed and T. Kailath) *Indefinite Quadratic Estimation and Control: A Unified Approach to H₂ and H₁ Theories* (New York: SIAM, 1999)

and *Linear Estimation* (Englewood Cliffs, NJ: Prentice Hall, 2000). He is a recipient of an Alborz Foundation Fellowship, the 1999 O. Hugo Schuck best paper award of the American Automatic Control Council (with H. Hindi and S.P. Boyd), the 2002 National Science Foundation Career Award, the 2002 Okawa Foundation Research Grant for Information and Telecommunications, the 2003 David and Lucille Packard Fellowship for Science and Engineering, the 2003 Presidential Early Career Award for Scientists and Engineers (PECASE), and the 2009 Al-Marai Award for Innovative Research in Communications, and was a participant in the 2004 National Academy of Engineering "Frontiers in Engineering" program.

He has been a Guest Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY special issue on "space-time transmission, reception, coding and signal processing" was an Associate Editor for Communications of the IEEE TRANSACTIONS ON INFORMATION THEORY during 2004-2006, and is currently an Editor for the Journal "*Foundations and Trends in Information and Communication*" and for the IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING.