

Harnessing Structures in Big Data via Guaranteed Low-Rank Matrix Estimation

Yudong Chen, *Member, IEEE*, Yuejie Chi, *Senior Member, IEEE*

Abstract—Low-rank modeling plays a pivotal role in signal processing and machine learning, with applications ranging from collaborative filtering, video surveillance, medical imaging, to dimensionality reduction and adaptive filtering. Many modern high-dimensional data and interactions thereof can be modeled as lying approximately in a low-dimensional subspace or manifold, possibly with additional structures, and its proper exploitations lead to significant reduction of costs in sensing, computation and storage. In recent years, there is a plethora of progress in understanding how to exploit low-rank structures using computationally efficient procedures in a provable manner, including both convex and nonconvex approaches. On one side, convex relaxations such as nuclear norm minimization often lead to statistically optimal procedures for estimating low-rank matrices, where first-order methods are developed to address the computational challenges; on the other side, there is emerging evidence that properly designed nonconvex procedures, such as projected gradient descent, often provide globally optimal solutions with a much lower computational cost in many problems. This survey article will provide a unified overview of these recent advances on low-rank matrix estimation from incomplete measurements. Attention is paid to rigorous characterization of the performance of these algorithms, and to problems where the low-rank matrix have additional structural properties that require new algorithmic designs and theoretical analysis.

Index Terms—low-rank matrix estimation, convex relaxation, non-convex matrix factorization, structured matrices, incomplete observations

I. INTRODUCTION

The ubiquity of advanced sensing and imaging technologies produce vast amounts of data at an unprecedented rate. A fundamental goal of signal processing is to extract, and possibly track the evolution of, the relevant structural information faithfully from such high-dimensional data, ideally with a minimal amount of computation, storage and human intervention. To overcome the curse of dimensionality, it is important to exploit the fact that real-world data often possess some *low-dimensional geometric structures*. In particular, such structures allow for a succinct description of the data by a number of parameters much smaller than the ambient dimension. One popular postulate of low-dimensional structures is *sparsity*, that is, a signal can be represented using a few nonzero coefficients in a proper domain. For instance, a natural image often has a sparse representation in the wavelet domain. The field of *compressed sensing* [1], [2] has made tremendous progress in capitalizing on the sparsity structures, particularly in solving

under-determined linear systems arising from *sample-starved* applications such as medical imaging, spectrum sensing and network monitoring. In these applications, compressed sensing techniques allow for faithful estimation of the signal of interest from a number of measurements that is proportional to the sparsity level — much fewer than is required by traditional techniques. The power of compressed sensing has made it a disruptive technology in many applications such as magnetic resonance imaging (MRI): a Cardiac Cine scan can now be performed within 25 seconds with the patients breathing freely. This is in sharp contrast to the previous status quo, where the scan takes up to six minutes and the patients need to hold their breaths several times [3].

While the sparsity model is powerful, the original framework of compressed sensing mainly focuses on vector-valued signals that admit sparse representations in an *a priori known* domain. However, knowledge of such sparsifying domains is not always available, thus limiting its applications. Fortunately, one can resort to a more general notion of sparsity that is more versatile when handling matrix-valued signals — or an ensemble of vector-valued signals — without the need of specifying a sparsifying basis. In this paper, we will review this powerful generalization of sparsity, termed the low-rank model, which captures a much broader class of low-dimensional structures. Roughly speaking, this model postulates that the matrix-valued signal is approximately low-rank. If we view each column of the matrix as a data vector, then this is equivalent to saying that the data approximately lies in a low-dimensional but unknown subspace. Historically, the exploitation of low-rank structures may begin even earlier than that of sparsity. In particular, the low-rank assumption is what underlies classical Principal Component Analysis (PCA) [4], which builds on the observation that real-world data has most of its variance in the first few top principal components. Such low-rank structures may arise due to various physical reasons and engineering designs. In face recognition, face images are found to trace out a 9-dimensional subspace if they are approximately convex and reflect light according to Lambert’s law [5]. In radar and sonar signal processing, the signals reside approximately in a low-dimensional subspace due to transmitting using a small set of waveforms to construct certain beam patterns. Low-rank structures also arise from modeling interactions between different objects. For example, in clustering or embedding, the pairwise interactions between objects can often be expressed as a low-rank matrix [6].

Given the collected data, the key problem is to infer the hidden low-dimensional subspace that captures most of the information relevant for subsequent tasks such as detection,

Authors are listed alphabetically.

Y. Chen is with the School of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA (email: yudong.chen@cornell.edu).

Y. Chi is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA (email: yuejie.chi@cmu.edu).

clustering, and parameter estimation. Traditional methods such as Singular Value Decomposition (SVD) for finding principal subspaces typically require the data to be fully observed. However, modern data applications often involve estimation problems with a number of measurements that is much smaller than the ambient dimension, a regime similar to the setting of compressed sensing. We refer to this problem as *low-rank matrix estimation*, emphasizing the fact that one only has under-sampled measurements or partial observations. Examples of such problems are abundant. In recommendation systems, the goal is to estimate the missing ratings given a small number of observed ones. In sensor networks, an important problem is to infer the locations of the sensors from pairwise distance measures, which are available only for sensors within a certain radius of each other. In wideband spectrum sensing, to reduce the sampling rate, a popular approach is to estimate the signal subspace and bearing parameters by randomly sub-sampling the outputs of the array.

In these applications, it is desirable to develop low-rank matrix estimation algorithms that are both *statistically efficient* — achieving low estimation errors with a minimal amount of (noisy) measurements — and *computationally efficient* — having low running time and storage cost. A particular focus of this paper is on algorithms that come with *provable* guarantees for their statistical and computation efficiency. The search for such algorithms is in part motivated by the remarkable success story of compressed sensing, for which many provable methods have been developed for sparse models. Handling the more general low-rank structures poses a new set of challenges as well as opportunities. The study of low-rank matrix estimation has attracted the attention of many researchers from diverse communities including signal processing, machine learning, statistics, mathematical programming and computer science [7]–[11]. As we elaborate below, this enterprise has been much fruitful, resulting in many powerful algorithms, novel analytical techniques, and deep theoretical insights.

This survey article is complementary to the nice overview article on low-rank matrix recovery by Davenport and Romberg [12], with different focuses. In particular, by focusing on recent algorithmic advancements with computational and statistical guarantees, we highlight the effectiveness of first-order methods in both convex and nonconvex optimization. We also put specific emphasis on a distinct set of applications involving structured matrix completion.

A. Paper Organizations

The rest of this paper is organized as follows. Section II motivates low-rank models from the perspectives of modeling data correlations and lifting vector-valued problems. Section III describes the basic mathematical setup of the low-rank estimation problem. Section IV discusses the theory and algorithms for low-rank matrix estimation via convex optimization. Section V discusses the theory and algorithms for low-rank matrix estimation via nonconvex optimization. Section VI discusses structured matrix completion, where the low-rank matrices have additional structural constraints, using several concrete examples. Numerical examples on a real-world

recommendation dataset are showcased in Section VII. The paper is concluded in Section VIII.

B. Notations

Throughout this paper, we use boldface capital letters such as \mathbf{A} to denote matrices, with \mathbf{A}^T being its transpose and A_{ij} being its (i, j) -th entry. Similarly, we use boldface lower-case letters such as \mathbf{a} to denote vectors, with \mathbf{a}^* being its conjugate transpose and a_i being its i -th entry. The expectation is denoted by \mathbb{E} . In addition, $\|\mathbf{A}\|$, $\|\mathbf{A}\|_F$, $\|\mathbf{A}\|_{2,\infty}$, $\text{Tr}(\mathbf{A})$, and $\|\mathbf{A}\|_*$, stand for the spectral norm (i.e. the largest singular value), the Frobenius norm, the ℓ_2/ℓ_∞ norm (i.e. the largest ℓ_2 norm of the rows), the trace, and the nuclear norm (i.e. the sum of singular values) of the matrix \mathbf{A} . For two matrices \mathbf{A} and \mathbf{B} of the same size, $\langle \mathbf{A}, \mathbf{B} \rangle \triangleq \text{Tr}(\mathbf{A}^T \mathbf{B})$ denotes their trace inner product. The notation $\text{diag}[\mathbf{c}]$ denotes a diagonal matrix whose diagonal entries are given by the vector \mathbf{c} . We use \mathbf{e}_i to denote the i -th standard basis vector of \mathbb{R}^n , for each $i = 1, 2, \dots, n$.

II. THE UBIQUITY OF LOW-RANK MODELS

In this section, we elucidate the motivations for studying low-rank modeling and low-rank matrix estimation problems. We start with a classical viewpoint and justify the low-rank priors of a data matrix from the perspective of bias-variance trade-offs for modeling correlations in data observations. We next argue that low-rank structures arise from a powerful reformulation of quadratic optimization problems by lifting them into a matrix space. Last but not least, we provide a list of other sources of low-rank structures in a wide range of science and engineering problems.

A. Correlation-Aware Modeling of Data Matrices

We first motivate the use of low-rank models as a general principle for bias-variance trade-off in signal estimation and processing given noisy data, which is a classical viewpoint articulated by Scharf and Tufts in [13]. Consider a stationary signal $\mathbf{x} \in \mathbb{R}^n$ with covariance matrix $\mathbf{\Sigma} = \mathbb{E}[\mathbf{x}\mathbf{x}^T] \in \mathbb{R}^{n \times n}$. Suppose that the eigenvalue decomposition of $\mathbf{\Sigma}$ is given by $\mathbf{\Sigma} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, where $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n] \in \mathbb{R}^{n \times n}$ are the eigenvectors, and $\mathbf{\Lambda} = \text{diag}[\lambda_1, \dots, \lambda_n]$ are the eigenvalues arranged as a non-increasing sequence. Define the matrix $\mathbf{P}_r \triangleq \mathbf{U}_r \mathbf{U}_r^T$, where $\mathbf{U}_r = [\mathbf{u}_1, \dots, \mathbf{u}_r] \in \mathbb{R}^{n \times r}$ is the matrix whose columns are the eigenvectors corresponding to the r largest eigenvalues. One sees that \mathbf{P}_r is the projection operator onto the subspace spanned by the top- r eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_r$.

Imagine that we receive a noisy copy of the signal \mathbf{x} as

$$\mathbf{y} = \mathbf{x} + \mathbf{w},$$

where the elements of the noise vector $\mathbf{w} \in \mathbb{R}^n$ are independent with variance σ^2 . Let us estimate the signal \mathbf{x} using a reduced-rank model with rank r by projecting observed data onto the r -dimensional principal subspace of $\mathbf{\Sigma}$. Such an estimate is given by

$$\hat{\mathbf{x}}_r = \mathbf{P}_r \mathbf{y} = \mathbf{P}_r \mathbf{x} + \mathbf{P}_r \mathbf{w}.$$

We make the crucial observation that one may decompose the mean squared error of the estimate $\hat{\mathbf{x}}_r$ into two terms:

$$\frac{1}{n} \mathbb{E} \|\hat{\mathbf{x}}_r - \mathbf{x}\|_2^2 = \frac{1}{n} \mathbb{E} \|\mathbf{x} - \mathbf{P}_r \mathbf{x}\|_2^2 + \frac{1}{n} \mathbb{E} \|\mathbf{P}_r \mathbf{w}\|_2^2. \quad (1)$$

Here the first term corresponds to the model *bias*,

$$b_r^2 = \frac{1}{n} \mathbb{E} \|\mathbf{x} - \mathbf{P}_r \mathbf{x}\|_2^2 = \frac{1}{n} \left(\sum_{i=r+1}^n \lambda_i^2 \right),$$

which arises due to approximating a full-rank signal by a rank- r model. The second term corresponds to the model *variance*,

$$v_r^2 = \frac{r}{n} \sigma^2,$$

which arises due to the presence of noise.

From the above decomposition, we see that as one increases the rank r , the bias of the estimate decreases, whereas the corresponding variance increases. Therefore, the choice of the rank controls the trade-off between the bias and the variance, whose sum constitutes the total estimation error. Importantly, many real-world datasets have a decaying spectrum, which in the above notation means that the eigenvalues of the covariance matrix decrease rapidly. As mentioned, this insight is the foundation of PCA [4] and moreover is observed across a wide range of applications including power systems, recommendation systems, Internet traffic and weather data. Consequently, it is beneficial to employ a small rank, so that the variance is controlled, while the bias remains small as long as the residual eigenvalues decreases quickly. With this in mind, we show in Fig. 1 the mean squared error as function of the rank, as well as the decomposition into the bias and the variance; here it is assumed that the spectrum decays at a rate of $\lambda_i = 1/i$, and the signal-to-noise ratio (SNR), defined as $\sum_{i=1}^n \lambda_i^2 / (n\sigma^2)$, is equal to 15dB with $n = 100$. It is clear from the figure that employing an appropriate low-rank estimator induces a much lower mean squared error than a full-rank one. In particular, the optimal rank may be much smaller than the ambient dimension n when the spectrum decays fast.

B. Lifting for Quadratic and Bilinear Optimization Problems

Another important source of low-rank structures is solving quadratic/bilinear optimization problems. As an example, consider the *phase retrieval* problem [14], an important routine in X-ray crystallography and optical imaging, where the goal is to recover a vector \mathbf{x} in \mathbb{C}^n or \mathbb{R}^n given only the *magnitudes* of its linear measurements, that is,

$$y_l = |\langle \mathbf{a}_l, \mathbf{x} \rangle|^2 = \langle \mathbf{a}_l \mathbf{a}_l^*, \mathbf{x} \mathbf{x}^* \rangle. \quad (2)$$

Due to the nonlinear nature of these equations, it is difficult to solve them directly, particularly when the problem size is large. A popular approach for solving such equations is called *lifting*: one rewrites the above equations in terms of the matrix variable $\mathbf{M} = \mathbf{x} \mathbf{x}^*$, and casts this problem as recovering the rank-one matrix \mathbf{M} from a set of linear measurements [15]. A similar formulation has been used for the blind deconvolution problem; cf. [16].

The lifting approach can be applied to other classes of quadratic equations, whose lifting formulations may lead to




Fig. 1. The mean squared error and its decomposition into bias and variance of signal estimation in additive Gaussian noise by assuming a low-rank model at a moderate SNR. This suggests that it is beneficial to apply a reduced-rank model when the data are correlated.

low-rank matrices of rank larger than one. For instance, in the problem of sensor network localization [17], the goal is to determine the locations of a set of n points/sensors $\{\mathbf{x}_i\}_{i=1}^n$ lying in an r -dimensional Euclidean space, where $r \ll n$, given a subset of their pairwise distances. The complete set of pairwise Euclidean distances can be arranged as a matrix $\mathbf{E} = [E_{ij}] \in \mathbb{R}^{n \times n}$, where $E_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|_2^2$, $1 \leq i, j \leq n$. Interestingly, each pairwise distance (that is, an entry of \mathbf{E}) is in fact a linear function of the rank- r , positive semidefinite (PSD) matrix $\mathbf{M} = \mathbf{X} \mathbf{X}^T$, where $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]^T \in \mathbb{R}^{n \times r}$; more precisely, one has

$$E_{ij} = \|\mathbf{X}^T (\mathbf{e}_i - \mathbf{e}_j)\|_2^2 = (\mathbf{e}_i - \mathbf{e}_j)^T \mathbf{M} (\mathbf{e}_i - \mathbf{e}_j). \quad (3)$$

Therefore, the problem of determining the locations \mathbf{X} of the sensors is equivalent to recovering the low-rank lifted matrix \mathbf{M} from a set of linear measurements in the form of (3); see [17] for a more detailed treatment of this powerful reformulation.

C. Other Sources of Low-rank Structures

There're many potential sources of low-rank structures. Below we provide a few further examples drawn from different science and engineering domains:

- In system identification and time series analysis, finding the minimum-order linear time-invariant system is equivalent to minimizing the rank of Hankel structured matrices [18] (cf. Section VI-A).
- In recommendation systems [19], the matrix of user ratings for a set of items is often approximately low-rank, as user preferences typically depend on a small number of underlying factors and hence their ratings correlate with each other.
- The background of a video usually changes slowly from frame to frame, hence stacking the frames as columns lead to an approximately low-rank matrix [20]. Similar low-rank structures arise from the smoothness properties of other visual and physical objects [5].

- In quantum state tomography, the density matrix of a pure or nearly pure quantum state is approximately low-rank, which can be exploited in the problem of state reconstruction from a small number of Pauli measurements [21].
- In a sparse graphical model with latent variables, one can show, using the Schur complement, that the inverse marginal covariance matrix of the observed variables can be approximated by a matrix with rank equal to the number of latent variables [22].
- Matrices with certain monotonicity properties can be well-approximated by a matrix with rank much smaller than the ambient dimension. Such matrices arise, for example, when measuring the pairwise comparison scores of a set of objects that possess an underlying ordering [23].
- The pairwise affinity matrix of a set of objects is often approximately low-rank due to the presence of clustering/community structures [24] (cf. Section VI-B).

The list continues for much longer. The ubiquity of these structures, either as a physical property or as an engineering choice, is what makes low-rank models useful, and motivates the extensive study of the low-rank matrix estimation problem.

III. LOW-RANK MATRIX ESTIMATION FROM INCOMPLETE OBSERVATIONS

In this section, we formally define the problem of *low-rank matrix estimation*, that is, recovery of a low-rank matrix from a number of measurements much smaller than the dimension of the matrix. Let $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ be the matrix-valued signal of interest.¹ Denote the SVD of \mathbf{X} by

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_{i=1}^{\min\{n_1, n_2\}} \sigma_i \mathbf{u}_i \mathbf{v}_i^T,$$

where the singular values $\sigma_1 \geq \sigma_2 \geq \dots$ are organized in an non-increasing order. The best rank- r approximation of \mathbf{X} is defined as

$$\mathbf{X}_r \triangleq \underset{\text{rank}(\mathbf{G}) \leq r}{\text{argmin}} \|\mathbf{X} - \mathbf{G}\|_F.$$

By the Eckart-Young theorem, the optimal approximation \mathbf{X}_r is given by

$$\mathbf{X}_r = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T. \quad (4)$$

Correspondingly, the rank- r approximation error is given by $\|\mathbf{X} - \mathbf{X}_r\|_F$, and we say that the matrix \mathbf{X} is approximately low-rank if its rank- r approximation error is small for some $r \ll \min\{n_1, n_2\}$.

As mentioned, in many modern applications, one does not directly observe \mathbf{X} , but rather is given an under-determined set of indirect noisy measurements of it. Here we assume that one has access to a set of linear measurements in the form

$$y_l = \langle \mathbf{A}_l, \mathbf{X} \rangle + w_l, \quad l = 1, \dots, m, \quad (5)$$

where $\mathbf{A}_l \in \mathbb{R}^{n_1 \times n_2}$ is the l -th measurement matrix, and $w_l \in \mathbb{R}$ is a noise term. We may rewrite these equations more compactly in a matrix form as

$$\mathbf{y} = \mathcal{A}(\mathbf{X}) + \mathbf{w}, \quad (6)$$

where $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^m$ is the linear measurement operator defined by $[\mathcal{A}(\mathbf{X})]_l = \langle \mathbf{A}_l, \mathbf{X} \rangle$, and $\mathbf{w} = [w_1, \dots, w_m]^T \in \mathbb{R}^m$ is the vector of noise terms. Denote by \mathcal{A}^* the conjugate operator of \mathcal{A} , where $\mathcal{A}^*(\mathbf{y}) = \sum_{l=1}^m y_l \mathbf{A}_l$.

As we are primarily concerned with estimating \mathbf{X} from $m \ll n_1 n_2$ measurements, direct approximation via SVD and the Eckart-Young theorem are impossible. Instead, we need to develop alternative methods to find an (approximate) low-rank solution that best fits the set of noisy under-determined linear equations (6). We further categorize the low-rank matrix estimation problem into two main types based on the structure of the measurement operator:

- *Low-Rank Matrix Sensing*, one observes linear combinations of the entries of \mathbf{X} , where each measurement matrix \mathbf{A}_l defining the linear combinations is typically dense.
- *Low-Rank Matrix Completion*, one directly observes a subset of the entries of \mathbf{X} , and aims to interpolate the missing entries. In this case, each \mathbf{A}_l is a sparse matrix with a single entry equal to 1 at the corresponding observed index.

For matrix completion, it is convenient to write the measurements in a matrix form as

$$\mathbf{Y} = \mathcal{P}_\Omega(\mathbf{X}) + \mathbf{W}, \quad (7)$$

where $\Omega \subset \{1, 2, \dots, n_1\} \times \{1, 2, \dots, n_2\}$ is the collection of indices of the observed entries, $\mathcal{P}_\Omega : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^{n_1 \times n_2}$ is the entry-wise partial observation operator defined by

$$[\mathcal{P}_\Omega(\mathbf{X})]_{ij} = \begin{cases} X_{ij}, & (i, j) \in \Omega, \\ 0, & \text{otherwise,} \end{cases}$$

and $\mathbf{W} \in \mathbb{R}^{n_1 \times n_2}$ is the noise matrix supported on Ω . With this notation, matrix completion is the problem of (approximately) recovering \mathbf{X} given \mathbf{Y} and Ω .

IV. THEORY AND ALGORITHMS FOR LOW-RANK MATRIX ESTIMATION VIA CONVEX OPTIMIZATION

The development of efficient algorithms for low-rank estimation owes much of its inspiration to the success of compressed sensing [1], [2]. There, the convex relaxation approach based on ℓ_1 -minimization is widely used for recovering sparse signals. For low-rank problems, the role of the ℓ_1 norm is replaced by its matrix counterpart, namely the nuclear norm (also known as the trace norm), which is a convex surrogate for the rank. This idea gives rise to convex optimization approaches for low-rank estimation based on nuclear norm minimization, an approach put forth by Fazel et al. in the seminal work [25]. This approach has since been extensively developed and expanded, which remains the most mature and well-understood method (though not the only one) for estimating low-rank matrices. In this section we provide a survey of this algorithmic approach and the associated theoretical results.

¹Our discussions can be extended complex-valued matrices straightforwardly.

A. Convex Relaxation via Nuclear Norm minimization

We begin by deriving the nuclear norm minimization algorithm as a convex relaxation for rank minimization. Recall our linear measurement model in (6), to which we seek a low-rank solution. A natural approach is to find the matrix with the minimum rank that is consistent with these measurements, which can be formulated as an optimization problem:

$$\min_{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}} \text{rank}(\mathbf{X}) \quad \text{subject to} \quad \mathbf{y} = \mathcal{A}(\mathbf{X}). \quad (8)$$

The rank, however, is a non-convex function of \mathbf{X} , and rank minimization (8) is known to be NP-hard in general. To develop a tractable formulation, one observes that the rank of \mathbf{X} is equal to the number of its nonzero singular values. Therefore, analogously to using the ℓ_1 norm as a convex surrogate of sparsity, we may replace the rank of \mathbf{X} by the sum of its singular values, a quantity known as the nuclear norm:

$$\|\mathbf{X}\|_* \triangleq \sum_{i=1}^{\min\{n_1, n_2\}} \sigma_i.$$

Then, instead of solving (8) directly, one solves for a matrix that minimizes the nuclear norm:

$$\hat{\mathbf{X}} = \underset{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}}{\text{argmin}} \|\mathbf{X}\|_* \quad \text{s.t.} \quad \mathbf{y} = \mathcal{A}(\mathbf{X}). \quad (9)$$

In the case where the measurements \mathbf{y} are noisy, one seeks a matrix with a small nuclear norm that is approximately consistent with the measurements, which can be formulated either as a regularized optimization problem:

$$\hat{\mathbf{X}} = \underset{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}}{\text{argmin}} \frac{1}{2} \|\mathbf{y} - \mathcal{A}(\mathbf{X})\|_2^2 + \tau \|\mathbf{X}\|_*, \quad (10)$$

or as a constrained optimization problem:

$$\hat{\mathbf{X}} = \underset{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}}{\text{argmin}} \|\mathbf{y} - \mathcal{A}(\mathbf{X})\|_2^2 \quad \text{subject to} \quad \|\mathbf{X}\|_* \leq \gamma, \quad (11)$$

where τ and γ are tuning parameters. Since the nuclear norm can be represented as the solution to a semidefinite program [25],

$$\|\mathbf{X}\|_* = \min_{\mathbf{W}_1, \mathbf{W}_2} \frac{1}{2} (\text{Tr}(\mathbf{W}_1) + \text{Tr}(\mathbf{W}_2)) \quad (12)$$

$$\text{subject to} \quad \begin{bmatrix} \mathbf{W}_1 & \mathbf{X} \\ \mathbf{X}^T & \mathbf{W}_2 \end{bmatrix} \succeq 0.$$

Consequently, the optimization problems (9)–(11) are convex, semidefinite programs.

B. Guarantees for Matrix Sensing via RIP

For there to be any hope of recovering \mathbf{X} from the output of the sensing process (6), the sensing operator \mathcal{A} needs to possess certain desirable properties so that it can distinguish different low-rank matrices. One such property is called the *restricted isometry property* (RIP). RIP stipulates that \mathcal{A} , viewed as a mapping to a lower-dimensional space, preserves the Euclidean distances between low-rank matrices. Below we give a general notion of RIP, where the distances after mapping may be measured in different norms:

Definition 1 (Restricted Isometry Property): The operator \mathcal{A} is said to satisfy the RIP- ℓ_2/ℓ_p property of rank r if for all matrices Φ of rank at most r , it satisfies the inequality

$$(1 - \underline{\delta}_r) \|\Phi\|_F \leq \|\mathcal{A}(\Phi)\|_p \leq (1 + \bar{\delta}_r) \|\Phi\|_F,$$

where $\underline{\delta}_r$ and $\bar{\delta}_r$ are some universal constants satisfying $0 < 1 - \underline{\delta}_r < 1 < 1 + \bar{\delta}_r$.

This definition is reminiscent of a similar notion with the same name used in the sparse signal recovery literature that is imposed on sparse vectors [1]. Certifying whether RIP holds for a given operator is known to be NP-hard [26]. Nevertheless, it turns out that a “generic” sensing operator, drawn from certain random distributions, satisfies RIP with high probability. For example:

- If the measurement matrix \mathbf{A}_l has i.i.d. Gaussian entries $\mathcal{N}(0, 1/m)$, then \mathcal{A} satisfies RIP- ℓ_2/ℓ_2 with high probability as long as $m \geq c(n_1 + n_2)r$ for some large enough constant $c > 0$ [25], [27].
- If the measurement matrix $\mathbf{A}_l = \mathbf{a}_l \mathbf{b}_l^T$ is rank-one with $\mathbf{a}_l, \mathbf{b}_l$ composed with i.i.d. Gaussian entries $\mathcal{N}(0, 1/m)$, then \mathcal{A} satisfies RIP- ℓ_2/ℓ_1 with high probability as long as $m \geq c(n_1 + n_2)r$ for some large enough constant $c > 0$ [28], [29].

When RIP- ℓ_2/ℓ_p holds, the nuclear norm minimization approach guarantees exact and stable recovery of the low-rank matrix in both noise-free and noisy cases, as shown in the following theorem adapted from [27], [28], [29].

Theorem 1 (Recovery Guarantee via RIP): Suppose that the noise satisfies $\|\mathbf{w}\|_p \leq \epsilon$. If \mathcal{A} satisfies RIP- ℓ_2/ℓ_p ,² then the solution to the nuclear norm minimization algorithms (9)–(11) (with appropriate values for the tuning parameters) satisfies the error bound

$$\|\hat{\mathbf{X}} - \mathbf{X}\|_F \leq C_1 \frac{\|\mathbf{X} - \mathbf{X}_r\|_*}{\sqrt{r}} + C_2 \epsilon \quad (13)$$

simultaneously for all $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$, where C_1 and C_2 are positive numerical constants.

For studying the performance of nuclear norm minimization via other notions such as the null space property, we refer interested readers to [30].

C. Guarantees for Matrix Completion via Incoherence

In the case of matrix completion, an additional complication arises: it is impossible to recover a low-rank matrix that is also sparse. In particular, when one only samples a small subset of the entries of \mathbf{X} , it is very likely that most, if not all, of the nonzero entries of \mathbf{X} are missed. This means that the sensing operator $\mathcal{A} = \mathcal{P}_\Omega$ used for matrix completion cannot satisfy the RIP. Therefore, for the problem to be well-posed, we need to restrict attention to low-rank matrices whose mass does not concentrate on a few entries. This property can be formalized by the notion of *incoherence*, which measures the alignment

²When $p = 2$, we require $\bar{\delta}_{4r} < 0.1892$, $\underline{\delta}_{4r} < 0.2346$ [27]; when $p = 1$, we require there exists a universal constant $k \geq 2$ such that $(1 + \bar{\delta}_{kr})/(1 - \underline{\delta}_{kr}) < \sqrt{k}$ [28], [29]. Both requirements can be met with a sample complexity of $\mathcal{O}((n_1 + n_2)r)$.

between the column/row spaces of the low-rank matrix with the standard basis vectors:

Definition 2 (Incoherence): For a matrix $\mathbf{U} \in \mathbb{R}^{n \times r}$ with orthonormal columns, let \mathbf{P}_U be the orthogonal projection onto the column space of \mathbf{U} . The incoherence parameter of \mathbf{U} is defined as

$$\mu(\mathbf{U}) = \frac{n}{r} \max_{1 \leq i \leq n} \|\mathbf{P}_U \mathbf{e}_i\|_2^2. \quad (14)$$

For a matrix with the SVD $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, the incoherence parameter of \mathbf{X} is defined as

$$\mu_0 = \max\{\mu(\mathbf{U}), \mu(\mathbf{V})\}.$$

It is easy to see that the incoherence parameter satisfies the bound $1 \leq \mu(\mathbf{U}) \leq n/r$. With a smaller $\mu(\mathbf{U})$, the column space of \mathbf{U} is more spread out over its coordinates. For a matrix \mathbf{X} , its incoherence parameters μ_0 is determined by its singular vectors and is independent of its singular values. In the noiseless setting, nuclear norm minimization can perfectly recover an incoherent low-rank matrix as soon as the number of measurements is slightly larger than the degrees of freedom of the matrix. Such recovery guarantees were proved and refined in a series of work in [8], [10], [11], [31], [32], [33]. The theorem below is adapted from [31], which is state-of-the-art.

Theorem 2 (Recovery Guarantee via Incoherence): [31] Suppose that each entry of \mathbf{X} is observed independently with probability $p \in (0, 1)$. If p satisfies

$$p \geq C\mu_0 \frac{r \log^2(n_1 + n_2)}{(n_1 + n_2)},$$

for some constant C , then with high probability, the nuclear norm minimization algorithm (9) exactly recovers \mathbf{X} as the unique optimal solution.

By a coupon-collecting argument [8], one can in fact show that it is impossible to recover the matrix with less than $(n_1 + n_2)r \log(n_1 + n_2)$ measurements using any algorithm. Therefore, Theorem 2 shows that nuclear norm minimization is near-optimal in terms of sample complexity — off by only a logarithmic factor — a remarkable fact considering that we are using a convex relaxation of the rank.

In the noisy case, one can study the performance of nuclear norm minimization in terms of its recovery error, which is done for example in [9], [34]. Here we state one such performance guarantee taken from [9]. Let us assume that the entries of the noise \mathbf{w} are independent with variance that scaled as $\nu/\sqrt{n_1 n_2}$. Then, by solving the regularized nuclear norm minimization problem (10) with parameter $\tau = 4\nu\sqrt{\frac{(n_1+n_2)\log(n_1+n_2)}{m}}$, we obtain a solution $\hat{\mathbf{X}}$ satisfying

$$\|\hat{\mathbf{X}} - \mathbf{X}\|_F^2 \leq \nu^2 \frac{(n_1 + n_2)r \log(n_1 + n_2)}{m}$$

with high probability in the moderate to low SNR regime.

D. First-Order Algorithms for Nuclear Norm Minimization

In principle, it is possible to solve the nuclear norm minimization problems in (9)–(11) to high numerical accuracy using off-the-shelf semidefinite programming solvers (such

Algorithm 1 FISTA for Low-Rank Matrix Estimation

Parameters: T, L, τ in (10);

Initialization: Let $\mathbf{R}_1 = \mathbf{0}, t_1 = 1$;

Updates: for $k = 1, 2, \dots, T$ do

$$1) \mathbf{X}_k = \mathcal{D}_{\tau/L}(\mathbf{R}_k - \frac{1}{L}\mathcal{A}^*(\mathcal{A}(\mathbf{R}_k) - \mathbf{y}));$$

$$2) t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2};$$

$$3) \mathbf{R}_{k+1} = \mathbf{X}_k + \left(\frac{t_k - 1}{t_{k+1}}\right)(\mathbf{X}_k - \mathbf{X}_{k-1}).$$

Output: \mathbf{X}_T .

as SDPT3 [35]). However, these solvers, typically based on interior-point methods, can be extremely slow when the size of the matrix is large. For example, SDPT3 can only handle matrices with dimensions no larger than a few thousands due to memory requirements. This computational issue motivates the development of fast alternatives that can handle significantly larger problems. First-order algorithms become an appealing candidate due to their low per-iteration cost, as well as the flexibility to incorporate the specific structures of the semidefinite programs that arise in low-rank matrix estimation. There is a long and still growing list of such algorithms, including singular value thresholding [36], accelerated proximal gradient descent [37], which is variant of FISTA for matrix completion [38], Augmented Lagrangian Multiplier methods [39], Frank-Wolfe [40], [41], CoGENT [42], and ADCG [43], just to name a few. Below, we discuss two representative algorithms: FISTA for solving the regularized problem (10), and Frank-Wolfe for solving the constrained problem (11). These two algorithms provide the stage for understanding many other algorithms.

An important subroutine in many of the aforementioned algorithms is the Singular Value Thresholding (SVT) operator $\mathcal{D}_\tau(\cdot)$ [36]. Mathematically, $\mathcal{D}_\tau(\mathbf{Y})$ is defined as the proximal mapping of \mathbf{Y} with respect to the nuclear norm:

$$\mathcal{D}_\tau(\mathbf{Y}) = \underset{\mathbf{Z}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{Y} - \mathbf{Z}\|_F^2 + \tau\|\mathbf{Y}\|_*, \quad (15)$$

which admits a closed-form expression. In particular, if the SVD of \mathbf{Y} is $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ with $\mathbf{\Sigma} = \operatorname{diag}[\sigma_1, \sigma_2, \dots]$, then $\mathcal{D}_\tau(\mathbf{Y}) = \mathbf{U}\mathbf{\Sigma}'\mathbf{V}^T$, where $\mathbf{\Sigma}' = \operatorname{diag}[\sigma'_1, \sigma'_2, \dots]$ with

$$\sigma'_k = \begin{cases} \sigma_k - \tau, & \sigma_k \geq \tau, \\ 0, & \sigma_k < \tau. \end{cases} \quad (16)$$

The fact that the SVT operator can be efficiently computed via SVD is leveraged in many of the first-order algorithms.

The FISTA algorithm for the regularized problem (10) is given in Algorithm 1, where L is an upper bound of the Lipschitz constant of $\nabla f(\mathbf{X})$, with $f(\mathbf{X}) := \frac{1}{2}\|\mathbf{y} - \mathcal{A}(\mathbf{X})\|_2^2$. FISTA makes use of Nesterov's momentum acceleration to speed up the convergence. Denote the objective function of (10) as $g(\mathbf{X})$, then to achieve ϵ -accuracy, i.e. $g(\mathbf{X}_T) - g(\hat{\mathbf{X}}) \leq \epsilon$, we need $T = \mathcal{O}\left(\sqrt{L/\epsilon}\right)$ iterations. In each iteration, only a partial SVD is needed to evaluate the SVT operator. Doing so for large-scale problems may still be too slow and require large memory. In this case, one can make use of modern randomized techniques from numerical linear algebra to further speed up the computation of SVD [44], [45].

Algorithm 2 Frank-Wolfe for Low-Rank Matrix Estimation

Input: T, γ in (11);

Initialization: Let $\mathbf{X}_1 = \mathbf{0}$,

Updates: for $k = 0, 1, 2, \dots, T - 1$ do

- 1) $\mathbf{S}_k = \gamma \mathbf{u} \mathbf{v}^T$, where \mathbf{u} and \mathbf{v} are the left and right top singular vector of $\mathcal{A}^*(\mathbf{y} - \mathcal{A}(\mathbf{X}_k))$;
- 2) $\mathbf{X}_{k+1} = (1 - \frac{2}{k+1})\mathbf{X}_k + \frac{2}{k+1}\mathbf{S}_k$;

Output: \mathbf{X}_T .

The standard Frank-Wolfe method [40] (also known as conditional gradient descent) for solving the constrained problem (11) is presented in Algorithm 2. Each iteration of algorithm only requires computing a rank-one SVD, which can be done using power iteration or Lanczos methods. Therefore, Frank-Wolfe typically has a much lower computational cost per iteration than methods based on the SVT operation. However, standard Frank-Wolfe may converge very slowly in practice. To achieve ϵ -accuracy, i.e. $f(\mathbf{X}_T) - f(\hat{\mathbf{X}}) \leq \epsilon$, Frank-Wolfe requires $T = \mathcal{O}(1/\epsilon)$ iterations, which can be quite slow. Variants of Frank-Wolfe with faster convergence or lower memory footprint have been actively developed recently by exploiting the problem structures. The list includes CoGENT [42], In-Face Extended Frank-Wolf [41], and ADCG [43], Block Frank-Wolfe [46], and sketchyCGM [47], is still growing.

V. PROVABLE AND FAST LOW-RANK MATRIX ESTIMATION VIA NONCONVEX FACTORIZATION

As we have seen, the computational concern of solving rank minimization problems is assuaged to some extent by the use of convex relaxation — the resulting semidefinite programs can be solved in time polynomial in the matrix dimension. However, for large-scale problems where the dimension is on the order of millions, solving these semidefinite programs, even using first-order methods, can still be computationally infeasible due to the fundamental bottleneck of storing and optimizing over a matrix variable. This issue severely limits the applicability of the convex relaxation methods.

To overcome this difficulty, a recent line of work studies more computationally efficient methods that are based on nonconvex optimization. These methods work directly with the original nonconvex, rank-constrained optimization problem, which can be generally written as

$$\min_{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}} F(\mathbf{X}) \quad \text{subject to} \quad \text{rank}(\mathbf{X}) \leq r, \quad (17)$$

where $F : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}$ is a given loss function, which typically is convex in \mathbf{X} . The key idea is to use a reparametrization trick: by writing a rank- r matrix in its factorization form $\mathbf{X} = \mathbf{L}\mathbf{R}^T$, where $\mathbf{L} \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{R} \in \mathbb{R}^{n_2 \times r}$, we enforce the low-rank constraint directly, leading to the following equivalent formulation of (17):

$$\min_{\mathbf{L} \in \mathbb{R}^{n_1 \times r}, \mathbf{R} \in \mathbb{R}^{n_2 \times r}} f(\mathbf{L}, \mathbf{R}) := F(\mathbf{L}\mathbf{R}^T). \quad (18)$$

We refer to this formulation as the *Burer-Monteiro* factorization, after the seminal work [48]. The optimization problem (18) can then be solved over the factor variables \mathbf{L} and \mathbf{R} . The

low-rank factorization $\mathbf{X} = \mathbf{L}\mathbf{R}^T$ is in general not unique; in fact, any pair $\tilde{\mathbf{L}} = \mathbf{L}\mathbf{Q}$ and $\tilde{\mathbf{R}} = \mathbf{R}\mathbf{Q}$ with $\mathbf{Q} \in \mathbb{R}^{r \times r}$ being an orthonormal matrix also corresponds to the same matrix \mathbf{X} , since $\tilde{\mathbf{L}}\tilde{\mathbf{R}}^T = \mathbf{L}\mathbf{Q}\mathbf{Q}^T\mathbf{R} = \mathbf{L}\mathbf{R}$. These pairs are all global optima of the problem. This reformulation brings a significant computational gain: since the rank r is often much smaller than $\min\{n_1, n_2\}$, the size of the variables (\mathbf{L}, \mathbf{R}) is roughly linear in $(n_1 + n_2)$ rather than quadratic, leading to the possibility of designing *linear-time* algorithms that are amenable to problems of very large scale.

Surprisingly, even though the Burer-Monteiro formulation (18) is nonconvex, global optima can sometimes be found (or approximated) efficiently using various iterative procedures; moreover, rigorous guarantees can be derived for the statistical accuracy of the resulting solution. Indeed, several iterative schemes has a computational cost proportional to $(n_1 + n_2) \text{poly}(r)$ and the size of the input, at least per iteration, which is typically much lower than $n_1 \times n_2$. These results are developed in a still growing line of recent work [49]–[69], and we devote the rest of this section to presenting the most representative results therein.

This line of work considers three major classes of iterative schemes for solving the Burer-Monteiro formulation (18):

- **(Projected) gradient descent** [48], [49], [67]: One runs (projected) gradient descent directly on the loss function $f(\mathbf{L}, \mathbf{R})$ with respect to the factor variables (\mathbf{L}, \mathbf{R}) :

$$\mathbf{L}^{t+1} = \mathcal{P}_{\mathcal{L}} \left[\mathbf{L}^t - \eta^t \nabla_{\mathbf{L}} f(\mathbf{L}^t, \mathbf{R}^t) \right], \quad (19a)$$

$$\mathbf{R}^{t+1} = \mathcal{P}_{\mathcal{R}} \left[\mathbf{R}^t - \eta^t \nabla_{\mathbf{R}} f(\mathbf{L}^t, \mathbf{R}^t) \right]. \quad (19b)$$

where η^t is the step size and $\mathcal{P}_{\mathcal{L}}, \mathcal{P}_{\mathcal{R}}$ denote the Euclidean projection onto the sets \mathcal{L} and \mathcal{R} , which are constraint sets that encode additional structures of the desired low-rank factors.

- **Alternating minimization** [59], [60]: One optimizes the loss function $f(\mathbf{L}, \mathbf{R})$ alternatively over one of the factors while fixing the other, which is a convex problem. In particular, each iteration takes the form

$$\mathbf{L}^{t+1} = \underset{\mathbf{L} \in \mathbb{R}^{n_1 \times r}}{\text{argmin}} f(\mathbf{L}, \mathbf{R}^t), \quad (20a)$$

$$\mathbf{R}^{t+1} = \underset{\mathbf{R} \in \mathbb{R}^{n_2 \times r}}{\text{argmin}} f(\mathbf{L}^{t+1}, \mathbf{R}). \quad (20b)$$

- **Singular value projection (SVP)** [62], [63], [64]: One performs a gradient descent step of $F(\mathbf{L}\mathbf{R}^T)$ on the “full” $n_1 \times n_2$ matrix space, then projects back to the factor space via SVD:

$$(\mathbf{L}^{t+1}, \mathbf{R}^{t+1}) = \text{SVD}_r \left[\mathbf{L}^t \mathbf{R}^{tT} - \eta^t \nabla F(\mathbf{L}^t \mathbf{R}^{tT}) \right]. \quad (21)$$

where η^t is the step size, and $\text{SVD}_r(\mathbf{Z})$ returns the top rank- r factors of \mathbf{Z} , that is, the pair $(\mathbf{U}\mathbf{\Sigma}^{1/2}, \mathbf{V}\mathbf{\Sigma}^{1/2})$ assuming that $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \mathbf{Z}_r = \mathcal{P}_r(\mathbf{Z})$ is SVD of the best rank- r approximation of \mathbf{Z} .

Because neither the function $f(\mathbf{L}, \mathbf{R})$ nor the set of low-rank matrices are convex, standard global convergence theory for convex optimization does not apply here. The recent breakthrough is based on the realization that convexity is

in fact not necessary for the convergence of these iterative schemes; instead, as long as the gradients of function $f(\mathbf{L}, \mathbf{R})$ always point (approximately) towards the desired solution, the iterates will make progress along the right direction. Note that this property concerns the *geometry* of $f(\mathbf{L}, \mathbf{R})$ itself, and is largely independent of the specific choice of the algorithm [50]. Among the above three options, the projected gradient descent approach stands out due to its simple form, cheap per-iteration cost (no SVD or inner optimization is needed) and efficiency with constrained problems. We thus use projected gradient descent as the focal point of our survey.

Playing a key role here is the use of statistical modeling and probabilistic analysis: we will show that $f(\mathbf{L}, \mathbf{R})$ has the desired geometric properties with high probability under probabilistic generative models of the data, thereby circumventing the worst-case hardness of the low-rank matrix estimation problem and instead focusing on its average-case behavior. Existing results in this direction can be divided into two categories. In the first line of work reviewed in Section V-A, one shows that iterative algorithms converge to the desired solution rapidly when initialized within a large neighborhood around the ground truth; moreover, a good initial solution can be obtained efficiently by simple procedures (which typically involve computing a partial SVD). The second line of work, reviewed in Section V-B, concerns the global landscape of the loss function, and aims to show that in spite of the non-convexity of $f(\mathbf{L}, \mathbf{R})$, all of its local minima are in fact close to the desired solution in an appropriate sense whereas all other stationary points (e.g., saddle points) possess a descent direction; these properties guarantee the convergence of iterative algorithms from any initial solution. Below we review the most representative results in each of these two categories for noiseless matrix sensing and matrix completion, both of which have their own merits and hence are complementary to each other. The readers are referred to [49], [67], [70] for extensions to the noisy case.

A. Convergence Guarantees with Proper Initialization

For simplicity, we assume that the truth \mathbf{X} is exactly rank- r and has a bounded condition number $\kappa = \sigma_1/\sigma_r$, thus effectively hiding the dependence on κ . Moreover, to measure the convergence of the algorithms for the purpose of reconstructing \mathbf{X} , we shall consider directly the *reconstruction error* $\|\mathbf{L}^t \mathbf{R}^{tT} - \mathbf{X}\|_{\text{F}}$ with respect to \mathbf{X} .

For matrix sensing, we take the loss function $f(\mathbf{L}, \mathbf{R})$ as

$$f_{\mathcal{A}}(\mathbf{L}, \mathbf{R}) = \|\mathcal{A}(\mathbf{L}\mathbf{R}^T) - \mathbf{y}\|_2^2 + \frac{1}{8}\|\mathbf{L}^T \mathbf{L} - \mathbf{R}^T \mathbf{R}\|_{\text{F}}^2, \quad (22)$$

where the second regularization term encourages \mathbf{L} and \mathbf{R} to have the same scale and ensures algorithmic stability. We can perform gradient descent specified in (19) with $\mathcal{L} = \mathbb{R}^{n_1 \times r}$ and $\mathcal{R} = \mathbb{R}^{n_2 \times r}$ (that is, without additional projections).

Similarly, for matrix completion, we use the loss function

$$f_{\Omega}(\mathbf{L}, \mathbf{R}) = \frac{1}{p}\|\mathcal{P}_{\Omega}(\mathbf{L}\mathbf{R}^T - \mathbf{Y})\|_{\text{F}}^2 + \frac{1}{32}\|\mathbf{L}^T \mathbf{L} - \mathbf{R}^T \mathbf{R}\|_{\text{F}}^2. \quad (23)$$

Since we can only hope to recover matrices that satisfy the incoherence property, we perform projected gradient descent by projecting to the constraint set:

$$\mathcal{L} := \left\{ \mathbf{L} \in \mathbb{R}^{n_1 \times r} \mid \|\mathbf{L}\|_{2,\infty} \leq \sqrt{\frac{2\mu r}{n_1}} \|\mathbf{L}^0\| \right\},$$

with \mathcal{R} defined similarly. Note that \mathcal{L} is convex, and depends on the initial solution \mathbf{L}^0 . The projection $\mathcal{P}_{\mathcal{L}}$ is given by the row-wise ‘‘clipping’’ operation

$$[\mathcal{P}_{\mathcal{L}}(\mathbf{L})]_i = \begin{cases} \mathbf{L}_i, & \|\mathbf{L}_i\|_2 \leq \sqrt{\frac{2\mu r}{n_1}} \|\mathbf{L}^0\|, \\ \mathbf{L}_i \cdot \sqrt{\frac{2\mu r}{n_1}} \frac{\|\mathbf{L}^0\|}{\|\mathbf{L}_i\|_2}, & \|\mathbf{L}_i\|_2 > \sqrt{\frac{2\mu r}{n_1}} \|\mathbf{L}^0\|, \end{cases}$$

for $i = 1, 2, \dots, n_1$, where \mathbf{L}_i is the i -th row of \mathbf{L} ; the projection $\mathcal{P}_{\mathcal{R}}$ is given by a similar formula. This projection ensures that the iterates of projected gradient descent (19) remain incoherent.

The following theorems guarantee that if the initial solution $(\mathbf{L}^0, \mathbf{R}^0)$ is reasonably close to the desired solution, then the iterates converge linearly to the ground truth (in terms of the reconstruction error), under conditions similar to nuclear norm minimization. Moreover, we can find a provably good initial solution using the so-called spectral method, which involves performing a partial SVD. In particular, we compute $(\mathbf{L}^0, \mathbf{R}^0) = \text{SVD}_r[\mathcal{A}^*(\mathbf{y})]$ for matrix sensing, and $(\mathbf{L}^0, \mathbf{R}^0) = \text{SVD}_r[p^{-1}\mathcal{P}_{\Omega}(\mathbf{Y})]$ for matrix completion.

Theorem 3 (Matrix Sensing): [49], [62], [51] Suppose that the sensing operator \mathcal{A} satisfies RIP- ℓ_2/ℓ_2 with parameter $\delta_{4r} = \max\{\bar{\delta}_{4r}, \delta_{4r}\} \leq c_1/r$ for some sufficiently small constant c_1 . Then, the initial solution $(\mathbf{L}^0, \mathbf{R}^0) = \text{SVD}_r[\mathcal{A}^*(\mathbf{y})]$ satisfies

$$\|\mathbf{L}^0 \mathbf{R}^{0T} - \mathbf{X}\|_{\text{F}} \leq c_0 \sigma_r, \quad (24)$$

where c_0 is some sufficiently small constant. Furthermore, starting from any $(\mathbf{L}^0, \mathbf{R}^0)$ that satisfies (24), the gradient descent iterates $\{(\mathbf{L}^t, \mathbf{R}^t)\}_{t=1}^{\infty}$ with an appropriate step size satisfy the bound

$$\|\mathbf{L}^t \mathbf{R}^{tT} - \mathbf{X}\|_{\text{F}}^2 \leq (1 - \delta)^t \|\mathbf{L}^0 \mathbf{R}^{0T} - \mathbf{X}\|_{\text{F}}^2,$$

where $0 < \delta < 1$ is a universal constant.

Theorem 4 (Matrix Completion): [71] There exists a positive constant C such that if

$$p \geq C \frac{\mu_0^2 r^2 \log(n_1 + n_2)}{(n_1 + n_2)}, \quad (25)$$

then with probability at least $1 - c(n_1 + n_2)^{-1}$ for some positive constant c , the initial solution satisfies

$$\|\mathbf{L}^0 \mathbf{R}^{0T} - \mathbf{X}\|_{\text{F}} \leq c_0 \sigma_r. \quad (26)$$

Furthermore, starting from any $(\mathbf{L}^0, \mathbf{R}^0)$ that satisfies (26), the projected gradient descent iterates $\{(\mathbf{L}_t, \mathbf{R}_t)\}_{t=0}^{\infty}$ with an appropriate step size satisfy the bound

$$\|\mathbf{L}^t \mathbf{R}^{tT} - \mathbf{X}\|_{\text{F}}^2 \leq \left(1 - \frac{\delta}{\mu_0 r}\right)^t \|\mathbf{L}^0 \mathbf{R}^{0T} - \mathbf{X}\|_{\text{F}}^2,$$

where $0 < \delta < 1$ is a universal constant.

The above theorems guarantees that the gradient descent iterates enjoy geometric convergence to a global optimum when

the initial solution is sufficiently close to the ground truth. Comparing with the guarantees for nuclear norm minimization, (projected) gradient descent succeeds under a similar sample complexity condition (up to a polynomial term in r and $\log n$), but the computational cost is significantly lower. To obtain ε -accuracy, meaning that the final estimate $(\widehat{\mathbf{L}}, \widehat{\mathbf{R}})$ satisfies

$$\|\widehat{\mathbf{L}}\widehat{\mathbf{R}}^T - \mathbf{X}\|_F \leq \varepsilon \cdot \sigma_r, \quad (27)$$

we only need to run a total of $T = \mathcal{O}(\log(1/\varepsilon))$ iterations for matrix sensing, and $T = \mathcal{O}(\mu_0 r \log(1/\varepsilon))$ iterations for matrix completion.

We now briefly discuss the overall computational complexity, and for simplicity we shall assume $n_1 = n_2 = n$. For matrix sensing, let T_0 be the maximum time of multiplying the matrix \mathbf{A}_l with a vector of compatible dimension. Each gradient step (19) can be performed in time $\mathcal{O}(mrT_0 + nr)$. The complexity of computing the initial solution is a bit subtler. To this end, we first note that by standard matrix perturbation bounds, one can show that the matrix $\mathcal{A}^*(\mathbf{y})$ and its singular values are sufficiently close to those of \mathbf{X} under the condition of Theorem 3; in particular, one has $\|\mathcal{A}^*(\mathbf{y}) - \mathbf{X}\| \leq c_0 \sigma_r / \sqrt{r}$, and the r -th and $(r+1)$ -th singular values of $\mathcal{A}^*(\mathbf{y})$ are at most $c_0 \sigma_r / \sqrt{r}$ away from the corresponding singular values of \mathbf{X} , where c_0 is a small constant. With such properties of $\mathcal{A}^*(\mathbf{y})$, one does *not* need to compute the exact *singular values/vectors* of $\mathcal{A}^*(\mathbf{y})$ in order to meet the initialization condition (24); rather, it suffices to find a *rank- r approximation* of $\mathcal{A}^*(\mathbf{y})$ with the property $\|\mathbf{L}^0 \mathbf{R}^{0T} - \mathcal{A}^*(\mathbf{y})\| \leq c_0 \sigma_r / \sqrt{r}$. This can be done using for example the Randomized SVD procedure in [45], which takes $\mathcal{O}(mrT_0 \log n + nr^2)$ time to compute; see [45, Theorem 1.2 and Eq. (1.11)] with $q = \log n$. Putting together, the overall time complexity is $\mathcal{O}(mrT_0 \log(n/\varepsilon))$ for achieving ε -accuracy.

For matrix completion, to obtain the initial solution, we can again follow similar arguments as above to show that it suffices to compute a rank- r approximation of the matrix $p^{-1}\mathcal{P}_\Omega(\mathbf{Y})$, which is close to \mathbf{X} and has a sufficiently large spectral gap. Since $p^{-1}\mathcal{P}_\Omega(\mathbf{Y})$ is a sparse matrix with support Ω , computing such an approximation can be done in time $\mathcal{O}(r|\Omega| \log n + nr^2)$ using the Randomize SVD procedure in [45]. Each step of gradient descent requires computing the gradient and the projection onto \mathcal{L} and \mathcal{R} . Both of them only involve operations on sparse matrices supported on Ω and thin matrices, and can be done in time $\mathcal{O}(r|\Omega| + nr^2)$. Therefore, projected gradient descent achieves ε -accuracy with running time $\mathcal{O}(r|\Omega| \log n \log(1/\varepsilon))$.

Remark 1: Via a refined analysis of gradient descent, it is in fact possible to drop the projection step onto \mathcal{L} and \mathcal{R} in matrix completion without performance loss; cf. [67]. In particular, as long as $p \geq C \frac{\mu_0^3 r^3 \log^3(n_1 + n_2)}{(n_1 + n_2)}$ for some constant C , gradient descent converges geometrically, and needs $T = \mathcal{O}(\log(1/\varepsilon))$ iterations to reach ε -accuracy for the reconstruction error measured not only in the Frobenius norm $\|\mathbf{L}^t \mathbf{R}^{tT} - \mathbf{X}\|_F$, but also in the spectral norm $\|\mathbf{L}^t \mathbf{R}^{tT} - \mathbf{X}\|$ and the entry-wise infinity norm $\|\mathbf{L}^t \mathbf{R}^{tT} - \mathbf{X}\|_\infty$.

Remark 2: In the noisy setting, the algorithms can be applied without change, and the same error bounds hold with

an additional term that depends on the noise. For matrix completion [49], this term is $\frac{\sqrt{r}\|\mathcal{P}_\Omega(\mathbf{W})\|}{p}$, where $\mathcal{P}_\Omega(\mathbf{W})$ is the noise matrix supported on the observed indices. This term can be bounded under various noise models. For example, when \mathbf{W} has i.i.d. Gaussian or $\pm\sigma$ entries with zero mean and variance σ^2 , then $\|\mathcal{P}_\Omega(\mathbf{W})\| \lesssim \sigma \sqrt{p(n_1 + n_2)}$ with high probability. The resulting error bound is optimal in an information-theoretic sense [9]. See also [67] for the near-optimal error control in the spectral norm and the entry-wise infinity norm.

B. Global Geometry and Saddle-Point Escaping Algorithms

A very recent line of work studies the *global* geometry of the Burer-Monteiro formulation (18), as well as its computational implications for algorithms starting at an *arbitrary* initial solution [54], [55], [69], [57]. These results are based on a geometric notion called the *strict saddle property* [54], [70].

Definition 3 (Strict Saddle): A function $g(\mathbf{x})$ is said to be $(\varepsilon, \gamma, \zeta)$ -strict saddle, if for each \mathbf{x} at least one of the following holds:

- $\|\nabla g(\mathbf{x})\|_2 \geq \varepsilon > 0$;
- $\lambda_{\min}(\nabla^2 g(\mathbf{x})) \leq -\gamma < 0$;
- there exists a local minimum \mathbf{x}^* such that $\|\mathbf{x} - \mathbf{x}^*\|_2 \leq \zeta$.

In Figure 2 and Figure 3 we provide examples of a one-dimensional function and a two-dimensional function, respectively, that satisfy the strict saddle property in Definition 3. Intuitively, the strict saddle property of $g(\mathbf{x})$ ensures that whenever one is not already close to a local minimum, the current solution will have either a large gradient, or a descent direction due to the Hessian having a negative eigenvalue. Therefore, any local search algorithms that are capable of finding such a descent direction, will make progress in decreasing the value of $g(\mathbf{x})$ and eventually converge to a local minimum. Many algorithms have been shown to enjoy this property, include cubic regularization [72], trust-region algorithms [73], stochastic gradient descent [74], and other more recent variants [75], [76], [77]. In Algorithm 3, we describe one such algorithm, namely the Perturbed Gradient Descent (PGD) algorithm from [75]. PGD is based on the standard gradient descent algorithm with the following additional steps: (i) when the gradient is small, indicating potential closeness to a saddle point, PGD adds a random perturbation to the current iterate (which is done at most once every t_{thres} iterations); (ii) if the last perturbation occurs t_{thres} iterations ago and the function value does not decrease sufficiently since, then PGD terminates and outputs the iterate before the last perturbation.

We do not delve further into the details of these saddle-escaping algorithms, as their parameter choices and run-time guarantees are somewhat technical. Rather, for the purpose of analyzing low-rank matrix estimation problems, we simply rely on the existence of such algorithms and the fact that their running time depends polynomially on the problem parameters. This is summarized in the following theorem, which abstracts out the key results in the work cited above.

Theorem 5 (Optimizing Strict Saddle Functions): Assume that $g: \mathbb{R}^N \rightarrow \mathbb{R}$ is β -smooth and $(\varepsilon, \gamma, \zeta)$ -strict saddle. There exist algorithms (such as PGD in Algorithm 3 with appropriate choices of the parameters) that output a solution that is ζ -close

Algorithm 3 Perturbed Gradient Descent [75]

Input: algorithm parameters $d_{\text{thres}}, t_{\text{thres}}, g_{\text{thres}}, \eta, R$;
Initialization: Let $\mathbf{x}_0 = \mathbf{0}$;
for $t = 0, 1, 2, \dots$ **do**
 1) **if** $\|\nabla g(\mathbf{x}_t)\|_2 \leq d_{\text{thres}}$ and $t - t_{\text{last}} > t_{\text{thres}}$ **then**
 $\mathbf{x}_{\text{last}} = \mathbf{x}_t, t_{\text{last}} = t$;
 $\mathbf{x}_t = \mathbf{x}_{\text{last}} + \xi$, where ξ is sampled uniformly from
 the unit ball centered at zero with radius R ;
 2) **if** $t - t_{\text{last}} = t_{\text{thres}}$ and $g(\mathbf{x}_t) - g(\mathbf{x}_{\text{last}}) > -g_{\text{thres}}$ **then**
 return \mathbf{x}_{last} ;
 3) $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \nabla g(\mathbf{x}_t)$.




Fig. 2. Example of a one-dimensional strict saddle function: $g(x) = (x^2 - 1)^2$. This function has two local minima $x^* = 1$ and $x^* = -1$ that are also global minima, as well as a local maximum $\bar{x} = 0$. In the region R_1 , the function satisfies $\lambda_{\min}(\nabla^2 g(x)) \leq -\gamma < 0$. In the regions R_2 and R_3 , the function satisfies $\|x - x^*\|_2 \leq \zeta$ for one of the local minima x^* . On the rest of the real line, the function satisfies $|\nabla g(x)| \geq \epsilon > 0$.

to a local minimum of g , with the required number of iterations upper bounded by a polynomial function of $N, \beta, 1/\epsilon, \gamma$ and $1/\zeta$.

Specializing to the low-rank matrix estimation problem, it remains to verify that (i) the loss function defined on (\mathbf{L}, \mathbf{R}) is strict saddle and (ii) all its local minima are “good”, in the sense that they correspond to a low-rank matrix equal to (in the noisy case, close to) the true matrix \mathbf{X} . To this end, we consider the same loss function as before for matrix sensing:

$$g_{\mathcal{A}}(\mathbf{L}, \mathbf{R}) = f_{\mathcal{A}}(\mathbf{L}, \mathbf{R}),$$

where $f_{\mathcal{A}}$ is defined in (22). For matrix completion, we consider a regularized loss function:

$$g_{\Omega}(\mathbf{U}, \mathbf{V}) = f_{\Omega}(\mathbf{L}, \mathbf{R}) + \lambda Q_{\alpha}(\mathbf{L}, \mathbf{R}),$$

where f_{Ω} is given in (23), the regularizer is given by [50]

$$Q_{\alpha}(\mathbf{L}, \mathbf{R}) = \sum_{i=1}^{n_1} (\|e_i^T \mathbf{L}\|_2 - \alpha)_+^4 + \sum_{j=1}^{n_2} (\|e_j^T \mathbf{R}\|_2 - \alpha)_+^4,$$

α and λ are regularization parameters, and $(x)_+ = \max\{x, 0\}$. The regularization plays a similar role as the projections $\mathcal{P}_{\mathcal{L}}, \mathcal{P}_{\mathcal{R}}$ previously used: it encourages incoherence of \mathbf{L} and \mathbf{R} . Replacing projections with regularization leads to




Fig. 3. Example of a two-dimensional strict saddle function $g(\mathbf{x}) = \|\mathbf{x}\mathbf{x}^T - \mathbf{1}\mathbf{1}^T\|_F^2$, where $\mathbf{x} = [x_1, x_2]^T$ and $\mathbf{1} = [1, 1]^T$. This function has two local minima $\mathbf{x}^* = [1, 1]^T$ and $\mathbf{x}^* = [-1, -1]^T$, and a strict saddle point $\mathbf{x}_{\text{saddle}} = [0, 0]^T$. The Hessian $\nabla^2 g(\mathbf{x}_{\text{saddle}})$ at the saddle point has a strictly negative eigenvalue with the eigenvector $[\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}]^T$, which corresponds to the descent direction $\mathbf{x}^* - \mathbf{x}_{\text{saddle}}$.

an unconstrained formulation that fits into the strict-saddle framework above. The follow theorems show that these loss functions indeed have the desired strict-saddle property with high probability under sample complexity conditions similar to before.

Theorem 6 (Matrix Sensing): [55], [70] Suppose that the measurement operator \mathcal{A} satisfy RIP- ℓ_2/ℓ_2 with parameter $\delta_{2r} = \max\{\delta_{2r}, \delta_r\} = \frac{1}{20}$. For any $\epsilon > 0$, the above loss function $g_{\mathcal{A}}$ satisfies the following: (i) it is $(\epsilon, \Omega(\sigma_r), \mathcal{O}(\frac{\epsilon}{\sigma_r}))$ -strict saddle, and (ii) all its local minima satisfying $\mathbf{L}\mathbf{R}^T = \mathbf{X}$.

Theorem 7 (Matrix Completion): [54], [70] Suppose that the observation probability satisfies $p \geq \Omega(\frac{\mu^4 r^6 \log(n_1 + n_2)}{(n_1 + n_2)})$, and one chooses $\alpha^2 = \Theta(\frac{\mu r \sigma_1}{(n_1 + n_2)})$ and $\lambda = \Theta(\frac{(n_1 + n_2)}{\mu r})$. Then, with probability at least $1 - (n_1 + n_2)^{-1}$, the above loss function g_{Ω} satisfies the following: (i) it is $(\epsilon, \Omega(\sigma_r), \mathcal{O}(\frac{\epsilon}{\sigma_r}))$ -strict saddle for any $\epsilon \leq \text{poly}(\frac{\mu^4 r^4 \sigma_1^4}{(n_1 + n_2)^2})$, and (ii) all its local minima satisfy $\mathbf{L}\mathbf{R}^T = \mathbf{X}$.

Combining Theorem 5 with Theorem 6 and Theorem 7, we conclude that iterative algorithms optimizing over the factor variables (\mathbf{L}, \mathbf{R}) converge globally to some pair satisfying $\mathbf{L}\mathbf{R}^T = \mathbf{X}$ in a polynomially number of iterations from any arbitrary initial solutions, as long as they can escape saddle points. We refer the readers to [70], [75] for more discussions.

C. Perspectives

Combining the discussions in the last two subsections, we obtain the following general picture for low-rank matrix estimation reformulated in the factor space of (\mathbf{L}, \mathbf{R}) :

- All the local minima of the loss function are in fact global minima, and correspond to some factorization $\mathbf{L}\mathbf{R}^T$ of the true low-rank matrix \mathbf{X} .
- In a neighborhood of each global minimum (\mathbf{L}, \mathbf{R}) , the loss function is essentially strongly convex, and has no saddle point. Within this neighborhood, gradient descent and other iterative algorithms converge geometrically.

- Any point outside such neighborhoods either has a strictly positive gradient, or is a saddle point with a descent direction corresponding to a strictly negative eigenvalue of the Hessian.
- Iterative algorithms escape all saddle points and enter a neighborhood of the global minima in polynomial time. Alternatively, one can find a solution in this neighborhood by performing one SVD of a matrix appropriately constructed from the observations.

Comparing the two approaches to non-convex matrix estimation discussed in the last two subsections, we also see that each of them has its own strengths. The first approach, taken in Section V-A, focuses on the convergence of algorithms with a proper initialization procedure. This approach immediately leads to simple, efficient algorithms, with provably geometric convergence and linear time-complexity. It also readily extends to problems that have additional structural constraints (such as sparsity, Hankel and discrete structures), or involve more complicated loss functions (such as robust PCA and matrix completion with quantized observations), some of which may involve a non-smooth loss function whose Hessian is not defined. However, finding a good initialization scheme is non-trivial, and in some settings is actually the harder part of the problem. The second approach, taken in Section V-B, instead focuses on the global geometric landscape of the problem. This approach is conceptually elegant, nicely decomposing the geometric aspect (properties of the local minima) and the algorithmic aspect (how to find the local minima) of the problem. Computationally, it eliminates the need of careful initialization, but the resulting run-time guarantees are somewhat weaker, which may be super-linear in the dimensions. Of course, we made the distinction between these two approaches mostly for ease of review of state-of-the-art; given the rapid developments in this area, we expect that both approaches will be improved, expanded, and eventually merged.

Before concluding this section, we add that there is a deeper reason for low-rank matrix estimation being such a benign nonconvex problem. The loss function of the low-rank matrix estimation problem can often be viewed, in a certain precise sense, as a perturbed version of the objective function of PCA, i.e., finding the best rank- r approximation in Frobenius norm in a factorized form:

$$\min_{\mathbf{L}, \mathbf{R} \in \mathbb{R}^{n \times r}} \|\mathbf{L}\mathbf{R}^T - \mathbf{X}\|_{\mathbb{F}}^2. \quad (28)$$

For example, the matrix completion loss function (23), with the regularization omitted, is exactly equal to the above objective *in expectation*. The PCA problem (28) is arguably the most well-understood tractable non-convex problem: in addition to having a closed-form solution (4), this problem satisfies all the geometric properties mentioned in the last two subsections [78], [79]; in particular, its local minima and saddle points can be expressed in terms of the top and non-top eigen components of \mathbf{X} , respectively. Under the probabilistic or RIP assumptions on the sensing operators, the geometric and algorithmic properties of the PCA problem (28) are essentially preserved under incomplete observations, with high probability, as long as the issue of incoherence is appropriately accounted for.

VI. STRUCTURED LOW-RANK MATRIX ESTIMATION

In many applications, the low-rank matrix estimation problems possess additional structures that need to be carefully exploited, and we present two such examples in this section: Hankel matrix completion, and the recovery of clustering matrices.

A. Hankel Matrix Completion

Imagine that one is interested in estimating the spectrum of time series, or the direction-of-arrivals from returns from a sensor array. One can model the signal of interest as a weighted sum of complex exponentials, i.e.,

$$\mathbf{x} = \sum_{i=1}^r c_i \mathbf{v}(z_i), \quad (29)$$

where $c_i \in \mathbb{C}$ represents the complex amplitude, $z_i \in \mathbb{C}$ is the complex frequency, and $\mathbf{v}(z_i)$ is of the form

$$\mathbf{v}(z) = [1 \quad z \quad \cdots \quad z^{n-1}]^T \in \mathbb{C}^n. \quad (30)$$

One can view the atom $\mathbf{v}(z)$ as an eigenvector of a linear time-invariant system, with z being the corresponding pole.

Spectral compressed sensing concerns with the problem of recovering the signal \mathbf{x} from only a subset of its entries. This is an important problem in super resolution [80], MRI imaging [81], and system identification [18]. Denoting the index set of the observed entries by $\Omega \subset \{1, \dots, n\}$, our goal is to recover \mathbf{x} given $\mathcal{P}_\Omega(\mathbf{x})$. In the literature of compressed sensing, this problem is typically approached by first discretizing the parameter space of z , and then solving an ℓ_1 minimization problem as in standard sparse recovery. However, this approach is sensitive to the discretization used, and cannot resolve parameters that fall off the grid [82].

It turns out that under certain mild conditions, this problem can be solved *exactly* as a structured matrix completion problem, without assuming any discretization/grid. The insight lies in exploiting the shift invariance property embedded in the structure of complex harmonics. This is done by constructing an n_1 -by- $(n - n_1 + 1)$ Hankel matrix spanned by the signal vector $\mathbf{x} \in \mathbb{C}^n$ as

$$\mathcal{H}(\mathbf{x}) = \begin{bmatrix} x_1 & x_2 & & & \\ x_2 & & \ddots & & \\ \vdots & & \ddots & & \\ x_{n_1} & x_{n_1+1} & \cdots & x_n & \end{bmatrix}, \quad (31)$$

where n_1 is commonly selected as $\lfloor n/2 \rfloor$ to make the matrix $\mathcal{H}(\mathbf{x})$ as square as possible. The important observation is that $\mathcal{H}(\mathbf{x})$ admits the following low-rank decomposition:

$$\mathcal{H}(\mathbf{x}) = \mathbf{V}_{n_1} \mathbf{C} \mathbf{V}_{n-n_1+1}^T, \quad (32)$$

where

$$\mathbf{V}_{n_1} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ z_1 & z_2 & \cdots & z_r \\ \vdots & \vdots & \ddots & \vdots \\ z_1^{n_1-1} & z_2^{n_1-1} & \cdots & z_r^{n_1-1} \end{bmatrix}, \quad (33)$$




Fig. 4. Illustration of structured matrices considered in Section VI. (a) The observation pattern in a Hankel matrix completion problem. (b) The cluster matrix, and (c) the affinity matrix in a cluster matrix recovery problem, when the nodes are ordered according to the cluster structure. (d) Same as (c), except that the nodes are randomly permuted.

$\mathbf{C} = \text{diag}[c_1, c_2, \dots, c_r]$, and \mathbf{V}_{n-n_1+1} is defined in a way similar to (33). This decomposition shows that $\text{rank}(\mathcal{H}(\mathbf{x})) \leq r$, and equality holds when all the poles are distinct. This representation of \mathbf{x} as a structured low-rank matrix can be leveraged to facilitate recovery of the un-measured entries of \mathbf{x} . In particular, one can try to recover the missing measurements by seeking a Hankel matrix with the smallest nuclear norm and consistent with the available measurements. This idea gives rise to the following algorithm, termed Enhanced Matrix Completion (EMaC) [80]:

$$\min_{\mathbf{g} \in \mathbb{C}^n} \|\mathcal{H}(\mathbf{g})\|_* \quad \text{subject to} \quad \mathcal{P}_\Omega(\mathbf{g}) = \mathcal{P}_\Omega(\mathbf{x}). \quad (34)$$

Figure 4 (a) illustrates the observation pattern in a Hankel matrix recovery problem, which is highly structured.

Under the parametric model (29), we define a new notion of incoherence that bears an interesting physical interpretation. Let the Dirichlet kernel be

$$\mathcal{D}_{n_1}(z) := \frac{1}{n_1} \left(\frac{1 - z^{n_1}}{1 - z} \right), \quad (35)$$

whose absolute value decays inverse proportionally with respect to $|z|$. Given r poles, one can construct two $r \times r$ Gram matrices \mathbf{G}_L and \mathbf{G}_R , corresponding to the column space and row space of $\mathcal{H}(\mathbf{x})$, where the entries of these matrices are specified by

$$\begin{aligned} (\mathbf{G}_L)_{i,l} &= \mathcal{D}_{n_1}(z_i - z_l), \quad 1 \leq i, l \leq r; \\ (\mathbf{G}_R)_{i,l} &= \mathcal{D}_{n-n_1+1}(z_i - z_l), \quad 1 \leq i, l \leq r. \end{aligned}$$

The incoherence parameter is then defined as follows.

Definition 4 (Incoherence): The incoherence parameter of a signal \mathbf{x} of the form (29) is defined as the smallest number μ satisfying the bounds

$$\sigma_{\min}(\mathbf{G}_L) \geq \frac{1}{\mu} \quad \text{and} \quad \sigma_{\min}(\mathbf{G}_R) \geq \frac{1}{\mu}, \quad (36)$$

where $\sigma_{\min}(\mathbf{G}_L)$ and $\sigma_{\min}(\mathbf{G}_R)$ denote the smallest singular values of \mathbf{G}_L and \mathbf{G}_R , respectively.

If all poles are well-separated by $2/n$, the incoherence parameter μ can be bounded by a small constant [83]. As the poles get closer, the Gram matrices become poorly-conditioned, resulting in a large μ . Therefore, the incoherence parameter provides a measure of the hardness of the recovery problem in terms of the relative positions of the poles. The theorem below summarizes the performance guarantees of the EMaC algorithm.

Theorem 8 (Structured Matrix Completion): [80] Suppose that each entry of \mathbf{x} is observed independently with probability p . As long as

$$p \geq C \frac{\mu r \log^4 n}{n},$$

for some sufficiently large constant C , the signal \mathbf{x} can be exactly recovered with high probability via EMaC.

Theorem 8 suggests that a Hankel-structured low-rank matrix can be faithfully recovered using a number of measurements much smaller than its dimension n . Recently it has been shown that Hankel matrix completion can also be efficiently solved using the non-convex Burer-Monteiro factorization and projected gradient descent approach described in Section V, under similar conditions [84]. Similar results can be obtained for block Hankel or Toeplitz low-rank matrix completion (for multi-dimensional data) as well. Interestingly, if the Toeplitz matrix is additionally positive-semidefinite, the incoherence condition can be relaxed by exploring the connection to Carathéodory's theorem; see [28], [85].

B. Cluster Matrices

Suppose that we are given an affinity matrix $\mathbf{Y} \in \mathbb{R}^{n \times n}$ between n nodes, where Y_{ij} is a measure of the pairwise similarity/affinity between nodes i and j . For example, each Y_{ij} may be an indicator of the friendship of two Facebook users, or the similarity between two movies on Netflix. Our goal is to partition these n nodes into several clusters such that nodes within the same clusters have high affinity values. This problem is known as *clustering* or *community detection*.

One may represent an ideal partition by a so-called *cluster matrix* $\mathbf{X}^* \in \{0, 1\}^{n \times n}$ defined as

$$X_{ij}^* = \begin{cases} 1 & \text{if nodes } i \text{ and } j \text{ are in the same cluster,} \\ 0 & \text{if nodes } i \text{ and } j \text{ are in different clusters.} \end{cases}$$

With an appropriate ordering of the rows and columns, the matrix \mathbf{X}^* takes the form of a block-diagonal matrix:

$$\mathbf{X}^* = \begin{bmatrix} \mathbf{J}_{\ell_1 \times \ell_1} & & & \\ & \mathbf{J}_{\ell_2 \times \ell_2} & & \\ & & \ddots & \\ & & & \mathbf{J}_{\ell_r \times \ell_r} \end{bmatrix},$$

where r is the number of clusters, ℓ_k is the size of the k -th cluster, and $\mathbf{J}_{\ell \times \ell}$ denotes the ℓ -by- ℓ all-one matrix. It is

clear that the rank of \mathbf{X}^* is equal to the number of clusters r . Moreover, the matrix \mathbf{X}^* has several additional structural properties: it is binary, block-diagonal, positive-semidefinite, and has all diagonal entries equal to one.

The fact that the nodes in the same cluster tend to have high affinity values, can be captured by the model $Y_{ij} = X_{ij}^* + W_{ij}$, where W_{ij} is some form of noise that encapsulates the inherent randomness/uncertainty in the pairwise affinity measure. In many applications, the affinity values between some pairs of nodes are unknown, or are costly to measure, in which case we only observe a subset Ω of the entries of \mathbf{Y} . Under this setup, the clustering problem can be cast as a noisy low-rank matrix completion problem, and the algorithms and theory in the last two sections can be immediately applied.

Notably, one can take advantage of the additional structures of the cluster matrix \mathbf{X}^* to obtain stronger performance guarantees. In particular, it is sometimes possible to *recover* \mathbf{X}^* *exactly even in the presence of noise*. We briefly review one such result from [24]. Consider the setting where the noise term W_{ij} is such that $Y_{ij} \sim \text{Bernoulli}(\tau_{\text{in}})$ if $X_{ij}^* = 1$ and $Y_{ij} \sim \text{Bernoulli}(\tau_{\text{out}})$ if $X_{ij}^* = 0$, where $\tau_{\text{in}} > \tau_{\text{out}}$; in this case, nodes in the same clusters have a higher probability of having a high (non-zero) affinity value. Figure 4 (b) and (c) illustrate the cluster matrix and the affinity matrix where the nodes are ordered according to the cluster structure, and (d) shows the affinity matrix except that the nodes are randomly permuted, as how it is typically observed in practice. As before, we assume that each entry of \mathbf{Y} is observed with some probability p . This model is sometimes referred to as the (*censored*) *stochastic block model* or *planted partition model* in the literature [86], [87]. For this model, we consider the maximum likelihood estimator of \mathbf{X}^* , and derive a convex relaxation of it by replacing the non-convex constraints on \mathbf{X}^* (low-rank, binary and block-diagonal) with a nuclear norm regularizer and linear inequality constraints $X_{ij}^* \in [0, 1]$. Doing so leads to the semidefinite program

$$\begin{aligned} \min_{\mathbf{X} \in \mathbb{R}^{n \times n}} \quad & -\alpha \langle \mathcal{P}_\Omega(\mathbf{Y}), \mathbf{X} \rangle + \alpha^{-1} \langle \mathbf{J} - \mathcal{P}_\Omega(\mathbf{Y}), \mathbf{X} \rangle \\ & + 48\sqrt{n} \|\mathbf{X}\|_* \end{aligned} \quad (37)$$

subject to $X_{ij} \in [0, 1], \forall i, j$,

where $\alpha = \sqrt{\frac{2-p(\tau_{\text{in}}+\tau_{\text{out}})}{p(\tau_{\text{in}}+\tau_{\text{out}})}}$ and $\mathbf{J} = \mathbf{J}_{n \times n}$ is the $n \times n$ all-one matrix; see [24] for the details. This approach enjoys the following guarantees.

Theorem 9 (Cluster Matrix Recovery): [24] Suppose that the minimum cluster size is ℓ_{min} . As long as

$$\frac{p(\tau_{\text{in}} - \tau_{\text{out}})^2}{\tau_{\text{in}}} \geq C \frac{n \log^2 n}{\ell_{\text{min}}^2}$$

for some constant C , the convex relaxation (37) recovers \mathbf{X}^* exactly as the unique minimizer with high probability.

In words, provided that the observation probability and the difference between τ_{in} and τ_{out} are large enough, the solution of the convex relaxation formulation is guaranteed to have the structures of a cluster matrix and equal \mathbf{X}^* exactly.

Clustering is a classical problem that has been studied extensively, with a huge body of literature. The above perspective of

casting it as a low-rank matrix recovery problem is relatively recent, and proves to be fruitful, leading to a number of new algorithms and theoretical results. A detailed account of these developments is outside the scope of this article, and we refer the readers to the recent surveys [88], [89] and the references therein.

VII. NUMERICAL EXAMPLES ON MOVIELENS DATA

In this section, we showcase some numerical results of applying the matrix completion approach to a real dataset, namely, the MovieLens 100K Dataset [90]. The dataset consists of 100,000 ratings, taking values of 1, 2, 3, 4 or 5, from 943 users on 1682 movies. We work with a popular version of the data that partitions the entire dataset ‘u.data’ into a training set ‘ua.base’ and a test set ‘ua.test’, where the test set contains exactly 10 ratings per user. We further delete two movies with no ratings available. Let \mathbf{M} denote the original incomplete rating matrix generated with rows corresponding to users and columns corresponding to movies, and let Φ be the index set of the available ratings contained in ‘u.data’. With this notation, we have $|\Phi| = 10^5$, and \mathbf{M} is a 943-by-1680 matrix where for each $(i, j) \in \Phi$, $M_{ij} \in \{1, 2, 3, 4, 5\}$. Further, let Ω and Ω^c denote the disjoint index sets of ratings for training and testing, respectively, with $\Omega \cup \Omega^c = \Phi$. In the experiment, we use the training data set Ω as the input to a matrix completion algorithm. The completed matrix produced by the algorithm, denoted by $\hat{\mathbf{M}} \in \mathbb{R}^{943 \times 1680}$, is used as an estimate of the unobserved ratings and will be evaluated on the test set Ω^c .

We demonstrate performance of three matrix completion algorithms: accelerated proximal gradient [37], singular value projection [62], and bi-factored gradient descent [91]. For these algorithms we use existing implementations with publicly available codes, and mostly adopt their default settings with only a few adjustments (detailed below) tailored to this specific dataset. For the error metric, we use the normalized mean absolute errors (NMAE) over the training set and test set, defined respectively as

$$\begin{aligned} \text{NMAE}_{\text{train}} &= \frac{1}{(5-1)|\Omega|} \sum_{(i,j) \in \Omega} \left| \hat{M}_{ij} - M_{ij} \right|, \\ \text{NMAE}_{\text{test}} &= \frac{1}{(5-1)|\Omega^c|} \sum_{(i,j) \in \Omega^c} \left| \hat{M}_{ij} - M_{i,j} \right|. \end{aligned}$$

We first employ the accelerated proximal gradient (APG) algorithm proposed in [37].³ We disable the adaptive updating of the regularization parameter in the implementation, and instead use a fixed one. We set the maximum rank to 100 and the maximum number of iterations to 1500. In our experiment, we observe that the algorithm usually meets the stop criteria after just a few of iterations, and hence stops early before reaching the maximum number of iterations. When using different values for the regularization parameter, APG outputs an estimate matrix $\hat{\mathbf{M}}$ with different ranks. Fig. 5 (a) shows the relation between the regularization parameter and rank, and Fig. 5 (b) shows NMAEs for the training data and test data against different ranks. The minimum NMAE on the test data

³<http://www.math.nus.edu.sg/~mattokc/NLS.html>

is 0.1924, which is achieved when the regularization parameter is set to 2.61 with the rank of the estimate being 5.

We next consider the singular value projection (SVP) algorithm proposed in [62].⁴ The stopping criteria are set as $\text{tol} = 10^{-3}$ and $\text{vtol} = 10^{-4}$, and the maximum number of iterations is set to 1000. Again, the algorithm usually stops early before reaching the max iterations. The step size is chosen to be $\eta = 0.1 \times 3 / (4 \times p)$, where $p = |\Omega| / (943 \times 1680)$ is the fraction of available ratings in the training set. The rank of the estimate matrix is itself a user-specified tuning parameter for SVP. The NMAEs of SVP on the training data and test data are shown in Fig. 6. The minimum NMAE for test data is 0.1929, achieved when the rank is set to 3.

Lastly, we apply the bi-factored gradient descent (BFGD) algorithm proposed in [91],⁵ which is a variant of the projected gradient descent algorithm applied to the non-convex Burer-Monteiro factorization formulation as described in Section V. We set the maximum number of iterations to 4000, and the convergence tolerance to 5×10^{-6} . Similarly as before, BFGD typically terminates early in our experiment. For the step size we use the default setting of the above implementation. The NMAEs of BFGD for the training data and test data are shown in Fig. 7. The minimum NMAE for test data is 0.1895, achieved by setting the rank to 2.

We make several observations from the above experiment results. First, we see that consistently across the three algorithms, the training error generally goes down as the rank becomes larger, whereas the test error exhibits a U-shape behavior, decreasing first and then increasing later. This phenomenon is in accordance with the bias-variance tradeoff principle described in Section II, and in particular shows that using a low-rank model is helpful in reducing the variance and prevents overfitting. Second, all three algorithms achieve a minimum test NMAE around 0.19, using a rank no more than 5. The small optimal values for the rank are likely due to the highly noisy nature of the MovieLens dataset, for which suppressing variance is crucial to good performance on the test set. Finally, while the estimation/prediction performance of these algorithms is similar, their computational costs, such as running times and memory usage, vary. These costs depend heavily on the specific implementations and termination criteria used, so we do not provide a detailed comparison here.

VIII. CONCLUDING REMARKS

Low-rank matrices represent an important class of signals with low-dimensional intrinsic structures. In this article, we have presented some recent developments on low-rank matrix estimation, focusing on the setting with incomplete measurements and additional structural constraints. We have particularly emphasized the remarkable modeling power of low-rank matrices, which are useful in a range of problems much wider than the name may suggest, including those where the presence of low-rank structures are not obvious at all. In terms of algorithms and theory, attention is paid to the integration

of statistical and computational considerations: fast algorithms have been developed that are applicable to large-scale problems, and at the same time enjoy provable performance guarantees under mild assumptions. As we have seen, such recent progress is made possible by combining techniques from diverse fields; in particular, convex and nonconvex optimization, as well as probabilistic analysis, play a key role.

We conclude by mentioning a few topics and future directions that are not covered in this article. We have focused on the matrix sensing and completion problems with linear measurements. There are many other low-rank estimation problems that are amenable to convex and nonconvex optimization-based algorithms, and enjoy similar geometric properties and performance guarantees. A partial list of such problems includes phase retrieval [57], blind deconvolution [56], robust PCA [70], dictionary learning [73], lifting for mixture problems [92], low-rank phase retrieval [93], community detection [89], and synchronization problems [69]. More broadly, applications of low-rank matrix recovery go well beyond the setting of linear measurements and least-squares objectives. Prime examples include low-rank matrix recovery with quantized, categorical and non-Gaussian data [94], [95], and ranking from comparison-based observations [96]. These problems involve more general objective functions (such as the log-likelihood) and constraints that depend on the specific observation schemes and noise structures. Another promising line of research aims at exploiting hidden low-rank structures in settings where the problem on the surface has nothing to do with low-rank matrices, yet such structures reveal themselves under suitable transformation and approximation. Problems of this type include latent variable models with certain smoothness/monotonicity properties [23].

Another topic of much interest is how to select the model rank automatically and robustly, and how to quantify the effect of model mismatch. These are important issues even in standard matrix sensing and completion; we have not discussed these issues in detail in this survey. Finally, we have omitted many other low-rank recovery algorithms that are not directly based on (continuous) optimization, including various spectral methods, kernel and nearest neighbor type methods, and algorithms with a more combinatorial flavor. Some of these algorithms are particularly useful in problems involving complicated discrete and time-evolving structures and active/adaptive sampling procedures. All of these topics are the subject of active research with tremendous potential.

ACKNOWLEDGMENT

The authors thank Mr. Yuanxin Li for preparing the numerical experiments in this paper. The work of Y. Chen is supported in part by NSF under the CRII award 1657420 and grant CCF-1704828. The work of Y. Chi is supported in part by AFOSR under the grant FA9550-15-1-0205, by ONR under the grant N00014-18-1-2142, and by NSF under the grants CAREER ECCS-1818571 and CCF-1806154.

REFERENCES

- [1] E. J. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.

⁴<http://www.cs.utexas.edu/~pjain/svp/>

⁵<http://akyrellidis.github.io/projects/>




Fig. 5. Performance of APG on MovieLens 100K dataset: (a) The rank of estimated matrix via APG with respect to the regularization parameter. (b) The NMAEs for training data and test data via APG with respect to the regularization parameter.




Fig. 6. Performance of SVP [62] on MovieLens 100K dataset: The NMAEs for training data and test data via SVP with respect to the rank.




Fig. 7. Performance of BFGD [91] on MovieLens 100K dataset: The NMAE for training data and test data via BFGD with respect to the rank.

[2] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, April 2006.

[3] S. Sudarski, T. Henzler, H. Haubenreisser, C. Dösch, M. O. Zenge, M. Schmidt, M. S. Nadar, M. Borggreffe, S. O. Schoenberg, and T. Papavassiliu, "Free-breathing sparse sampling cine MR imaging with iterative reconstruction for the assessment of left ventricular function and mass at 3.0 T," *Radiology*, vol. 282, no. 1, pp. 74–83, 2016.

[4] I. Jolliffe, *Principal Component Analysis*. Springer-Verlag, 1986.

[5] R. Basri and D. W. Jacobs, "Lambertian reflectance and linear subspaces," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 2, pp. 218–233, 2003.

[6] B. Kulis, A. C. Surendran, and J. C. Platt, "Fast low-rank semidefinite programming for embedding and clustering," in *International Conference on Artificial Intelligence and Statistics*, 2007, pp. 235–242.

[7] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2980–2998, 2010.

[8] E. J. Candès and T. Tao, "The power of convex relaxation: Near-optimal matrix completion," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2053–2080, 2010.

[9] S. Negahban and M. Wainwright, "Restricted strong convexity and

weighted matrix completion: Optimal bounds with noise," *The Journal of Machine Learning Research*, vol. 98888, pp. 1665–1697, May 2012.

[10] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717–772, April 2009.

[11] D. Gross, "Recovering low-rank matrices from few coefficients in any basis," *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1548–1566, March 2011.

[12] M. A. Davenport and J. Romberg, "An overview of low-rank matrix recovery from incomplete observations," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 4, pp. 608–622, 2016.

[13] L. Scharf and D. Tufts, "Rank reduction for modeling stationary signals," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 35, no. 3, pp. 350–355, 1987.

[14] R. Balan, P. Casazza, and D. Edidin, "On signal reconstruction without phase," *Applied and Computational Harmonic Analysis*, vol. 20, no. 3, pp. 345–356, 2006.

[15] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.

[16] A. Ahmed, B. Recht, and J. Romberg, "Blind deconvolution using convex programming," *IEEE Transactions on Information Theory*, vol. 60, no. 3, pp. 1711–1732, 2014.

[17] A. Javanmard and A. Montanari, "Localization from incomplete noisy distance measurements," *Foundations of Computational Mathematics*, vol. 13, no. 3, pp. 297–345, 2013.

[18] M. Fazel, H. Hindi, and S. P. Boyd, "Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices," *American Control Conference*, vol. 3, pp. 2156 – 2162 vol.3, June 2003.

[19] J. Bennett and S. Lanning, "The Netflix prize," in *Proceedings of KDD Cup and Workshop*, 2007, p. 35.

[20] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of the ACM*, vol. 58, no. 3, pp. 11:1–11:37, 2011.

[21] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, "Quantum state tomography via compressed sensing," *Physical review letters*, vol. 105, no. 15, p. 150401, 2010.

[22] V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky, "Latent variable graphical model selection via convex optimization," *The Annals of Statistics*, vol. 40, no. 4, pp. 1935–1967, 2012.

[23] S. Chatterjee, "Matrix estimation by universal singular value thresholding," *The Annals of Statistics*, vol. 43, no. 1, pp. 177–214, 2014.

[24] Y. Chen, S. Sanghavi, and H. Xu, "Improved graph clustering," *IEEE Transactions on Information Theory*, vol. 60, no. 10, pp. 6440–6455, 2014.

[25] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Review*, vol. 52, no. 471, 2010.

[26] A. S. Bandeira, E. Dobriban, D. G. Mixon, and W. F. Sawin, "Certifying the restricted isometry property is hard," *IEEE Transactions on Information Theory*, vol. 59, no. 6, pp. 3448–3450, 2013.

[27] E. J. Candès and Y. Plan, "Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements," *IEEE Transactions on Information Theory*, vol. 57, no. 4, pp. 2342–2359, 2011.

[28] Y. Chen, Y. Chi, and A. Goldsmith, "Exact and stable covariance estimation from quadratic sampling via convex programming," *IEEE*

- Transactions on Information Theory*, vol. 61, no. 7, pp. 4034–4059, July 2015.
- [29] T. T. Cai and A. Zhang, “ROP: Matrix recovery via rank-one projections,” *The Annals of Statistics*, vol. 43, no. 1, pp. 102–138, 2015.
- [30] B. Recht, W. Xu, and B. Hassibi, “Null space conditions and thresholds for rank minimization,” *Mathematical Programming*, vol. 127, no. 1, pp. 175–202, 2011.
- [31] Y. Chen, “Incoherence-optimal matrix completion,” *IEEE Transactions on Information Theory*, vol. 61, no. 5, pp. 2909–2923, 2015.
- [32] B. Recht, “A simpler approach to matrix completion,” *Journal of Machine Learning Research*, vol. 12, pp. 3413–3430, February 2011.
- [33] R. H. Keshavan, A. Montanari, and S. Oh, “Matrix completion from a few entries,” *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2980–2998, June 2010.
- [34] V. Koltchinskii, K. Lounici, and A. B. Tsybakov, “Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion,” *The Annals of Statistics*, vol. 39, no. 5, pp. 2302–2329, 2011.
- [35] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, “SDPT3—MATLAB software package for semidefinite programming,” *Optimization methods and software*, vol. 11, no. 1–4, pp. 545–581, 1999.
- [36] J.-F. Cai, E. J. Candès, and Z. Shen, “A singular value thresholding algorithm for matrix completion,” *SIAM Journal on Optimization*, vol. 20, no. 4, pp. 1956–1982, 2010.
- [37] K.-C. Toh and S. Yun, “An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems,” *Pacific Journal of Optimization*, vol. 6, no. 615–640, p. 15, 2010.
- [38] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [39] Z. Lin, M. Chen, L. Wu, and Y. Ma, “The Augmented Lagrange Multiplier Method for Exact Recovery of Corrupted Low-Rank Matrices,” *UIUC Technical Report UILU-ENG-09-2215*, 2009.
- [40] M. Jaggi, “Revisiting Frank-Wolfe: Projection-free sparse convex optimization,” in *International Conference on Machine Learning*, 2013, pp. 427–435.
- [41] R. M. Freund, P. Grigas, and R. Mazumder, “An extended Frank-Wolfe method with in-face directions, and its application to low-rank matrix completion,” *SIAM Journal on Optimization*, vol. 27, no. 1, pp. 319–346, 2017.
- [42] N. Rao, P. Shah, and S. Wright, “Forward-backward greedy algorithms for atomic norm regularization,” *IEEE Transactions on Signal Processing*, vol. 63, no. 21, pp. 5798–5811, 2015.
- [43] N. Boyd, G. Schiebinger, and B. Recht, “The alternating descent conditional gradient method for sparse inverse problems,” *SIAM Journal on Optimization*, vol. 27, no. 2, pp. 616–639, 2017.
- [44] V. Cevher, S. Becker, and M. Schmidt, “Convex optimization for big data: Scalable, randomized, and parallel algorithms for big data analytics,” *IEEE Signal Processing Magazine*, vol. 31, no. 5, pp. 32–43, 2014.
- [45] N. Halko, P.-G. Martinsson, and J. A. Tropp, “Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions,” *SIAM Review*, vol. 53, no. 2, pp. 217–288, 2011.
- [46] Z. Allen-Zhu, E. Hazan, W. Hu, and Y. Li, “Linear convergence of a Frank-Wolfe type algorithm over trace-norm balls,” in *Advances in Neural Information Processing Systems*, 2017, pp. 6192–6201.
- [47] A. Yurtsever, M. Udell, J. Tropp, and V. Cevher, “Sketchy decisions: Convex low-rank matrix optimization with optimal storage,” in *Artificial Intelligence and Statistics*, 2017, pp. 1188–1196.
- [48] S. Burer and R. D. C. Monteiro, “Local minima and convergence in low-rank semidefinite programming,” *Mathematical Programming*, vol. 103, no. 3, pp. 427–444, 2005.
- [49] Y. Chen and M. J. Wainwright, “Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees,” *arXiv preprint 1509.03025*, 2015.
- [50] R. Sun and Z.-Q. Luo, “Guaranteed matrix completion via non-convex factorization,” *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 6535–6579, 2016.
- [51] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht, “Low-rank solutions of linear matrix equations via Procrustes flow,” in *International Conference on Machine Learning*, 2016, pp. 964–973.
- [52] Q. Zheng and J. Lafferty, “A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements,” in *Advances in Neural Information Processing Systems*, 2015, pp. 109–117.
- [53] —, “Convergence analysis for rectangular matrix completion using Burer-Monteiro factorization and gradient descent,” *arXiv preprint arXiv:1605.07051*, 2016.
- [54] R. Ge, J. D. Lee, and T. Ma, “Matrix completion has no spurious local minimum,” in *Advances in Neural Information Processing Systems* 29, 2016.
- [55] S. Bhojanapalli, B. Neyshabur, and N. Srebro, “Global optimality of local search for low rank matrix recovery,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3873–3881.
- [56] X. Li, S. Ling, T. Strohmer, and K. Wei, “Rapid, robust, and reliable blind deconvolution via nonconvex optimization,” to appear in *Applied and Computational Harmonic Analysis*, 2018.
- [57] J. Sun, Q. Qu, and J. Wright, “A geometric analysis of phase retrieval,” in *IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2016, pp. 2379–2383.
- [58] Y. Chen and E. J. Candès, “Solving random quadratic systems of equations is nearly as easy as solving linear systems,” *Communications on Pure and Applied Mathematics*, vol. 5, no. 70, pp. 822–883, 2017.
- [59] P. Jain, P. Netrapalli, and S. Sanghavi, “Low-rank matrix completion using alternating minimization,” in *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, 2013, pp. 665–674.
- [60] M. Hardt, “Understanding alternating minimization for matrix completion,” in *IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2014, pp. 651–660.
- [61] B. D. Haeffele and R. Vidal, “Global optimality in tensor factorization, deep learning, and beyond,” *arXiv preprint arXiv:1506.07540*, 2015.
- [62] P. Jain, R. Meka, and I. S. Dhillon, “Guaranteed rank minimization via singular value projection,” in *Advances in Neural Information Processing Systems*, 2010, pp. 937–945.
- [63] P. Jain and P. Netrapalli, “Fast exact matrix completion with finite samples,” in *Conference on Learning Theory*, 2015, pp. 1007–1034.
- [64] P. Netrapalli, U. Niranjan, S. Sanghavi, A. Anandkumar, and P. Jain, “Non-convex robust PCA,” in *Advances in Neural Information Processing Systems*, 2014, pp. 1107–1115.
- [65] E. J. Candès, X. Li, and M. Soltanolkotabi, “Phase retrieval via Wirtinger flow: Theory and algorithms,” *IEEE Transactions on Information Theory*, vol. 61, no. 4, pp. 1985–2007, 2015.
- [66] H. Zhang, Y. Chi, and Y. Liang, “Provable non-convex phase retrieval with outliers: Median truncated Wirtinger flow,” in *International Conference on Machine Learning (ICML)*, 2016, pp. 1022–1031.
- [67] C. Ma, K. Wang, Y. Chi, and Y. Chen, “Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution,” *arXiv preprint arXiv:1711.10467*, 2017.
- [68] Y. Li, C. Ma, Y. Chen, and Y. Chi, “Nonconvex matrix factorization from rank-one measurements,” *arXiv preprint arXiv:1802.06286*, 2018.
- [69] A. S. Bandeira, N. Boumal, and V. Voroninski, “On the low-rank approach for semidefinite programs arising in synchronization and community detection,” in *29th Annual Conference on Learning Theory*, 2016, pp. 361–382.
- [70] R. Ge, C. Jin, and Y. Zheng, “No spurious local minima in nonconvex low rank problems: A unified geometric analysis,” *arXiv preprint arXiv:1704.00708*, 2017.
- [71] X. Yi, D. Park, Y. Chen, and C. Caramanis, “Fast algorithms for robust PCA via gradient descent,” in *Neural Information Processing Systems Conference (NIPS)*, 2016.
- [72] Y. Nesterov and B. T. Polyak, “Cubic regularization of newton method and its global performance,” *Mathematical Programming*, vol. 108, no. 1, pp. 177–205, 2006.
- [73] J. Sun, Q. Qu, and J. Wright, “Complete dictionary recovery using nonconvex optimization,” in *Proceedings of the 32nd International Conference on Machine Learning*, 2015, pp. 2351–2360.
- [74] R. Ge, F. Huang, C. Jin, and Y. Yuan, “Escaping from saddle points — online stochastic gradient for tensor decomposition,” in *Conference on Learning Theory*, 2015, pp. 797–842.
- [75] C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, “How to escape saddle points efficiently,” in *International Conference on Machine Learning (ICML)*, 2017.
- [76] Y. Carmon, J. C. Duchi, O. Hinder, and A. Sidford, “Accelerated methods for non-convex optimization,” *arXiv preprint arXiv:1611.00756*, 2016.
- [77] N. Agarwal, Z. Allen-Zhu, B. Bullins, E. Hazan, and T. Ma, “Finding approximate local minima for nonconvex optimization in linear time,” *arXiv preprint arXiv:1611.01146*, 2016.
- [78] B. Yang, “Projection approximation subspace tracking,” *IEEE Transactions on Signal processing*, vol. 43, no. 1, pp. 95–107, 1995.
- [79] X. Li, Z. Wang, J. Lu, R. Arora, J. Haupt, H. Liu, and T. Zhao, “Symmetry, saddle points, and global geometry of nonconvex matrix factorization,” *arXiv preprint arXiv:1612.09296*, 2016.

- [80] Y. Chen and Y. Chi, "Robust spectral compressed sensing via structured matrix completion," *IEEE Transactions on Information Theory*, vol. 60, no. 10, pp. 6576–6601, 2014.
- [81] K. H. Jin, D. Lee, and J. C. Ye, "A general framework for compressed sensing and parallel mri using annihilating filter based low-rank hankel matrix," *IEEE Transactions on Computational Imaging*, vol. 2, no. 4, pp. 480–495, 2016.
- [82] Y. Chi, L. Scharf, A. Pezeshki, and A. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2182–2195, May 2011.
- [83] W. Liao and A. Fannjiang, "Music for single-snapshot spectral estimation: Stability and super-resolution," *Applied and Computational Harmonic Analysis*, vol. 40, no. 1, pp. 33–67, 2016.
- [84] J.-F. Cai, T. Wang, and K. Wei, "Spectral compressed sensing via projected gradient descent," *arXiv preprint arXiv:1707.09726*, 2017.
- [85] H. Qiao and P. Pal, "Gridless line spectrum estimation and low-rank Toeplitz matrix compression using structured samplers: A regularization-free approach," *IEEE Transactions on Signal Processing*, vol. 65, no. 9, pp. 2221–2236, 2017.
- [86] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: Some first steps," *Social Networks*, vol. 5, pp. 109–137, 1983.
- [87] A. Condon and R. M. Karp, "Algorithms for graph partitioning on the planted partition model," *Random Structures and Algorithms*, vol. 18, no. 2, pp. 116–140, 2001.
- [88] E. Abbe, "Community detection and the stochastic block model: recent developments," *Journal of Machine Learning Research*, to appear, 2017. [Online]. Available: http://www.princeton.edu/~eabbe/publications/sbm_jmlr_4.pdf
- [89] Y. Chen and J. Xu, "Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices," *Journal of Machine Learning Research*, vol. 17, no. 27, pp. 1–57, 2016.
- [90] F. M. Harper and J. A. Konstan, "The MovieLens datasets: History and context," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 5, no. 4, p. 19, 2016.
- [91] D. Park, A. Kyriillidis, C. Caramanis, and S. Sanghavi, "Finding low-rank solutions to matrix problems, efficiently and provably," *arXiv preprint arXiv:1606.03168*, 2016.
- [92] Y. Chen, X. Yi, and C. Caramanis, "Convex and nonconvex formulations for mixed regression with two components: Minimax optimal rates," *IEEE Transactions on Information Theory*, to appear, 2017.
- [93] N. Vaswani, S. Nayer, and Y. C. Eldar, "Low-rank phase retrieval," *IEEE Transactions on Signal Processing*, vol. 65, no. 15, pp. 4059–4074, 2017.
- [94] M. A. Davenport, Y. Plan, E. van den Berg, and M. Wootters, "1-bit matrix completion," *Information and Inference*, vol. 3, no. 3, pp. 189–223, 2014.
- [95] J. Lafond, "Low rank matrix completion with exponential family noise," in *Conference on Learning Theory*, 2015, pp. 1224–1243.
- [96] Y. Lu and S. N. Negahban, "Individualized rank aggregation using nuclear norm regularization," in *53rd Annual Allerton Conference on Communication, Control, and Computing*. IEEE, 2015, pp. 1473–1479.

Yuejie Chi (yuejie.chi@cmu.edu) (S'09-M'12-SM'17) received the Ph.D. degree in Electrical Engineering from Princeton University in 2012, and the B.E. (Hon.) degree in Electrical Engineering from Tsinghua University in 2007. She is currently an Associate Professor with the department of Electrical and Computer Engineering at Carnegie Mellon University. She is the recipient of the IEEE Signal Processing Society Young Author Best Paper Award, NSF CAREER Award, Young Investigator Program Awards from AFOSR and ONR, and Google Faculty Research Award. Her research interests include signal processing, statistical inference, machine learning and their applications in high-dimensional data analysis, network inference, active sensing and bioinformatics.

Yudong Chen (yudong.chen@cornell.edu) is an Assistant Professor with the School of Operations Research and Information Engineering at Cornell University. He obtained his Ph.D. degree in Electrical and Computer Engineering in 2013 from The University of Texas at Austin, and M.S. and B.S. degrees in Control Science and Engineering from Tsinghua University. He was a postdoc in the Electrical Engineering and Computer Sciences department at the University of California, Berkeley. He has served on the senior program committees of AAAI and AISTATS. His research work lies in machine learning, high-dimensional statistics, and optimization, with applications in network scheduling, wireless communication, and financial systems.