

Efficient Matrix Sensing Using Rank-1 Gaussian Measurements

Kai Zhong¹, Prateek Jain², and Inderjit S. Dhillon^{1(✉)}

¹ University of Texas at Austin, Austin, USA
zhongkai@ices.utexas.edu, nderjit@cs.utexas.edu

² Microsoft Research, Bangalore, India
prajain@microsoft.com

Abstract. In this paper, we study the problem of low-rank matrix sensing where the goal is to reconstruct a matrix *exactly* using a small number of linear measurements. Existing methods for the problem either rely on measurement operators such as random element-wise sampling which cannot recover arbitrary low-rank matrices or require the measurement operator to satisfy the Restricted Isometry Property (RIP). However, RIP based linear operators are generally full rank and require large computation/storage cost for both measurement (encoding) as well as reconstruction (decoding).

In this paper, we propose simple rank-one Gaussian measurement operators for matrix sensing that are significantly less expensive in terms of memory and computation for both encoding and decoding. Moreover, we show that the matrix can be reconstructed *exactly* using a simple alternating minimization method as well as a nuclear-norm minimization method. Finally, we demonstrate the effectiveness of the measurement scheme vis-a-vis existing RIP based methods.

Keywords: Matrix sensing · Matrix completion · Inductive learning · Alternating minimization

1 Introduction

In this paper, we consider the matrix sensing problem, where the goal is to recover a low-rank matrix using a small number of linear measurements. The matrix sensing process contains two phases: a) compression phase (encoding), and b) reconstruction phase (decoding).

In the compression phase, a sketch/measurement of the given low-rank matrix is obtained by applying a linear operator $\mathcal{A} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$. That is, given a rank- k matrix, $W_* \in \mathbb{R}^{d_1 \times d_2}$, its linear measurements are computed by:

$$\mathbf{b} = \mathcal{A}(W_*) = [\text{Tr}(A_1^T W_*) \quad \text{Tr}(A_2^T W_*) \quad \dots \quad \text{Tr}(A_m^T W_*)]^T, \quad (1)$$

Electronic Supplementary Materials The online version of this chapter (doi:[10.1007/978-3-319-24486-0_1](https://doi.org/10.1007/978-3-319-24486-0_1)) contains supplementary material, which is available to authorized users.

where $\{A_l \in \mathbb{R}^{d_1 \times d_2}\}_{l=1,2,\dots,m}$ parameterize the linear operator \mathcal{A} and Tr denotes the trace operator. Then, in the reconstruction phase, the underlying low-rank matrix is reconstructed using the given measurements \mathbf{b} . That is, W_* is obtained by solving the following optimization problem:

$$\min_W \text{rank}(W) \quad \text{s.t.} \quad \mathcal{A}(W) = \mathbf{b}. \quad (2)$$

The matrix sensing problem is a matrix generalization of the popular compressive sensing problem and has several real-world applications in the areas of system identification and control, Euclidean embedding, and computer vision (see [20] for a detailed list of references).

Naturally, the design of the measurement operator \mathcal{A} is critical for the success of matrix sensing as it dictates cost of both the compression as well as the reconstruction phase. Most popular operators for this task come from a family of operators that satisfy a certain Restricted Isometry Property (RIP). However, these operators require each A_l , that parameterizes \mathcal{A} , to be a full rank matrix. That is, cost of compression as well as storage of \mathcal{A} is $O(md_1d_2)$, which is infeasible for large matrices. Reconstruction of the low-rank matrix W_* is also expensive, requiring $O(md_1d_2 + d_1^2d_2)$ computation steps. Moreover, m is typically at least $O(k \cdot d \log(d))$ where $d = d_1 + d_2$. But, these operators are universal, i.e., every rank- k W_* can be compressed and recovered using such RIP based operators.

In this paper, we seek to reduce the computational/storage cost of such operators but at the cost of the universality property. That is, we propose to use simple rank-one operators, i.e., where each A_l is a rank one matrix. We show that using similar number of measurements as the RIP operators, i.e., $m = O(k \cdot d \log(d))$, we can recover a fixed rank W_* exactly.

In particular, we propose two measurement schemes: a) *rank-one independent* measurements, b) *rank-one dependent* measurements. In rank-one independent measurement, we use $A_l = \mathbf{x}_l \mathbf{y}_l^T$, where $\mathbf{x}_l \in \mathbb{R}^{d_1}$, $\mathbf{y}_l \in \mathbb{R}^{d_2}$ are both sampled from zero mean sub-Gaussian product distributions, i.e., each element of \mathbf{x}_l and \mathbf{y}_l is sampled from a fixed zero-mean univariate sub-Gaussian distribution. Rank-one dependent measurements combine the above rank-one measurements with element-wise sampling, i.e., $A_l = \mathbf{x}_{i_l} \mathbf{y}_{j_l}^T$ where $\mathbf{x}_{i_l}, \mathbf{y}_{j_l}$ are sampled as above. Also, $(i_l, j_l) \in [n_1] \times [n_2]$ is a randomly sampled index, where $n_1 \geq d_1$, $n_2 \geq d_2$. These measurements can also be viewed as the inductive version of the matrix completion problem (see Section 3), where \mathbf{x}_i represents features of the i -th user (row) and \mathbf{y}_j represents features of the j -th movie (column). In fact, an additional contribution of our work is that we can show that the inductive matrix completion problem can also be solved using only $O(k(d_1 + d_2) \log(d_1 + d_2) \log(n_1 + n_2))$ samples, as long as X, Y, W_* satisfy certain incoherence style assumptions (see Section 5 for more details)¹.

Next, we provide two recovery algorithms for both of the above measurement operators: a) alternating minimization, b) nuclear-norm minimization.

¹ A preliminary version of this work appeared in [11]. Since the above mentioned work, several similar rank-one measurement operators have been studied [2, 15, 25].

Table 1. Comparison of sample complexity and computational complexity for different approaches and different measurements

METHODS		SAMPLE COMPLEXITY	COMPUTATIONAL COMPLEXITY
ALS	RANK-1 INDEP.	$O(k^4 \beta^2 d \log^2 d)$	$O(kdm)$
	RANK-1 DEP.	$O(k^4 \beta^2 d \log d)$	$O(dm + knd)$
	RIP	$O(k^4 d \log d)$	$O(d^2 m)$
NUCLEAR	RANK-1 INDEP.	$O(kd)$	$O(kdm)$
	RANK-1 DEP.	$O(kd \log n \log d)$	$O(k^2 m + knd)$
	RIP	$O(kd \log d)$	$O(d^2 m)$

Note that, in general, the recovery problem (2) is NP-hard to solve. However, for the RIP based operators, both alternating minimization as well as nuclear norm minimization methods are known to solve the problem exactly in polynomial time [13]. Note that the existing analysis of both the methods crucially uses RIP and hence does not extend to the proposed operators.

We show that if $m = O(k^4 \cdot \beta^2 \cdot (d_1 + d_2) \log^2(d_1 + d_2))$, where β is the condition number of W_* then alternating minimization for the rank-one independent measurements recovers W_* in time $O(kdm)$, where $d = d_1 + d_2$. Similarly, if $m = O(k \cdot (d_1 + d_2) \cdot \log^2(d_1 + d_2))$ then the nuclear norm minimization based method recovers W_* in time $O(d^2 m)$ in the worst case. Note that alternating minimization has slightly higher sample complexity requirement but is significantly faster than the nuclear norm minimization method. Due to this, most practical low-rank estimation systems in fact use alternating minimization method for recovery. We obtain similar results for the rank-one dependent measurements.

We summarize the sample complexity and computational complexity for different approaches and different measurements in Table 1. In the table, ALS refers to alternating least squares, i.e., alternating minimization. For the symbols, $d = d_1 + d_2$, $n = n_1 + n_2$, and k is the maximum of rank estimate used as an input in the nuclear-norm solver [10][9]. This number can be as large as $\min\{d_1, d_2\}$ when we have no prior knowledge of the rank.

Paper Organization: We summarize related work in Section 2. In Section 3 we formally introduce the matrix sensing problem and our proposed rank-one measurement operators. In Section 4, we present the alternating minimization method for matrix reconstruction. We then present a *generic* analysis for alternating minimization when applied to the proposed rank-one measurement operators. We present the nuclear-norm minimization based method in Section 5 and present its analysis for the rank-one dependent measurements. Finally, we provide empirical validation of our methods in Section 6.

2 Related Work

Matrix Sensing: Matrix sensing[20][12][16] is a generalization of the popular compressive sensing problem for the sparse vectors and has applications in several

domains such as control, vision etc. [20] introduced measurement operators that satisfy RIP and showed that using only $O(kd \log d)$ measurements, a rank- k $W_* \in \mathbb{R}^{d_1 \times d_2}$ can be recovered. Recently, a set of universal Pauli measurements, used in quantum state tomography, have been shown to satisfy the RIP condition [17]. These measurement operators are Kronecker products of 2×2 matrices, thus, they have appealing computation and memory efficiency. In concurrent work, [15][2] also proposed an independent rank-one measurement using nuclear-norm minimization. In contrast, we use two different measurement operators and show that the popular alternating minimization method also solves the problem exactly.

Matrix Completion and Inductive Matrix Completion: Matrix completion [3][14][13] is a special case of rank-one matrix sensing problem when the operator takes a subset of the entries. However, to guarantee exact recovery, the target matrix has to satisfy the incoherence condition. Using our rank-one Gaussian operators, we don't require any condition on the target matrix. For inductive matrix completion (IMC), which is a generalization of matrix completion utilizing movies' and users' features, the authors of [22] provided the theoretical recovery guarantee for nuclear-norm minimization. In this paper, we show that IMC is equivalent to the matrix sensing problem using dependent rank-one measurements, and provide a similar result for nuclear-norm based methods for IMC but eliminate a "joint" incoherent condition on the rank-one measurements and an upper bound condition on the sample complexity. Moreover, we give a theoretical guarantee using alternating minimization methods.

Alternating Minimization: Although nuclear-norm minimization enjoys nice recovery guarantees, it usually doesn't scale well. In practice, alternating minimization is employed to solve problem (2) by assuming the rank is known. Alternating minimization solves two least square problems alternatively in each iteration, thus is very computationally efficient[23]. Although widely used in practice, its theoretical guarantees are relatively less understood due to non-convexity. [13] first showed optimality of alternating minimization in the matrix sensing/low-rank estimation setting under the RIP setting. Subsequently, several other papers have also shown global convergence guarantees for alternating minimization, e.g. matrix completion [8][7], robust PCA [18] and dictionary learning [1]. In this paper, we provide a *generic* analysis for alternating minimization applied to the proposed rank-one measurement operators. Our results distill out certain key problem specific properties that would imply global optimality of alternating minimization. We then show that the rank-one Gaussian measurements satisfy those properties.

3 Problem Formulation

The goal of matrix sensing is to design a linear operator $\mathcal{A} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ and a recovery algorithm so that a low-rank matrix $W_* \in \mathbb{R}^{d_1 \times d_2}$ can be recovered exactly using $\mathcal{A}(W_*)$. In this paper, we focus on rank-one measurement operators, $A_l = \mathbf{x}_l \mathbf{y}_l^T$, and call such problems as *Low-Rank matrix estimation using*

Rank One Measurements (LRROM): recover the rank- k matrix $W_* \in \mathbb{R}^{d_1 \times d_2}$ by using rank-1 measurements of the form:

$$\mathbf{b} = [\mathbf{x}_1^T W_* \mathbf{y}_1 \quad \mathbf{x}_2^T W_* \mathbf{y}_2 \quad \dots \quad \mathbf{x}_m^T W_* \mathbf{y}_m]^T, \quad (3)$$

where $\mathbf{x}_l, \mathbf{y}_l$ are “feature” vectors and are provided along with the measurements \mathbf{b} .

We propose two different kinds of rank-one measurement operators based on Gaussian distribution.

3.1 Rank-one Independent Gaussian Operator

Our first measurement operator is a simple rank-one Gaussian operator, $\mathcal{A}_{GI} = [A_1 \dots A_m]$, where, $\{A_l = \mathbf{x}_l \mathbf{y}_l^T\}_{l=1,2,\dots,m}$, and $\mathbf{x}_l, \mathbf{y}_l$ are sampled i.i.d. from spherical Gaussian distribution.

3.2 Rank-one Dependent Gaussian Operator

Our second operator can introduce certain “dependencies” in our measurement and has in fact interesting connections to the matrix completion problem. We provide the operator as well as the connection to matrix completion in this sub-section. To generate the rank-one dependent Gaussian operator, we first sample two Gaussian matrices $X \in \mathbb{R}^{n_1 \times d_1}$ and $Y \in \mathbb{R}^{n_2 \times d_2}$, where each entry of both X, Y is sampled independently from Gaussian distribution and $n_1 \geq C d_1$, $n_2 \geq C d_2$ for a global constant $C \geq 1$. Then, the Gaussian dependent operator $\mathcal{A}_{GD} = [A_1, \dots, A_m]$ where $\{A_l = \mathbf{x}_{i_l} \mathbf{y}_{j_l}^T\}_{(i_l, j_l) \in \Omega}$. Here \mathbf{x}_i^T is the i -th row of X and \mathbf{y}_j^T is the j -th row of Y . Ω is a uniformly random subset of $[n_1] \times [n_2]$ such that $E[|\Omega|] = m$. For simplicity, we assume that each entry $(i_l, j_l) \in [n_1] \times [n_2]$ is sampled i.i.d. with probability $p = m/(n_1 \times n_2)$. Therefore, the measurements using the above operator are given by: $b_l = \mathbf{x}_{i_l}^T W \mathbf{y}_{j_l}, (i_l, j_l) \in \Omega$.

Connections to Inductive Matrix Completion (IMC): Note that the above measurements are inspired by matrix completion style sampling operator. However, here we first multiply W with X, Y and then sample the obtained matrix XWY^T . In the domain of recommendation systems (say user-movie system), the corresponding reconstruction problem can also be thought as the inductive matrix completion problem. That is, let there be n_1 users, n_2 movies, X represents user features, and Y represents the movie features. Then, the true ratings matrix is given by $R = XWY^T \in \mathbb{R}^{n_1 \times n_2}$.

That is, given the user/movie feature vectors $\mathbf{x}_i \in \mathbb{R}^{d_1}$ for $i = 1, 2, \dots, n_1$ and $\mathbf{y}_j \in \mathbb{R}^{d_2}$ for $j = 1, 2, \dots, n_2$, our goal is to recover a rank- k matrix W_* of size $d_1 \times d_2$ from a few observed entries $R_{ij} = \mathbf{x}_i^T W_* \mathbf{y}_j$, for $(i, j) \in \Omega \subset [n_1] \times [n_2]$. Because of the equivalence between the dependent rank-one measurements and the entries of the rating matrix, in the rest of the paper, we will use $\{R_{ij}\}_{(i,j) \in \Omega}$ as the dependent rank-one measurements.

Now, if we can reconstruct W_* from the above measurements, we can predict ratings *inductively* for new users/movies, provided their feature vectors are given.

Hence, our reconstruction procedure also solves the IMC problem. However, there is a key difference: in matrix sensing, we can select X, Y according to our convenience, while in IMC, X and Y are provided a priori. But, for general X, Y one cannot solve the problem because if say $R = XW_*Y^T$ is a 1-sparse matrix, then W_* cannot be reconstructed even with a large number of samples.

Interestingly, our proof for reconstruction using nuclear-norm based method does not require Gaussian X, Y . Instead, we can distill out two key properties of R, X, Y ensuring that using only $O(k(d_1 + d_2) \log(d_1 + d_2) \log(n_1 + n_2))$ samples, we can reconstruct W_* . Note that a direct application of matrix completion results [3][4] would require $O(k(n_1 + n_2) \log(n_1 + n_2))$ samples which can be much large if $n_1 \gg d_1$ or $n_2 \gg d_2$. See Section 5 for more details on the assumptions that we require for the nuclear-norm minimization method to solve the IMC problem exactly.

4 Rank-one Matrix Sensing via Alternating Minimization

We now present the alternating minimization approach for solving the reconstruction problem (2) with rank-one measurements (3). Since W to be recovered is restricted to have at most rank- k , (2) can be reformulated as the following non-convex optimization problem:

$$\min_{U \in \mathbb{R}^{d_1 \times k}, V \in \mathbb{R}^{d_2 \times k}} \sum_{l=1}^m (b_l - \mathbf{x}_l^T U V^T \mathbf{y}_l)^2. \quad (4)$$

Alternating minimization is an iterative procedure that alternately optimizes for U and V while keeping the other factor fixed. As optimizing for U (or V) involves solving just a least squares problem, so each individual iteration of the algorithm is linear in matrix dimensions. For the rank-one measurement operator, we use a particular initialization method to initialize U (see line 3 of Algorithm 1). See Algorithm 1 for a pseudo-code of the algorithm.

4.1 General Theoretical Guarantee for Alternating Minimization

As mentioned above, (4) is non-convex in U, V and hence standard analysis would only ensure convergence to a local minima. However, [13] recently showed that the alternating minimization method in fact converges to the global minima of two low-rank estimation problems: matrix sensing with RIP matrices and matrix completion.

The rank-one operator given above does not satisfy RIP (see Definition 1), even when the vectors $\mathbf{x}_l, \mathbf{y}_l$ are sampled from the normal distribution (see Claim 4.2). Furthermore, each measurement need not reveal exactly one entry of W_* as in the case of matrix completion. Hence, the proof of [13] does not apply directly. However, inspired by the proof of [13], we distill out three key properties that the operator should satisfy, so that alternating minimization would converge to the global optimum.

Algorithm 1. AltMin-LRRM : Alternating Minimization for LRRM

-
- 1: **Input:** Measurements: \mathbf{b}_{all} , Measurement matrices: \mathcal{A}_{all} , Number of iterations: H
 - 2: Divide $(\mathcal{A}_{all}, \mathbf{b}_{all})$ into $2H + 1$ sets (each of size m) with h -th set being $\mathcal{A}^h = \{A_1^h, A_2^h, \dots, A_m^h\}$ and $\mathbf{b}^h = [b_1^h \ b_2^h \ \dots \ b_m^h]^T$
 - 3: **Initialization:** $U_0 = \text{top-}k$ left singular vectors of $\frac{1}{m} \sum_{l=1}^m b_l^0 A_l^0$
 - 4: **for** $h = 0$ to $H - 1$ **do**
 - 5: $b \leftarrow b^{2h+1}, \mathcal{A} \leftarrow \mathcal{A}^{2h+1}$
 - 6: $\widehat{V}_{h+1} \leftarrow \operatorname{argmin}_{V \in \mathbb{R}^{d_2 \times k}} \sum_l (b_l - \mathbf{x}_l^T U_h V^T \mathbf{y}_l)^2$
 - 7: $V_{h+1} = QR(\widehat{V}_{h+1})$ // orthonormalization of \widehat{V}_{h+1}
 - 8: $b \leftarrow b^{2h+2}, \mathcal{A} \leftarrow \mathcal{A}^{2h+2}$
 - 9: $\widehat{U}_{h+1} \leftarrow \operatorname{argmin}_{U \in \mathbb{R}^{d_1 \times k}} \sum_l (b_l - \mathbf{x}_l^T U V_{h+1}^T \mathbf{y}_l)^2$
 - 10: $U_{h+1} = QR(\widehat{U}_{h+1})$ // orthonormalization of \widehat{U}_{h+1}
 - 11: **end for**
 - 12: **Output:** $W_H = U_H (\widehat{V}_H)^T$
-

Theorem 1. Let $W_* = U_* \Sigma_* V_*^T \in \mathbb{R}^{d_1 \times d_2}$ be a rank- k matrix with k -singular values $\sigma_*^1 \geq \sigma_*^2 \cdots \geq \sigma_*^k$. Also, let $\mathcal{A} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ be a linear measurement operator parameterized by m matrices, i.e., $\mathcal{A} = \{A_1, A_2, \dots, A_m\}$ where $A_l = \mathbf{x}_l \mathbf{y}_l^T$. Let $\mathcal{A}(W)$ be as given by (1).

Now, let \mathcal{A} satisfy the following properties with parameter $\delta = \frac{1}{k^{3/2} \cdot \beta \cdot 100}$ ($\beta = \sigma_*^1 / \sigma_*^k$):

1. **Initialization:** $\|\frac{1}{m} \sum_l b_l A_l - W_*\|_2 \leq \|W_*\|_2 \cdot \delta$.
2. **Concentration of operators B_x, B_y :** Let $B_x = \frac{1}{m} \sum_{l=1}^m (\mathbf{y}_l^T \mathbf{v})^2 \mathbf{x}_l \mathbf{x}_l^T$ and $B_y = \frac{1}{m} \sum_{l=1}^m (\mathbf{x}_l^T \mathbf{u})^2 \mathbf{y}_l \mathbf{y}_l^T$, where $\mathbf{u} \in \mathbb{R}^{d_1}, \mathbf{v} \in \mathbb{R}^{d_2}$ are two unit vectors that are independent of randomness in $\mathbf{x}_l, \mathbf{y}_l, \forall i$. Then the following holds: $\|B_x - I\|_2 \leq \delta$ and $\|B_y - I\|_2 \leq \delta$.
3. **Concentration of operators G_x, G_y :** Let $G_x = \frac{1}{m} \sum_l (\mathbf{y}_l^T \mathbf{v})(\mathbf{y}_l \mathbf{v}_\perp) \mathbf{x}_l \mathbf{x}_l^T$, $G_y = \frac{1}{m} \sum_l (\mathbf{x}_l^T \mathbf{u})(\mathbf{u}_\perp^T \mathbf{x}_l) \mathbf{y}_l \mathbf{y}_l^T$, where $\mathbf{u}, \mathbf{u}_\perp \in \mathbb{R}^{d_1}, \mathbf{v}, \mathbf{v}_\perp \in \mathbb{R}^{d_2}$ are unit vectors, s.t., $\mathbf{u}^T \mathbf{u}_\perp = 0$ and $\mathbf{v}^T \mathbf{v}_\perp = 0$. Furthermore, let $\mathbf{u}, \mathbf{u}_\perp, \mathbf{v}, \mathbf{v}_\perp$ be independent of randomness in $\mathbf{x}_l, \mathbf{y}_l, \forall i$. Then, $\|G_x\|_2 \leq \delta$ and $\|G_y\|_2 \leq \delta$.

Then, after H -iterations of the alternating minimization method (Algorithm 1), we obtain $W_H = U_H V_H^T$ s.t., $\|W_H - W_*\|_2 \leq \epsilon$, where $H \leq 100 \log(\|W_*\|_F / \epsilon)$.

See Supplementary A for a detailed proof. Note that we require intermediate vectors $\mathbf{u}, \mathbf{v}, \mathbf{u}_\perp, \mathbf{v}_\perp$ to be independent of randomness in A_l 's. Hence, we partition \mathcal{A}_{all} into $2H + 1$ partitions and at each step $(\mathcal{A}^h, \mathbf{b}^h)$ and $(\mathcal{A}^{h+1}, \mathbf{b}^{h+1})$ are supplied to the algorithm. This implies that the measurement complexity of the algorithm is given by $m \cdot H = m \log(\|W_*\|_F / \epsilon)$. That is, given $O(m \log(\|(d_1 + d_2)W_*\|_F))$ samples, we can estimate matrix W_H , s.t., $\|W_H - W_*\|_2 \leq \frac{1}{(d_1 + d_2)^c}$, where $c > 0$ is any constant.

4.2 Independent Gaussian Measurements

In this subsection, we consider the rank-one independent measurement operator \mathcal{A}_{GI} specified in Section 3. Now, for this operator \mathcal{A}_{GI} , we show that if

$m = O(k^4 \beta^2 \cdot (d_1 + d_2) \cdot \log^2(d_1 + d_2))$, then w.p. $\geq 1 - 1/(d_1 + d_2)^{100}$, any fixed rank- k matrix W_* can be recovered by AltMin-LRRM (Algorithm 1). Here $\beta = \sigma_*^1/\sigma_*^k$ is the condition number of W_* . That is, using nearly linear number of measurements in d_1, d_2 , one can exactly recover the $d_1 \times d_2$ rank- k matrix W_* .

As mentioned in the previous section, the existing matrix sensing results typically assume that the measurement operator \mathcal{A} satisfies the Restricted Isometry Property (RIP) defined below:

Definition 1. A linear operator $\mathcal{A} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ satisfies RIP iff, for $\forall W$ s.t. $\text{rank}(W) \leq k$, the following holds:

$$(1 - \delta_k) \|W\|_F^2 \leq \|\mathcal{A}(W)\|_F^2 \leq (1 + \delta_k) \|W\|_F^2,$$

where $\delta_k > 0$ is a constant dependent only on k .

Naturally, this begs the question whether we can show that our rank-1 measurement operator \mathcal{A}_{GI} satisfies RIP, so that the existing analysis for RIP based low-rank matrix sensing can be used [13]. We answer this question in the negative, i.e., for $m = O((d_1 + d_2) \log(d_1 + d_2))$, \mathcal{A}_{GI} does not satisfy RIP even for rank-1 matrices (with high probability):

Claim. 4.2. Let $\mathcal{A}_{GI} = \{A_1, A_2, \dots, A_m\}$ be a measurement operator with each $A_l = \mathbf{x}_l \mathbf{y}_l^T$, where $\mathbf{x}_l \in \mathbb{R}^{d_1} \sim \mathcal{N}(0, I)$, $\mathbf{y}_l \in \mathbb{R}^{d_2} \sim \mathcal{N}(0, I)$, $1 \leq l \leq m$. Let $m = O((d_1 + d_2) \log^c(d_1 + d_2))$, for any constant $c > 0$. Then, with probability at least $1 - 1/m^{10}$, \mathcal{A}_{GI} does not satisfy RIP for rank-1 matrices with a constant δ .

See Supplementary B for a detailed proof of the above claim. Now, even though \mathcal{A}_{GI} does not satisfy RIP, we can still show that \mathcal{A}_{GI} satisfies the three properties mentioned in Theorem 1. and hence we can use Theorem 1 to obtain the exact recovery result.

Theorem 2 (Rank-One Independent Gaussian Measurements using ALS). Let $\mathcal{A}_{GI} = \{A_1, A_2, \dots, A_m\}$ be a measurement operator with each $A_l = \mathbf{x}_l \mathbf{y}_l^T$, where $\mathbf{x}_l \in \mathbb{R}^{d_1} \sim \mathcal{N}(0, I)$, $\mathbf{y}_l \in \mathbb{R}^{d_2} \sim \mathcal{N}(0, I)$, $1 \leq l \leq m$. Let $m = O(k^4 \beta^2 (d_1 + d_2) \log^2(d_1 + d_2))$. Then, Property 1, 2, 3 required by Theorem 1 are satisfied with probability at least $1 - 1/(d_1 + d_2)^{100}$.

Proof. Here, we provide a brief proof sketch. See Supplementary B for a detailed proof.

Initialization: Note that,

$$\frac{1}{m} \sum_{l=1}^m b_l \mathbf{x}_l \mathbf{y}_l^T = \frac{1}{m} \sum_{l=1}^m \mathbf{x}_l \mathbf{x}_l^T U_* \Sigma_* V_*^T \mathbf{y}_l \mathbf{y}_l^T = \frac{1}{m} \sum_{l=1}^m Z_l,$$

where $Z_l = \mathbf{x}_l \mathbf{x}_l^T U_* \Sigma_* V_*^T \mathbf{y}_l \mathbf{y}_l^T$. Note that $\mathbb{E}[Z_l] = U_* \Sigma_* V_*^T$. Hence, to prove the initialization result, we need a tail bound for sums of random matrices. To this end, we use Theorem 1.6 in [21]. However, Theorem 1.6 in [21] requires a bounded random variable while Z_l is an unbounded variable. We handle this issue

by clipping Z_l to ensure that its spectral norm is always bounded. Furthermore, by using properties of normal distribution, we can ensure that w.p. $\geq 1 - 1/m^3$, Z_l 's do not require clipping and the new “clipped” variables converge to nearly the same quantity as the original “non-clipped” Z_l 's. See Supplementary B for more details.

Concentration of B_x, B_y, G_x, G_y : Consider $G_x = \frac{1}{m} \sum_{l=1}^m \mathbf{x}_l \mathbf{x}_l^T \mathbf{y}_l^T \mathbf{v} \mathbf{v}_\perp^T \mathbf{y}_l$. As, $\mathbf{v}, \mathbf{v}_\perp$ are unit-norm vectors, $\mathbf{y}_l^T \mathbf{v} \sim \mathcal{N}(0, 1)$ and $\mathbf{v}_\perp^T \mathbf{x}_l \sim \mathcal{N}(0, 1)$. Also, since \mathbf{v} and \mathbf{v}_\perp are orthogonal, $\mathbf{y}_l^T \mathbf{v}$ and $\mathbf{v}_\perp^T \mathbf{x}_l$ are independent variables. Hence, $G_x = \frac{1}{m} \sum_{l=1}^m Z_l$ where $\mathbb{E}[Z_l] = 0$. Here again, we apply Theorem 1.6 in [21] after using a clipping argument. We can obtain the required bounds for B_x, B_y, G_y also in a similar manner.

Note that the clipping procedure ensures that Z_l 's don't need to be clipped with probability $\geq 1 - 1/m^3$ only. That is, we cannot apply the *union* bound to ensure that the concentration result holds *for all* $\mathbf{v}, \mathbf{v}_\perp$. Hence, we need a fresh set of measurements after each iteration to ensure concentration.

Global optimality of the rate of convergence of the Alternating Minimization procedure for this problem now follows directly by using Theorem 1. We would like to note that while the above result shows that the \mathcal{A}_{GI} operator is almost as powerful as the RIP based operators for matrix sensing, there is one critical drawback: while RIP based operators are universal that is they can be used to recover any rank- k W_* , \mathcal{A}_{GI} needs to be resampled for each W_* . We believe that the two operators are at two extreme ends of randomness vs universality trade-off and intermediate operators with higher success probability but using larger number of random bits should be possible.

4.3 Dependent Gaussian Measurements

For the dependent Gaussian measurements, the alternating minimization formulation is given by:

$$\min_{U \in \mathbb{R}^{d_1 \times k}, V \in \mathbb{R}^{d_2 \times k}} \sum_{(i,j) \in \Omega} (\mathbf{x}_i^T U V^T \mathbf{y}_j - R_{ij})^2. \quad (5)$$

Here again, we can solve the problem by alternatively optimizing for U and V . Later in Section 4.4, we show that using such dependent measurements leads to a faster recovery algorithm when compared to the recovery algorithm for independent measurements.

Note that both the measurement matrices X, Y can be thought of as orthonormal matrices. The reason being, $XW_*Y^T = U_X \Sigma_X V_X^T W_* V_Y \Sigma_Y U_Y^T$, where $X = U_X \Sigma_X V_X^T$ and $Y = U_Y \Sigma_Y V_Y^T$ is the SVD of X, Y respectively. Hence, $R = XW_*Y^T = U_X (\Sigma_X V_X^T W_* V_Y \Sigma_Y) U_Y^T$. Now U_X, U_Y can be treated as the true “X”, “Y” matrices and $W_* \leftarrow (\Sigma_X V_X^T W_* V_Y \Sigma_Y)$ can be thought of as W_* . Then the “true” W_* can be recovered using the obtained W_H as: $W_H \leftarrow V_X \Sigma_X^{-1} W_H \Sigma_Y^{-1} V_Y^T$. We also note that such a transformation implies that the condition number of R and that of $W_* \leftarrow (\Sigma_X V_X^T W_* V_Y \Sigma_Y)$ are exactly the same.

Similar to the previous section, we utilize our general theorem for optimality of the LRRM problem to provide a convergence analysis of rank-one Gaussian dependent operators \mathcal{A}_{GD} . We prove if X and Y are random orthogonal matrices, defined in [3], the above mentioned dependent measurement operator \mathcal{A}_{GD} generated from X, Y also satisfies Properties 1, 2, 3 in Theorem 1. Hence, AltMin-LRRM (Algorithm 1) converges to the global optimum in $O(\log(\|W_*\|_F/\epsilon))$ iterations.

Theorem 3 (Rank-One Dependent Gaussian Measurements using ALS). *Let $X_0 \in \mathbb{R}^{n_1 \times d_1}$ and $Y_0 \in \mathbb{R}^{n_2 \times d_2}$ be Gaussian matrices, i.e. every entry is sampled i.i.d from $\mathcal{N}(0, 1)$. Let $X_0 = X \Sigma_X V_X^T$ and $Y_0 = Y \Sigma_Y V_Y^T$ be the thin SVD of X_0 and Y_0 respectively. Then the rank-one dependent operator \mathcal{A}_{GD} formed by X, Y with $m \geq O(k^4 \beta^2 (d_1 + d_2) \log(d_1 + d_2))$ satisfy Property 1, 2, 3 required by Theorem 1 with high probability.*

See Supplementary C for a detailed proof. Interestingly, our proof does not require X, Y to be Gaussian. It instead utilizes only two key properties about X, Y which are given by:

1. **Incoherence:** For some constant μ, c ,

$$\max_{i \in [n]} \|\mathbf{x}_i\|_2^2 \leq \frac{\mu \bar{d}}{n}, \quad (6)$$

where $\bar{d} = \max(d, \log n)$

2. **Averaging Property:** For H different orthogonal matrices $U_h \in \mathbb{R}^{d \times k}$, $h = 1, 2, \dots, H$, the following hold for these U_h 's,

$$\max_{i \in [n]} \|U_h^T \mathbf{x}_i\|_2^2 \leq \frac{\mu_0 \bar{k}}{n}, \quad (7)$$

where μ_0, c are some constants and $\bar{k} = \max(k, \log n)$.

Hence, the above theorem can be easily generalized to solve the inductive matrix completion problem (IMC), i.e., solve (5) for arbitrary X, Y . Moreover, the sample complexity of the analysis would be nearly in $(d_1 + d_2)$, instead of $(n_1 + n_2)$ samples required by the standard matrix completion methods.

The following lemma shows that the above two properties hold w.h.p. for random orthogonal matrices .

Lemma 1. *If $X \in \mathbb{R}^{n \times d}$ is a random orthogonal matrix, then both Incoherence and Averaging properties are satisfied with probability $\geq 1 - (c/n^3) \log n$, where c is a constant.*

The proof of Lemma 1 can be found in Supplementary C.

4.4 Computational Complexity for Alternating Minimization

In this subsection, we briefly discuss the computational complexity for Algorithm 1. For simplicity, we set $d = d_1 + d_2$ and $n = n_1 + n_2$, and in practical implementation, we don't divide the measurements and use the whole measurement operator \mathcal{A} for every iteration. The most time-consuming part of Algorithm 1 is the step for solving the least square problem. Given $U = U_h$, V can be obtained by solving the following linear system,

$$\sum_{l=1}^m \langle V, A_l^T U_h \rangle A_l U_h = \sum_{l=1}^m b_l A_l^T U_h. \quad (8)$$

The dimension of this linear system is kd , which could be large, thus we use conjugate gradient (CG) method to solve it. In each CG iteration, different measurement operators have different computational complexity. For RIP-based full-rank operators, the computational complexity for each CG step is $O(d^2m)$ while it is $O(kdm)$ for rank-one independent operators. However, for rank-one dependent operators, using techniques introduced in [24], we can reduce the per iteration complexity to be $O(kdn + md)$. Furthermore, if $n = d$, the computational complexity of dependent operators is only $O(kd^2 + md)$, which is better than the complexity of rank-one independent operators in an order of k .

5 Rank-one Matrix Sensing via Nuclear Norm Minimization

In this section, we consider solving LRRM by nuclear norm relaxation. We first note that using nuclear norm relaxation, [15] provided the analysis for independent rank-one measurement operators when the underlying matrix is Hermitian. It can be shown that non-Hermitian matrices problem can be transformed to Hermitian cases. Their proof uses the bowling scheme and only requires $O(k(d_1 + d_2))$ measurements for Gaussian case and $O(k(d_1 + d_2) \log(d_1 + d_2) \log(n_1 + n_2))$ measurements for 4-designs case. In this paper, we consider dependent measurement operators which have a similar sample complexity as the independent operators, but less computational complexity and memory footprint than those of independent ones.

The nuclear norm minimization using rank-one dependent Gaussian operator is of form,

$$\begin{aligned} \min \|W\|_* \\ \text{s.t. } \mathbf{x}_i^T W \mathbf{y}_j = R_{ij}, (i, j) \in \Omega. \end{aligned} \quad (9)$$

(9) can be solved exactly by semi-definite programming or approximated by

$$\min_W \sum_{(i,j) \in \Omega} (\mathbf{x}_i^T W \mathbf{y}_j - R_{ij})^2 + \lambda \|W\|_*, \quad (10)$$

where λ is a constant and can be viewed as a Lagrange multiplier.

5.1 Recovery Guarantee for Nuclear-norm Minimization

In this subsection, we show that using rank-one dependent Gaussian operators, the nuclear-norm minimization can recover *any* low-rank matrix exactly with $O(k(d_1 + d_2) \log(d_1 + d_2) \log(n_1 + n_2))$ measurements. We also generalize the theorem to the IMC problem where the feature matrices X and Y can be arbitrary instead of being Gaussian. We show that as long as X, W_*, Y satisfy certain incoherence style properties, the nuclear norm minimization can guarantee exact recovery using only $O(k(d_1 + d_2) \log(d_1 + d_2))$ samples.

We first provide recovery guarantees for our rank-one dependent operator, i.e., when X, Y are sampled from the Gaussian distribution.

Theorem 4 (Rank-one Dependent Gaussian Measurements using Nuclear-norm Minimization). *Let $W_* = U_* \Sigma_* V_*^T \in \mathbb{R}^{d_1 \times d_2}$ be a rank- k matrix. Let $X \in \mathbb{R}^{n_1 \times d_1}$ and $Y \in \mathbb{R}^{n_2 \times d_2}$ be random orthogonal matrices. Assume each $(i, j) \in \Omega$ is sampled from $[n_1] \times [n_2]$ i.i.d.. Then if $m = |\Omega| \geq O(k(d_1 + d_2) \log(d_1 + d_2) \log(n_1 + n_2))$, the minimizer to the problem (9) is unique and equal to W_* with probability at least $1 - c_1(d_1 + d_2)^{-c_2}$, where c_1 and c_2 are universal constants.*

The above theorem is a directly corollary of Theorem 5 combined with Lemma 1. Lemma 1 shows that random orthonormal matrices X, Y (can be generated using Gaussian matrices as stated in Theorem 2) satisfy the requirements of Theorem 5.

Nuclear-norm minimization approach for inductive matrix completion (9) has also been studied by [22]. However, their recovery guarantee holds under a much more restrictive set of assumptions on X, W_*, Y and in fact requires that the number of samples is not only lower bounded by certain quantity but also upper bounded by some other quantity. Our general analysis below doesn't rely on this upper bound. Moreover, their proof also requires a "joint" incoherent condition, i.e., an upper bound on $\max_{i,j} |\mathbf{x}_i^T U_* V_*^T \mathbf{y}_j|$ which is not required by our method; to this end, we use a technique introduced by [5] to bound an $\ell_{\infty,2}$ -norm.

Theorem 5 (Inductive Matrix Completion using Nuclear-norm Minimization). *Let $W_* = U_* \Sigma_* V_*^T \in \mathbb{R}^{d_1 \times d_2}$ be a rank- k matrix. Assume X, Y are orthogonal matrices, and satisfy the following conditions with respect to W_* for some constant μ and μ_0 ,*

$$\begin{aligned} \text{C1. } & \max_{i \in [n_1]} \|\mathbf{x}_i\|_2^2 \leq \frac{\mu d_1}{n_1}, \quad \max_{j \in [n_2]} \|\mathbf{y}_j\|_2^2 \leq \frac{\mu d_2}{n_2}, \\ \text{C2. } & \max_{i \in [n_1]} \|U_*^T \mathbf{x}_i\|_2^2 \leq \frac{\mu_0 k}{n_1}, \quad \max_{j \in [n_2]} \|V_*^T \mathbf{y}_j\|_2^2 \leq \frac{\mu_0 k}{n_2}. \end{aligned}$$

Then if each observed entry $(i, j) \in \Omega$ is sampled from $[n_1] \times [n_2]$ i.i.d. with probability p ,

$$p \geq \max \left\{ \frac{c_0 \mu_0 \mu k d \log(d) \log(n)}{n_1 n_2}, \frac{1}{\min\{n_1, n_2\}^{10}} \right\}, \quad (11)$$

the minimizer to the problem (9) is unique and equal to W_* with probability at least $1 - c_1 d^{-c_2}$, where c_0 , c_1 and c_2 are universal constants, $d = d_1 + d_2$ and $n = n_1 + n_2$.

Note that the first condition **C1** is actually the incoherence condition on X, Y , while the second one **C2** is the incoherence of XU_*, YV_* . Additionally, **C2** is weaker than the *Averaging* property in Lemma 1, as it only asks for one U_* rather than H different U_h 's to satisfy the property.

Proof. We follow the popular proof ideas used by [3][19], that is, finding a dual feasible solution for (9) to certify the uniqueness of the minimizer of (9). Unlike the analysis in [22], we build our dual certificate in the $\mathbb{R}^{d_1 \times d_2}$ matrix space rather than the $\mathbb{R}^{n_1 \times n_2}$ space. This choice makes it easy to follow the analysis in standard matrix completion problem. In Proposition 1 in Supplementary D, we give certain conditions the dual certificate should satisfy for the uniqueness. Then we apply golfing scheme [6] to find such a certificate. When building the dual certificate, we use an $\ell_{\infty, 2}$ -norm adapted from [5]. This enables us to discard the assumption of ‘‘joint’’ incoherence. The details can be found in Supplementary D.


5.2 Computational Complexity for Nuclear-norm Minimization

The optimization for nuclear-norm formulation is much more complex. Recently [10] proposed an active subspace method to solve Problem (10). The computational bottleneck is the approxSVD step and the inner problem step, both of which involve calculating a similar equation as shown on the left hand side of Eq (8). However, the rank of U or V is not fixed in each iteration as that of ALS, and in the worst case, it can be as large as $\min\{d_1, d_2\}$. The computational complexity for this basic operation is shown in Table 1.


6 Experiments

In this section, we demonstrate empirically that our Gaussian rank-one linear operators are significantly more efficient for matrix sensing than the existing RIP based measurement operators. In particular, we apply the two recovery methods namely alternating minimization (ALS) and nuclear norm minimization (Nuclear) to the measurements obtained using three different operators: rank-one independent (Rank1 Indep), rank-one dependent (Rank1 Dep), and a RIP based operator generated using random Gaussian matrices (RIP).

The experiments are conducted on Matlab and the nuclear-norm solver is adapted from the code by [10]. We first generated a random rank-5 signal $W_* \in \mathbb{R}^{50 \times 50}$, and compute $m = 1000$ measurements using different measurement operators. Here, we fix a small $\lambda = 10^{-6}$ for solving Eq (10) in order to exactly recover the matrix. And we set the maximum possible rank $\hat{k} = k$ as the input of the nuclear-norm solver. Figure 1a plots the relative error in recovery, $err = \|W - W^*\|_F^2 / \|W^*\|_F^2$, against computational time required by each method. Clearly, recovery using rank-one measurements requires significantly



(a) Relative error in recovery v.s. computation time



(b) Recovery rate v.s. number of measurements

Fig. 1. Comparison of computational complexity and measurement complexity for different approaches and different operators




Fig. 2. Recovery rate for different matrix dimension d (x-axis) and different number of measurements m (y-axis). The color reflects the recovery rate scaled from 0 to 1. The white color indicates perfect recovery, while the black color denotes failure in all the experiments.

less time compared to the RIP based operator. Moreover, ALS in general seems to be significantly faster than Nuclear methods.

Next, we compare the measurement complexity (m) of each method. Here again, we first generate a random rank-5 signal $W_* \in \mathbb{R}^{50 \times 50}$ and its measurements using different operators. We then measure error in recovery by each of the method and consider success if the relative error $err \leq 0.05$. We repeat the experiment 10 times to obtain the recovery rate (number of success/10) for each value of m (number of measurements). Figure 1b plots the recovery rate of different approaches for different m . Clearly, the rank-one based measurements have similar recovery rate and measurement complexity as the RIP based operators. However, our rank-one operator based methods are much faster than the corresponding methods for the RIP-based measurement scheme.

Finally, in Figure 2, we validate our theoretical analysis on the measurement complexity by showing the recovery rate for different d and m . We fix the rank

$k = 5$, set $d = d_1 = d_2$ and $n_1 = d_1, n_2 = d_2$ for dependent operators. Figure 3 plots the recovery rate for various d and m . As shown in Figure 2, both independent and dependent operators using alternating minimization or nuclear-norm minimization require a number of measurements proportional to the dimension of d . We also see that dependent operators require a slight larger number of measurements than that of independent ones. Another interesting observation is that although our theoretical analysis requires a higher measurement complexity of ALS than that of Nuclear methods, the empirical results show their measurement complexities are almost identical for the same measurement operator.

References

1. Agarwal, A., Anandkumar, A., Jain, P., Netrapalli, P., Tandon, R.: Learning sparsely used overcomplete dictionaries via alternating minimization. COLT (2014)
2. Cai, T.T., Zhang, A., et al.: Rop: Matrix recovery via rank-one projections. *The Annals of Statistics* **43**(1), 102–138 (2015)
3. Candès, E.J., Recht, B.: Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* **9**(6), 717–772 (2009)
4. Candès, E.J., Tao, T.: The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inform. Theory* **56**(5), 2053–2080 (2009)
5. Chen, Y.: Incoherence-optimal matrix completion. arXiv preprint [arXiv:1310.0154](https://arxiv.org/abs/1310.0154) (2013)
6. Gross, D.: Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory* **57**(3), 1548–1566 (2011)
7. Hardt, M.: Understanding alternating minimization for matrix completion. In: *Foundations 2014 IEEE 55th Annual Symposium on of Computer Science (FOCS)*, pp. 651–660. IEEE (2014)
8. Hardt, M., Wootters, M.: Fast matrix completion without the condition number. In: *Proceedings of The 27th Conference on Learning Theory*, pp. 638–678 (2014)
9. Hsieh, C.J., Dhillon, I.S., Ravikumar, P.K., Becker, S., Olsen, P.A.: Quic & dirty: A quadratic approximation approach for dirty statistical models. In: *Advances in Neural Information Processing Systems*, pp. 2006–2014 (2014)
10. Hsieh, C.J., Olsen, P.: Nuclear norm minimization via active subspace selection. In: *Proceedings of The 31st International Conference on Machine Learning*, pp. 575–583 (2014)
11. Jain, P., Dhillon, I.S.: Provable inductive matrix completion (2013). CoRR. <http://arxiv.org/abs/1306.0626>
12. Jain, P., Meka, R., Dhillon, I.S.: Guaranteed rank minimization via singular value projection. In: *NIPS*, pp. 937–945 (2010)
13. Jain, P., Netrapalli, P., Sanghavi, S.: Low-rank matrix completion using alternating minimization. In: *STOC* (2013)
14. Keshavan, R.H., Montanari, A., Oh, S.: Matrix completion from a few entries. *IEEE Transactions on Information Theory* **56**(6), 2980–2998 (2010)
15. Kueng, R., Rauhut, H., Terstiege, U.: Low rank matrix recovery from rank one measurements. arXiv preprint [arXiv:1410.6913](https://arxiv.org/abs/1410.6913) (2014)
16. Lee, K., Bresler, Y.: Guaranteed minimum rank approximation from linear observations by nuclear norm minimization with an ellipsoidal constraint. arXiv preprint [arXiv:0903.4742](https://arxiv.org/abs/0903.4742) (2009)

17. Liu, Y.K.: Universal low-rank matrix recovery from pauli measurements. In: Advances in Neural Information Processing Systems, pp. 1638–1646 (2011)
18. Netrapalli, P., Niranjan, U., Sanghavi, S., Anandkumar, A., Jain, P.: Non-convex robust PCA. In: Advances in Neural Information Processing Systems, pp. 1107–1115 (2014)
19. Recht, B.: A simpler approach to matrix completion. *The Journal of Machine Learning Research* **12**, 3413–3430 (2011)
20. Recht, B., Fazel, M., Parrilo, P.A.: Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review* **52**(3), 471–501 (2010)
21. Tropp, J.A.: User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics* **12**(4), 389–434 (2012)
22. Xu, M., Jin, R., Zhou, Z.H.: Speedup matrix completion with side information: application to multi-label learning. In: Advances in Neural Information Processing Systems, pp. 2301–2309 (2013)
23. Yu, H.F., Hsieh, C.J., Si, S., Dhillon, I.S.: Scalable coordinate descent approaches to parallel matrix factorization for recommender systems. In: ICDM, pp. 765–774 (2012)
24. Yu, H.F., Jain, P., Kar, P., Dhillon, I.S.: Large-scale multi-label learning with missing labels. In: Proceedings of The 31st International Conference on Machine Learning, pp. 593–601 (2014)
25. Zuk, O., Wagner, A.: Low-rank matrix recovery from row-and-column affine measurements. In: Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6–11 July 2015, pp. 2012–2020 (2015)