UC Berkeley
UC Berkeley Previously Published Works

Title
A new complexity metric for nonconvex rank-one generalized matrix completion

Permalink
https://escholarship.org/uc/item/9t75j4dh

Journal
Mathematical Programming, 207(1-2)

ISSN
0025-5610

Authors
Zhang, Haixiang
Yalcin, Baturalp
Lavaei, Javad
et al.

Publication Date
2024-09-01

DOI
10.1007/s10107-023-02008-5

Copyright Information
This work is made available under the terms of a Creative Commons Attribution-NonCommercial-ShareAlike License, available at https://creativecommons.org/licenses/by-nc-sa/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Mathematical Programming (2024) 207:227–268
https://doi.org/10.1007/s10107-023-02008-5

FULL LENGTH PAPER

Series A

A new complexity metric for nonconvex rank-one generalized matrix completion

Haixiang Zhang[1] · Baturalp Yalcin[2] · Javad Lavaei[2] · Somayeh Sojoudi[3]

Received: 5 April 2022 / Accepted: 20 July 2023 / Published online: 16 August 2023
© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2023

Abstract

In this work, we develop a new complexity metric for an important class of low-rank matrix optimization problems in both symmetric and asymmetric cases, where the metric aims to quantify the complexity of the nonconvex optimization landscape of each problem and the success of local search methods in solving the problem. The existing literature has focused on two recovery guarantees. The RIP constant is commonly used to characterize the complexity of matrix sensing problems. On the other hand, the incoherence and the sampling rate are used when analyzing matrix completion problems. The proposed complexity metric has the potential to generalize these two notions and also applies to a much larger class of problems. To mathematically study the properties of this metric, we focus on the rank-1 generalized matrix completion problem and illustrate the usefulness of the new complexity metric on three types of instances, namely, instances with the RIP condition, instances obeying the Bernoulli sampling

We note that a similar complexity metric based on a special case of instances in Sect. 3.3 was proposed in our conference paper [56]. However, the complexity metric in this work has a different form and is proved to work on a broader set of applications. In addition, we prove several theoretical properties of the metric in this work, which are not included in [56].

✉ Javad Lavaei
lavaei@berkeley.edu

Haixiang Zhang
haixiang_zhang@berkeley.edu

Baturalp Yalcin
baturalp_yalcin@berkeley.edu

Somayeh Sojoudi
sojoudi@berkeley.edu

1 Department of Mathematics, University of California, Berkeley, CA, USA

2 Department of Industrial Engineering and Operations Research, University of California, Berkeley, CA, USA

3 Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA, USA

model, and a synthetic example. We show that the complexity metric exhibits a consistent behavior in the three cases, even when other existing conditions fail to provide theoretical guarantees. These observations provide a strong implication that the new complexity metric has the potential to generalize various conditions of optimization complexity proposed for different applications. Furthermore, we establish theoretical results to provide sufficient conditions and necessary conditions for the existence of spurious solutions in terms of the proposed complexity metric. This contrasts with the RIP and incoherence conditions that fail to provide any necessary condition.

Keywords Matrix completion · Complexity metric · Nonconvex optimization · Global convergence

Mathematics Subject Classification 05C90 · 65F55 · 90C26

1 Introduction

A variety of modern signal processing and machine learning applications require solving optimization problems that involve a low-rank matrix variable. More specifically, given measurements to some unknown ground truth matrix $M^* \in \mathbb{R}^{n\times n}$ of rank $r \ll n$, the *low-rank matrix optimization* problem can be formulated as

$$\min_{M\in\mathbb{R}^{n\times n}} f(M; M^*) \quad \text{s.t.} \quad M \succeq 0, \quad \text{rank}(M) \le r, \tag{1.1}$$

where $f(\cdot; M^*)$ is the loss function that penalizes the mismatch between the input matrix and M^*. The goal is to recover the matrix M^* via (1.1). Examples of this problem include matrix sensing [43, 59, 62], matrix completion [12, 13, 25], phase retrieval [10, 19, 48] and robust principle component analysis [9, 24]; see the review papers [17, 22] for more applications. The asymmetric version of problem (1.1) eliminates the condition $M \succeq 0$ and allows M to be a non-square matrix. To deal with the nonconvex rank constraint, there have been several works on the convex relaxations of problem (1.1). More concretely, one may replace the rank constraint with a nuclear norm regularizer [9, 12, 13, 35, 43]. The convex relaxation approach is proven to achieve the optimal sampling complexity for various statistical models. In the special case when $f(\cdot; M^*)$ is a linear function, the sketching method [58] can be applied to accelerate the computation. However, for most applications of problem (1.1), the convex relaxation approach needs to update a matrix variable in each iteration, which relies on the Singular Value Decomposition (SVD) of the matrix variable. This will lead to an $O(n^3)$ computational complexity in each iteration and an $O(n^2)$ space complexity, which are prohibitively high for large-scale problems; see the numerical comparison in [63].

To improve the computational efficiency, an alternative approach was proposed by Burer and Monteiro [8], which is named as the Burer-Monteiro factorization approach. The factorization approach is based on the fact that the mapping $U \mapsto UU^T$ is surjective onto the manifold of positive semi-definite matrices of rank at most r, where $U \in \mathbb{R}^{n\times r}$. Therefore, problem (1.1) is equivalent to

$$\min_{U \in \mathbb{R}^{n \times r}} f\left(UU^T; M^*\right), \tag{1.2}$$

which is an unconstrained nonconvex problem. A major difficulty about nonconvex optimization problems is the existence of spurious local minima.[1] In general, common local search methods are only able to guarantee a point approximately satisfying the first-order and the second-order necessary optimality conditions. Therefore, local search methods with a random initialization will likely be stuck at spurious local minima and unable to converge to the global solution. However, despite the aforementioned issue of nonconvex optimization problems, simple iterative algorithms such as gradient descent and alternating minimization have achieved empirical success in a wide range of applications. In recent years, substantial progress has been made on the theoretical understandings of these algorithms, which generally focused on proving the absence of spurious local minima. For example, the alternating minimization algorithm was first studied in [31, 41, 42]. The (stochastic) gradient descent algorithm, which is generally easier to implement than the alternating minimization algorithm, was analyzed in [10, 17, 19, 53, 57]. Besides algorithmic analysis, a critical geometric property named the strict-saddle property [48] was established in [25, 48, 59, 64], which can guarantee the polynomial-time global convergence of various saddle-escaping algorithms [5, 14, 33].

Complexity metrics are useful to characterize the behavior of local search methods for problem (1.2). A small complexity metric implies that the landscape of problem (1.2) is benign and thus, local search methods with random initialization converge to global solutions with high probability. Otherwise, if the complexity metric takes a large value, problem (1.2) may have spurious local minima, which will imply the failure of most local search methods. However, the existing so-called "complexity metrics" for problem (1.2) are only able to guarantee a benign landscape when the complexity is small and fail to prove the existence of spurious local minima when the complexity is large. To differentiate with true complexity metrics, we use the term *recovery guarantees* to reflect such weaker properties. In addition, the existing recovery guarantees were designed separately for different applications. As a result, several different bounds were proposed to characterize the optimization complexity of problem (1.2). For example, in the context of matrix sensing problems, the following Restrict Isometry Property (RIP) is usually assumed:

Definition 1.1 ([43, 64]) Given natural numbers r and s, the function $f(\cdot; M^*)$ is said to satisfy the **Restricted Isometry Property** (RIP) of rank $(2r, 2s)$ for a constant $\delta \in [0, 1)$, denoted as δ-RIP$_{2r,2s}$, if

$$(1-\delta)\|K\|_F^2 \leq \left[\nabla^2 f(M; M^*)\right](K, K) \leq (1+\delta)\|K\|_F^2 \tag{1.3}$$

holds for all matrices $M, K \in \mathbb{R}^{n \times n}$ such that $\operatorname{rank}(M) \leq 2r$, $\operatorname{rank}(K) \leq 2s$, where $\left[\nabla^2 f(M; M^*)\right](\cdot, \cdot)$ is the curvature of the Hessian at point M.

[1] A point U^0 is called a spurious local minimum if it is a local minimum of problem (1.2) and $U^0\left(U^0\right)^T \neq M^*$.

One important class of matrix sensing problems is the *linear matrix sensing problem*, which is induced by linear measurements of the ground truth matrix M^*. If the ℓ_2-loss is used, the linear matrix sensing problem can be formulated as

$$\min_{U \in \mathbb{R}^{n\times r}} \frac{1}{m} \sum_{i=1}^{m} \left\langle A_i, UU^T - M^* \right\rangle^2, \tag{1.4}$$

where $m \in \mathbb{N}$ is the number of measurements modeled by the known measurement matrices $A_i \in \mathbb{R}^{n\times n}$ for all $i \in [m]$. In the special case when each matrix A_i is an independently identically distributed Gaussian random matrix, the δ-RIP$_{2r,2s}$ condition holds with high probability if $m = O\left(nr\delta^{-2}\right)$ [11]. The RIP constant δ plays a critical role in bounding the optimization complexity of problem (1.2). In [7], the authors showed that the strict-saddle property holds for problem (1.2) if the δ-RIP$_{2r,2r}$ condition holds with $\delta < 1/2$ and the ground truth matrix satisfies $\text{rank}(M^*) = r$. On the other hand, counterexamples have been constructed in [59, 62] to illustrate that the strict-saddle property can fail under the δ-RIP$_{2r,2r}$ condition with $\delta \geq 1/2$.

Despite these strong theoretical results under the RIP assumption, there exists a large number of applications that do not satisfy the RIP condition. One of those applications without the RIP condition is the *matrix completion problem*. Given a set of indices $\Omega \subset [n] \times [n]$, the matrix completion problem aims at recovering the low-rank matrix M^* from the available entries M^*_{ij} for $(i, j) \in \Omega$. With the least squares loss function, the matrix completion problem can be formulated as

$$\min_{U \in \mathbb{R}^{n\times r}} \sum\nolimits_{(i,j)\in\Omega} \left[\left(UU^T\right)_{ij} - M^*_{ij} \right]^2. \tag{1.5}$$

The matrix completion problem (1.5) is a special case of the matrix sensing problem (1.4), where each measurement matrix A_i has exactly one nonzero entry. However, the RIP$_{2r,2r}$ condition does not hold for problem (1.5) unless all entries of M^* are observed, namely, when $\Omega = [n] \times [n]$. As an alternative to the RIP condition, the optimization complexity of problem (1.5) is closely related to the incoherence of M^*.

Definition 1.2 ([12]) Given a constant $\mu \in [1, n]$, the ground truth matrix M^* is said to be μ-**incoherent** if

$$\|e_i^T V^*\|_F \leq \sqrt{\mu r/n}, \quad \forall i \in [n], \tag{1.6}$$

where $V^*\Lambda^*(V^*)^T$ is the truncated SVD of M^* and e_i is the i-th standard basis of $\mathbb{R}^n$.

Intuitively, if the ground truth M^* is highly sparse, it is likely that only zero entries of M^* are observed and there is no chance to learn the other entries of the matrix M^*. A relatively small incoherence of M^* avoids this extreme case. The most popular statistical model of the measurements for problem (1.5) is the Bernoulli model, where each entry of M^* is observed independently with probability $p \in (0, 1]$.

Assuming the Bernoulli model, the incoherence of M^* and the sampling probability p can jointly characterize the complexity of the matrix completion problem. For example, the scaled gradient descent algorithm with a spectral initialization [50] converges linearly given the condition $p \geq O\left(\mu r^2 \kappa^2 \max\left(\mu\kappa^2, \log n\right)/n\right)$, where $\kappa := \sigma_1(M^*)/\sigma_r(M^*)$ is the condition number of M^*. In addition, under the assumption that $p \geq O\left(\mu^4 r^6 \kappa^6 \log n/n\right)$, the global convergence was established in [25] through the strict-saddle property of a regularized version of problem (1.5). We note that the dependence on the condition number κ may be unnecessary as shown in [28] and that the condition number is equal to 1 in the rank-1 case. On the other hand, the information-theoretical lower bound in [12] shows that $p \geq \Theta(\mu r \log(n/\delta)/n)$ is necessary for the exact completion with probability at least $1 - \delta$. Therefore, the complexity of problem (1.5) is closely related to the incoherence of M^* and the sampling probability p. In the remainder of this work, we refer to the conditions on the incoherence of M^* and sampling rate p as *incoherence conditions* when there is no confusion in the context.

To be more rigorous, the RIP condition and the incoherence condition may have a subtle difference in their nature. As a counterpart of the incoherence condition in other low-rank matrix optimization problems, one should consider conditions in terms of the sampling complexity. On the other hand, the RIP condition is a deterministic condition on the loss function and is not related to the underlying random model. However, there is a wide range of problems that satisfy the RIP condition when the sample complexity is sufficiently large. By considering the properties of the RIP condition, we are able to analyze a large number of low-rank matrix optimization problems simultaneously. Therefore, we use the RIP condition instead of conditions based on the sample complexity as a notion of the computational complexity for those problems.

The main issue with the notions of RIP and incoherence is that they require stringent conditions to guarantee the success of local search methods for recovering M^*. Whenever these conditions are violated, local search methods may still work successfully, which questions whether these customized notions designed for special cases of the problem truly capture the complexity of the problem in general. Hence, it is natural to ask:

> *Does there exist a complexity metric with two properties: (i) it is consistent with existing recovery guarantees designed for different applications, e.g., the RIP constant δ and the incoherence μ combined with the sampling rate p*, (ii) *even when the customized conditions for different applications are violated, it still quantifies the optimization complexity of the problem in the sense that the smaller the value of this metric is, the higher the success of local search methods with random initialization is in finding the ground truth M^**?

In this work, we provide a partial answer to the question by developing a powerful complexity metric. To analyze the usefulness of this new metric, we focus on the rank-1 generalized matrix completion problem

$$\min_{u \in \mathbb{R}^n} \sum\nolimits_{i,j \in [n]} C_{ij} \left(u_i u_j - M^*_{ij}\right)^2, \tag{1.7}$$

where the ground truth M^* is symmetric and has rank at most 1. The weights are $C_{ij} \geq 0$ for all $i, j \in [n]$. Without loss of generality, we can assume that the matrix $C := (C_{ij})_{i,j\in[n]}$ is symmetric since otherwise one can replace C with $\left(C + C^T\right)/2$, which will not change the optimization landscape. We use $\mathcal{MC}(C, u^*)$ to denote the instance of problem (1.7) with the weight matrix C and the ground truth $M^* = u^*(u^*)^T$, for all $C \in \mathbb{R}^{n\times n}$ and $u^* \in \mathbb{R}^n$. The matrix completion problem (1.5) is a special case of the generalized matrix completion problem (1.7), where $C_{ij} = 1$ if $(i, j) \in \Omega$ and $C_{ij} = 0$ otherwise.

Moreover, problem (1.7) is a special case of the matrix sensing problem (1.4), where each measurement only captures one entry of M^*. However, the problem (1.7) still contains difficult instances of the matrix sensing problem from the perspective of the RIP condition. In Sect. 3.3, we show that there exists an instance of problem (1.7) that satisfies the $1/2$-RIP$_{2,2}$ condition but has spurious local minima. This counterexample implies that the optimal RIP bound in [59, 62] still holds for problem (1.7) and thus, problem (1.7) contains difficult instances of the matrix sensing problem. Moreover, we show in Sect. 3.1 that some of the results to be developed for problem (1.7) can be extended to the general matrix sensing problem (1.4).

Now, we provide an intuition into the design of our complexity metric for problem (1.7). For a given problem instance of (1.7), if there exist global solutions u^1, u^2 such that $u^1\left(u^1\right)^T \neq u^2\left(u^2\right)^T$, it is impossible to decide which global solution corresponds to M^* from the observations. Intuitively, no matter what optimization algorithm we choose and how much computational effort is exerted, there is a chance that we could not recover M^* by solving problem (1.7). This observation motivates us to define the complexity metric to be the inverse of the infimum of the distance between any given instance and the set of instances with multiple global solutions. Since problem (1.7) is parameterized by the weight matrix C and the global solution M^*, we are able to define the metric through norms in Euclidean spaces and their Cartesian products. In addition, in the rank-1 case, (random) graph theory serves as an important tool in characterizing the solvability of problem (1.7). These two advantages enable a more thorough analysis of the new complexity metric. The formal definition of the metric is provided in Sect. 2. In this work, we exhibit several pieces of evidence to show that the proposed metric can serve as an alternative to the RIP constant and the incoherence, which are summarized below:

1. For problem instances that satisfy the δ-RIP$_{2,2}$ condition, we provide an upper bound on the complexity metric. The upper bound is tightened with extra information about the incoherence of M^*. Similarly, for matrix completion problems obeying the Bernoulli sampling model, an upper bound on the complexity metric in terms of the incoherence of M^* is derived.
2. We then construct a class of parameterized instances of problem (1.7), where the RIP condition fails to provide useful guarantees. A lower bound on the complexity metric is developed to prove that instances whose complexity metric is larger than the lower bound have an exponential number of spurious local minima. In addition, an upper bound that is consistent with the aforementioned two upper bounds is established to guarantee the absence of spurious local minima if the complexity metric is below this bound. The consistency of the upper bounds between different

types of models provides strong evidence that the new complexity metric is able to provide theoretical guarantees for different applications, even when the RIP condition or the incoherence condition fails.

3. We prove the existence of a non-trivial upper bound on the complexity metric. For all problem instances whose complexity metric is below this upper bound, problem (1.7) has no spurious local minima and M^* can be successfully found via local search methods with random initialization. In addition, under a standard bounded-away-from-zero assumption, we show that all instances with a larger complexity metric will possess spurious local minima.
4. We extend all results for the symmetric generalized matrix completion problem to the asymmetric case, where the low-rank matrix is decomposed into UV^T for some $U \in \mathbb{R}^{m\times r}$ and $V \in \mathbb{R}^{n\times r}$ in problem (1.2).

Based on the aforementioned results, we make some key conjectures and discuss the potential extensions of the proposed metric to more general cases of the low-rank matrix optimization problem (1.1).

1.1 Related works

Following the famous *Netflix prize*, the theoretical analysis of problem (1.1) has attracted a lot of attention in recent years; see the review papers [18, 22]. Early attempts mainly focused on the construction of convex relaxations to rank-constrained problems [9, 12, 13, 43], where the RIP condition and the incoherence condition were introduced. Recently, several modified RIP conditions were proposed to better characterize the landscapes of other classes of problems, e.g., the ℓ_1/ℓ_2-RIP condition [36], the sign-RIP condition [39], and the approximation and sharpness condition [15].

Although the convex relaxation is usually guaranteed to recover the exact ground truth with almost the optimal sample complexity, the associated algorithms operate in the space of matrix variables and, thus, are computationally inefficient for large-scale problems [63]. Similar issues are observed for algorithms based on the Singular Value Projection [30] and Riemannian optimization algorithms [2, 29, 38, 54, 55]. The analysis of the convex relaxation approach in the noisy case is recently conducted by bridging the convex and the nonconvex approaches [20, 21].

To deal with the difficulties in solving large-scale problems, an efficient alternative model (1.2) using the Burer–Monteiro factorization is considered. Despite the non-convexity, a growing number of works demonstrated that problem (1.2) has benign landscapes and, therefore, is amenable to efficient optimization. Theoretical analysis stems from the alternating minimization method [1, 27, 28, 31, 41, 42]. The alternating minimization method has the advantage that the number of iterations has only logarithmic dependence on the condition number of the ground truth [28]. More recently, this advantage is also achieved by the scaled (sub)gradient descent algorithm [50–52, 61].

The gradient descent algorithm has also gained significant attention due to its simplicity in implementation. In general, there are two ways to apply the gradient descent algorithm. First, the gradient descent algorithm can serve as the local refinement method after a suitable initialization [4, 10, 17, 49, 53, 57]. On the other hand, the gradient descent algorithm is proved to converge globally for the phase retrieval problem [19]. More generally, under the strict-saddle property, a number of saddle-escaping algorithms [5, 14, 33] converge to the global solution in polynomial time; see e.g., [6, 7, 16, 25, 26, 40, 47, 48, 59, 62, 64]. Moreover, the gradient descent algorithm is proved to have the implicit regularization phenomenon in the over-parameterization case [23, 37, 46].

1.2 Notation

The number of elements in a finite set $\mathcal{S}$ is denoted as $|\mathcal{S}|$. We use $\overline{\mathcal{S}}$ to denote the closure of a set $\mathcal{S} \subset \mathbb{R}^n$. The index set $\{1, \ldots, n\}$ is denoted as $[n]$ for all $n \in \mathbb{N}$. The entry-wise ℓ_1-norm and the Frobenius norm of a matrix M are denoted as $\|M\|_1$ and $\|M\|_F$, respectively. The unit sphere of matrices with non-negative entries, denoted as $\mathbb{S}^{n^2-1}_{+,1}$, is the set of all symmetric matrices $X \in \mathbb{R}^{n\times n}$ such that $\|X\|_1 = 1$ and $X_{ij} \geq 0$ for all $i, j \in [n]$. Similarly, the unit sphere of vectors, S^{n-1}_1, is the set of all vectors $x \in \mathbb{R}^n$ such that $\|x\|_1 = 1$. For every symmetric matrix $M \in \mathbb{R}^{n\times n}$, the minimum eigenvalue is denoted as $\lambda_{min}(M)$. The n-by-n identity matrix is denoted as $\mathcal{I}_n$. The notation $M \succeq 0$ means that the matrix M is symmetric and positive semi-definite. The sub-matrix $R_{i:j,k:\ell}$ consists of the i-th to the j-th rows and the k-th to the ℓ-th columns of matrix R. For every vector $x \in \mathbb{R}^n$, the sets of indices corresponding to zero and nonzero components of x are denoted as $\mathcal{I}_0(x)$ and $\mathcal{I}_1(x)$, respectively. For every instance $\mathcal{MC}(C, u^*)$, we use $\mathbb{G}(C, u^*) = \left[\mathbb{V}(C, u^*), \mathbb{E}(C, u^*), \mathbb{W}(C, u^*)\right]$ to denote the associated weighted graph, which is defined in Sect. 2. The unweighted undirected graph $\mathbb{G}$ with node set $\mathbb{V}$ and edge set $\mathbb{E}$ is denoted as $\mathbb{G} = (\mathbb{V}, \mathbb{E})$. The objective function of an instance $\mathcal{MC}(C, u^*)$ is shown as $g(u; C, u^*) := \sum_{i,j\in[n]} C_{ij}\left(u_i u_j - u^*_i u^*_j\right)^2$. We use $\left[\nabla^2 g(M; C, u^*)\right](K, L) := \sum_{i,j,k,\ell}\left[\nabla^2 g(M; C, u^*)\right]_{i,j,k,\ell} K_{i,j} L_{k,\ell}$ to denote the action of the Hessian $\nabla^2 g(M; C, u^*)$ on any two matrices K and L. The notations $a_n = O(b_n)$ and $a_n = \Theta(b_n)$ mean that there exist constants $c_1, c_2 > 0$ such that $a_n \leq c_2 b_n$ and $c_1 b_n \leq a_n \leq c_2 b_n$ hold for all $n \in \mathbb{Z}$, respectively.

1.3 Organization

In the remainder of this paper, we first define the proposed complexity metric and derive basic properties of the metric in Sect. 2. In Sect. 3, we analyze this metric under existing conditions, including the RIP condition and the incoherence condition. Section 4 is devoted to the theoretical guarantees provided by the new complexity metric on the general instances of problem (1.7). The results for the rank-1 asymmetric generalized matrix completion problem are essentially the same as those for the symmetric case, and therefore they are provided only in the online version [60]. Finally, we conclude the paper in Sect. 5.

2 New complexity metric and basic properties

In this section, we first provide the formal definition of the new complexity metric and investigate the properties of the proposed metric. More specifically, we show that we are able to utilize the graph theory to estimate the complexity metric and calculate the minimum possible value of the proposed complexity metric in closed form. Before proceeding to the definitions, we note that the problem (1.7) is “scale-free” in the sense that the instance $\mathcal{MC}(\eta_1 C, \eta_2 u^*)$ has the same landscape as $\mathcal{MC}(C, u^*)$ up to a scaling, where $C \in \mathbb{R}^{n\times n}$, $u^* \in \mathbb{R}^n$ and $\eta_1, \eta_2 > 0$ are constants. Therefore, we may normalize the parameters C and u^* without loss of generality, as follows:

Assumption 2.1 Assume that $C \in \mathbb{S}^{n^2-1}_{+,1}$ and $u^* \in \mathbb{S}^{n-1}_1$, i.e., $\|C\|_1 = \|u^*\|_1 = 1$.

The above assumption excludes the degenerate cases when $C = 0$ or $M^* = 0$. If $C = 0$, the objective function is always 0 and it is impossible to recover the ground truth. For the case when $M^* = 0$, we can prove that either $u = 0$ is the only stationary point or the instance $\mathcal{MC}(C, 0)$ has multiple different global solutions. In the first situation, the results in [34] imply that the randomly initialized gradient descent algorithm will converge to 0 with probability 1. In the second situation, the instance is information-theoretically unsolvable. We provide a more detailed analysis in the appendix and assume that Assumption 2.1 holds in the remainder of the paper.

The definition of the complexity metric is closely related to the set of instances with multiple “essentially different” global solutions. More specifically, the set of degenerate instances is defined as

$$\mathcal{D} := \Big\{ \left(C, u^*\right) \mid C \in \mathbb{S}^{n^2-1}_{+,1}, u^* \in \mathbb{S}^{n-1}_1, \\ \exists u \in \mathbb{R}^n \quad \text{s.t.} \quad g\left(u; C, u^*\right) = 0,\ uu^T \neq u^* \left(u^*\right)^T \Big\}.$$

Since there exist multiple global solutions to problem (1.7) if $(C, u^*) \in \mathcal{D}$, it is information-theoretically impossible to find the ground truth for any instance in $\mathcal{D}$. Intuitively, we say that the *optimization complexity* of all instances in $\mathcal{D}$ is infinity. Motivated by the above observation, we introduce the new complexity metric.

Definition 2.1 (Complexity Metric) Given arbitrary parameters $C \in \mathbb{S}^{n^2-1}_{+,1}$, $u^* \in \mathbb{S}^{n-1}_1$ and $\alpha \in [0, 1]$, the complexity of the instance $\mathcal{MC}(C, u^*)$ is defined as

$$\mathbb{D}_\alpha\left(C, u^*\right) := \left[\inf_{(\tilde{C},\tilde{u}^*)\in\mathcal{D}} \alpha\|C - \tilde{C}\|_1 + (1-\alpha)\|u^* - \tilde{u}^*\|_1 \right]^{-1}. \tag{2.1}$$

Since the set $\mathcal{D}$ is bounded, the infimum in the definition is finite. The term inside the inverse operation can be viewed as a weighted distance between the point (C, u^*) and the set $\mathcal{D}$. In addition, we take the convention that $1/0 = +\infty$ and thus, $\mathbb{D}_\alpha(C, u^*) = +\infty$ for all $(C, u^*) \in \mathcal{D}$. In this work, we choose the entry-wise ℓ_1-norm in (2.1) for the simplicity of calculations. We believe that a similar theory can still be derived for other choices of the norm. We note that a similar complexity was proposed in [44, 45] for conic optimization and to the best of authors' knowledge, there is no similar complexity metric for nonconvex optimization problems.

For the parameter α, we will discuss two potential choices in this section, namely α^* and $\alpha^\diamond$. In the case when $\alpha = \alpha^*$, the range of the complexity metric has the largest size. Intuitively, by choosing $\alpha = \alpha^*$, the difference between the complexities of two instances will be maximized and thus, it is easier to compare the complexities of different instances. On the other hand, when we choose $\alpha = \alpha^\diamond$, the complexity metric attains its minimum possible value if and only if the 0-RIP$_{2,2}$ condition holds. This is consistent with the intuition that instances with the RIP constant 0 are the easiest to solve. We note that both α^* and $\alpha^\diamond$ satisfy $1 - \alpha = \Theta(1/n)$. Moreover, in Sect. 3, we show that the parameter α strikes a balance between the RIP constant of the instance and the incoherence of the ground truth. It is still an open question what the optimal choice of parameter α is, which may depend on the class of problems under consideration. It may be needed to jointly consider the complexity metric with several different choices of α to determine the solvability of the instance.

2.1 Basic properties of the new complexity metric

We first provide a more concrete characterization of the set $\mathcal{D}$. In the rank-1 case, we are able to exactly describe the set $\mathcal{D}$ using graph-theoretic notations. We introduce the associated graphs of any instance of the problem. Given an instance $\mathcal{MC}(C, u^*)$, the weighted graph $\mathbb{G}(C, u^*) = \left[\mathbb{V}(C, u^*), \mathbb{E}(C, u^*), \mathbb{W}(C, u^*)\right]$ is defined by

$$\mathbb{V}\left(C, u^*\right) := [n], \quad \mathbb{E}\left(C, u^*\right) := \left\{\{i, j\} \mid C_{ij} > 0, i, j \in [n]\right\},$$
$$\left[\mathbb{W}\left(C, u^*\right)\right]_{ij} := C_{ij}, \quad \forall i, j \in [n] \quad \text{s.t.} \quad \{i, j\} \in \mathbb{E}\left(C, u^*\right).$$

To include the information of u^*, we define

$$\mathcal{I}_1\left(C, u^*\right) := \left\{i \in [n] \mid u_i^* \neq 0\right\}, \quad \mathcal{I}_0\left(C, u^*\right) := [n] \backslash \mathcal{I}_1\left(C, u^*\right),$$
$$\mathcal{I}_{00}\left(C, u^*\right) := \left\{i \in \mathcal{I}_0\left(C, u^*\right) \mid \{i, j\} \notin \mathbb{E}\left(C, u^*\right), \ \forall j \in \mathcal{I}_1\left(C, u^*\right)\right\}.$$

Intuitively, the sets $\mathcal{I}_1(C, u^*)$ and $\mathcal{I}_0(C, u^*)$ contain the locations of the nonzero and zero components of u^*. The subset $\mathcal{I}_{00}(C, u^*)$ corresponds to indices in $\mathcal{I}_0(C, u^*)$ that are not connected to any index in $\mathcal{I}_1(C, u^*)$. We denote the subgraph of $\mathbb{G}(C, u^*)$ induced by the index set $\mathcal{I}_1(C, u^*)$ as $\mathbb{G}_1(C, u^*) = [\mathcal{I}_1(C, u^*), \mathbb{E}_1(C, u^*), \mathbb{W}_1(C, u^*)]$, where $\mathbb{E}_1(C, u^*)$ and $\mathbb{W}_1(C, u^*)$ are the edge set and weight set of this subgraph. The following theorem provides an equivalent definition of $\mathcal{D}$ in terms of $\mathcal{I}_{00}(C, u^*)$ and $\mathbb{G}_1(C, u^*)$.

Theorem 2.2 *Given $C \in \mathbb{S}_{+,1}^{n^2-1}$ and $u^* \in \mathbb{S}_1^{n-1}$, it holds that $(C, u^*) \notin \mathcal{D}$ if and only if*

1. *$\mathbb{G}_1(C, u^*)$ is connected and not bipartite;*
2. *$\{i, i\} \in \mathbb{E}(C, u^*)$ for all $i \in \mathcal{I}_{00}(C, u^*)$.*

Proof We first construct counterexamples for the necessity part and then prove the uniqueness of the global minimum (up to a sign flip) for the sufficiency part in the online version [60]. For the notational simplicity, we fix the point (C, u^*) and omit them in the notations.

Necessity. In this part, our goal is to construct a solution $u \in \mathbb{R}^n$ such that

$$u_i u_j = u_i^* u_j^*, \quad \forall \{i, j\} \in \mathbb{E}; \quad uu^T \neq u^*(u^*)^T.$$

We denote $M^* := u^*(u^*)^T$ and analyze three different cases below.

Case I. First, we consider the case when $\mathbb{G}_1$ is disconnected, which means that there exist two non-empty subsets $\mathcal{I}$ and $\mathcal{J}$ such that

$$\mathcal{I} \cup \mathcal{J} = \mathcal{I}_1, \quad \mathcal{I} \cap \mathcal{J} = \emptyset; \quad \{i, j\} \notin \mathbb{E}_1, \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{J}.$$

We define the vector $u \in \mathbb{R}^n$ as

$$u_i := 0, \quad \forall i \in \mathcal{I}_0; \quad u_i = u_i^*, \quad \forall i \in \mathcal{I}; \quad u_i = -u_i^*, \quad \forall i \in \mathcal{J}.$$

The above definition leads to

$$u_i u_j = \begin{cases} -M_{ij}^* & \text{if } i \in \mathcal{I} \text{ and } j \in \mathcal{J} \\ M_{ij}^* & \text{otherwise.} \end{cases}$$

Since $u_i^* \neq 0$ for all $i \in \mathcal{I}_1$, it follows that $u_i u_j = -M_{ij}^* \neq M_{ij}^*$ for all $\{i, j\}$ such that $i \in \mathcal{I}$ and $j \in \mathcal{J}$.

Case II. Next, we consider the case when $\mathbb{G}_1$ is bipartite, which means that there exist two non-empty subsets $\mathcal{I}$ and $\mathcal{J}$ such that

$$\mathcal{I} \cup \mathcal{J} = \mathcal{I}_1, \quad \mathcal{I} \cap \mathcal{J} = \emptyset; \quad \{i, j\} \notin \mathbb{E}_1, \quad \forall i, j \in \mathcal{I}_1 \quad \text{s.t.} \quad i, j \in \mathcal{I} \text{ or } i, j \in \mathcal{J}.$$

In this case, we define the vector $u \in \mathbb{R}^n$ as

$$u_i := 0, \quad \forall i \in \mathcal{I}_0; \quad u_i := u_i^*/2, \quad \forall i \in \mathcal{I}; \quad u_i := 2u_i^*, \quad \forall i \in \mathcal{J}.$$

Now, we have

$$u_i u_j = \begin{cases} M_{ij}^*/4 & \text{if } i, j \in \mathcal{I} \\ 4M_{ij}^* & \text{if } i, j \in \mathcal{J} \\ M_{ij}^* & \text{otherwise.} \end{cases}$$

Since $M_{ij}^* \neq 0$ for all $i, j \in \mathcal{J}$, we have that $u_i u_j = 4M_{ij}^* \neq M_{ij}^*$ for all $i, j \in \mathcal{J}$.

Case III. Finally, we check the case when there exists a node $i_0 \in \mathcal{I}_{00}$ such that $\{i_0, i_0\} \notin \mathbb{E}$. In this case, we define the vector $u \in \mathbb{R}^n$ as

$$u_{i_0} := 1, \quad u_i := u_i^*, \quad \forall i \in [n] \backslash \{i_0\}.$$

Now, we have

$$u_{i_0}u_{i_0} = 1 \neq 0 = M^*_{i_0 i_0}, \quad u_i u_j = M^*_{ij}, \quad \forall \{i, j\} \in \mathbb{E}.$$

Combining the above three cases completes the proof of the necessity part. □

Since the set $\mathcal{D}$ is bounded, the infimum in the definition (2.1) can be attained by using the closure of $\mathcal{D}$, namely

$$\mathbb{D}_\alpha\left(C, u^*\right) = \left[\min_{(\tilde{C}, \tilde{u}^*) \in \overline{\mathcal{D}}} \alpha \|C - \tilde{C}\|_1 + (1 - \alpha)\|u^* - \tilde{u}^*\|_1\right]^{-1}. \tag{2.2}$$

The alternative Definition (2.2) simplifies the verification of parameters that attain the infimum. In addition, with the help of Theorem 2.2, we can exactly characterize the closure $\overline{\mathcal{D}}$, which has a slightly simpler form than $\mathcal{D}$.

Theorem 2.3 *We have the following relation:*

$$\overline{\mathcal{D}} = \left\{\left(C, u^*\right) \mid C \in \mathbb{S}^{n^2-1}_{+,1}, u^* \in \mathbb{S}^{n-1}_1, \mathbb{G}_1\left(C, u^*\right) \textit{ is disconnected or bipartite}\right\}$$
$$\cup \left\{\left(C, u^*\right) \mid C \in \mathbb{S}^{n^2-1}_{+,1}, u^* \in \mathbb{S}^{n-1}_1, \mathcal{I}_{00}\left(C, u^*\right) \textit{ is not empty}\right\}.$$

Let the set in the right-hand side of the above equation be called $\mathcal{D}'$. The proof of Theorem 2.3 is based on a standard technique that first shows $\bar{\mathcal{D}} \subset \mathcal{D}'$ and then shows $\mathcal{D}' \subset \bar{\mathcal{D}}$. The details can be found in the online version [60]. Using the results in Theorems 2.2 and 2.3, we provide an estimate on the scale of the new metric. Since $\mathcal{D}$ is a bounded set, there exists an upper bound on the minimum possible value of the complexity metric, which is defined below:

$$\mathbb{D}^{min}_\alpha := \min_{C \in \mathbb{S}^{n^2-1}_{+,1}, u^* \in \mathbb{S}^{n-1}_1} \mathbb{D}_\alpha\left(C, u^*\right).$$

The next theorem provides the expression of $\mathbb{D}^{min}_\alpha$.

Theorem 2.4 *Suppose that $n \geq 5$. Then, it holds that*

$$\mathbb{D}^{min}_\alpha = \begin{cases} \frac{n}{4\alpha} & \textit{if } \alpha \leq \frac{n^2-5n+4}{n^2-3n-2} \\ \frac{n^2}{2(1-\alpha)(n-2)n+4\alpha} & \textit{if } \frac{n}{n+2} \leq \alpha \leq \frac{n}{n+1} \\ \frac{n(n+1)}{2(1-\alpha)(n-2)(n+1)+4} & \textit{if } \alpha \geq \frac{n}{n+1}. \end{cases}$$

In the regime $\frac{n^2-5n+4}{n^2-3n-2} \leq \alpha \leq n/(n+2)$, we have the estimate

$$\mathbb{D}^{min}_\alpha \in \left[\frac{n}{4\alpha}, \frac{n^2}{4\alpha(n-1)}\right].$$

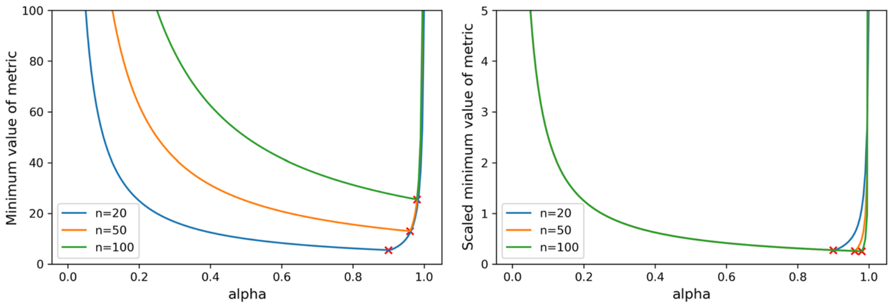

Fig. 1 Comparison of $\mathbb{D}_\alpha^{min}$ for $n = 20, 50, 100$. The red "×" sign refers to the value at α^*. In the right plot, the complexity metric is scaled by n^{-1}

The proof of Theorem 2.4 can be found in Appendix A.1. The results of Theorem 2.4 imply that in the regime where $\alpha \geq \Theta(1)$ and $1-\alpha \geq \Theta\left(n^{-1}\right)$, we have $\mathbb{D}_\alpha^{min} = O(n)$. This suggests that $n^{-1}\mathbb{D}_\alpha(C, u^*)$ may be a dimension-free complexity metric; see more examples supporting this claim in Sect. 3. In addition, the minimum possible value of the complexity is attained at

$$\alpha^* := \left(n^2 - 5n + 4\right) / \left(n^2 - 3n - 2\right).$$

Hence, the set of possible values of the complexity metric attains the maximum size by choosing $\alpha = \alpha^*$. This observation hints that α^* may be the optimal choice of α since it may enable the metric to differentiate instances with different complexities to the maximum degree. Using the exact formulation of $g(\alpha, c)$ in Lemma A.1, we plot the minimum possible value of the complexity metric both without scaling and after scaling by n^{-1} in Fig. 1.

From the numerical results, we can see that the complexity scales with n if α is smaller than α^*, which is consistent with Theorem 2.4. If α is larger than α^*, the complexity metric for different values of n approximately lies on the same curve.

In the following theorem, we show that if $\alpha = \alpha^*$, the instances that attain the minimum value of the complexity metric are unique up to sign flips to components of the global solution.

Theorem 2.5 *Suppose that $n \geq 5$ and the instance $\mathcal{MC}(C, u^*)$ satisfies*

$$\mathbb{D}_{\alpha^*}\left(C, u^*\right) = n/(4\alpha^*).$$

Then, it holds that

$$|u_i^*| = 1/n, \quad C_{ii} = 0, \quad \forall i \in [n]; \quad C_{ij} = 1/[n(n-1)], \quad \forall i, j \in [n], \ i \neq j.$$

The proof of Theorem 2.5 can be found in the online version [60]. The above theorem states that if we choose the weight $\alpha = \alpha^*$, the "easiest" instance is unique up to a change in the signs of the components of the global solution u^*. In the next theorem, we show that a similar property as α^* holds if we set α to be

$$\alpha^\diamond := n/(n+2).$$

Theorem 2.6 *Suppose that $n \geq 5$ and the instance $\mathcal{MC}\,(C, u^*)$ satisfies*

$$\mathbb{D}_{\alpha^\diamond}\left(C, u^*\right) = \mathbb{D}_\alpha^{min} = n(n+2)/[4(n-1)].$$

Then, it holds that

$$|u_i^*| = 1/n, \quad \forall i \in [n]; \quad C = n^{-2} I_n.$$

Since the proof is similar to that of Theorem 2.5, we omit it for brevity. The above theorem implies that the weight matrix C of the "easiest" instances is a constant multiple of the identity matrix I_n, which satisfies the δ-RIP$_{2,2}$ condition with $\delta = 0$. This is consistent with the common sense that the RIP constant δ being 0 is the optimal situation. Hence, Theorem 2.6 suggests that the choice $\alpha^\diamond = n/(n+2)$ may potentially be the optimal choice of α. Moreover, we will prove in Sect. 4.1 that the "easiest" instances in Theorems 2.5 and 2.6 all have a benign landscape in the sense that they satisfy the strict-saddle property [48], which guarantees the polynomial-time global convergence of various algorithms. If the weight α is different from α^* and $\alpha^\diamond$, there may exist multiple "essentially" different instances attaining the minimum complexity.

3 Connections to existing results

In this section, we provide estimates of the proposed complexity metric on two well-studied problem instances and a synthetic problem. More specifically, we consider matrix sensing problems satisfying the RIP condition and matrix completion problems under the Bernoulli sampling model. In addition, we construct a class of instances parameterized by a single parameter. We estimate the threshold of the parameter that separates instances with a desirable optimization landscape from those with a bad landscape. The results in the synthetic example show that our proposed complexity metric has the potential to provide guarantees on the optimization landscape when the RIP condition fails.

3.1 Matrix sensing problem: RIP condition

We first consider instances of problem (1.7) that satisfy the δ-RIP$_{2,2}$ condition, where $\delta \in [0, 1)$ is the RIP constant. However, the constraint that $C \in \mathbb{S}_{+,1}^{n^2-1}$ is inconsistent with the RIP condition (1.3) in the sense that the entries of C are averagely on the scale of n^{-2}, but the RIP condition requires that the entries of C be on the scale of $O(1)$. Therefore, we generalize the definition of the RIP condition to deal with the inconsistent scaling:

Definition 3.1 Given natural numbers r and s, the function $f(\cdot; M^*)$ is said to satisfy the **Restricted Isometry Property** (RIP) of rank $(2r, 2s)$ for a constant $\delta \in [0, 1)$, denoted as δ-RIP$_{2r,2s}$, if there exist constants $c_1, c_2 \geq 0$ such that $c_2/c_1 = (1+\delta)/(1-\delta)$ and

$$c_1 \|K\|_F^2 \leq \left[\nabla^2 f(M; M^*)\right](K, K) \leq c_2 \|K\|_F^2 \tag{3.1}$$

holds for all matrices $M, K \in \mathbb{R}^{n\times n}$ such that $\operatorname{rank}(M) \leq 2r$, $\operatorname{rank}(K) \leq 2s$.

The above definition of the RIP condition is scale-free in the sense that for any constant $c > 0$, the function $cf(\cdot; M^*)$ satisfies the δ-RIP$_{2r,2s}$ condition if and only if $f(\cdot; M^*)$ satisfies the same condition.

Since the instances satisfying the RIP condition have a benign optimization landscape, we expect that the complexity metric is upper-bounded for those instances. By suitably generalizing the definitions of $\mathbb{D}_\alpha(C, u^*)$ and $\mathcal{D}$, we provide an upper bound for problem (1.2) under the RIP condition. Note that the ground truth M^* is not necessarily rank-1 in this part. Instead, we assume that $M^* = U^*(U^*)^T$ is rank-r, where U^* belongs to $\mathbb{R}^{n\times r}$. For problem (1.2), each instance is defined by the loss function $f(\cdot;\cdot)$ and the ground truth M^*. We assume that the M^* is a global optimum of the loss function, namely,

$$f(M^*; M^*) = \min_{K\in\mathbb{R}^{n\times n}} f(K; M^*), \quad \forall M^* \in \mathbb{R}^{n\times n} \quad \text{s.t.} \quad M^* \succeq 0,\ \operatorname{rank}(M^*) = r. \tag{3.2}$$

In the special case when $f(\cdot;\cdot)$ is the weighted ℓ_2-loss function in (1.7), the above condition implies that $C_{ij} \geq 0$ for all $i, j \in [n]$. Similar to the normalization constraint $C \in \mathbb{S}_{+,1}^{n^2-1}$, we assume that objective function $f(\cdot; M^*)$ is normalized in the sense that

$$\sum\nolimits_{i,j\in[n]} \left[f\left(M^* + E_{ij}; M^*\right) - f\left(M^*; M^*\right)\right] = 1. \tag{3.3}$$

For the normalization constraint $u^* \in \mathbb{S}_1^{n-1}$, we assume that the global truth M^* satisfies

$$\|U^*\|_1 = 1. \tag{3.4}$$

The set of degenerate instances is given by

$$\mathcal{D} := \left\{(f, M^*) \,\middle|\, f(\cdot;\cdot) \text{ and } M^* \text{ satisfy}(3.2) - (3.4), \right.$$
$$\left. \exists M \neq M^* \text{s.t.} f(M; M^*) = f(M^*; M^*),\ M^* \succeq 0,\ \operatorname{rank}(M^*) = r \right\}.$$

The "entry-wise ℓ_1-norm" between two arbitrary functions $h^1(\cdot)$ and $h^2(\cdot)$ with the domain $\mathbb{R}^{n\times n}$ is defined as the restricted ℓ_∞-Lipschitz constant of $h^1 - h^2$. Namely, we define $\|h^1 - h^2\|_1$ to be

$$\|h^1 - h^2\|_1 := \sup_{K,L\in\mathbb{R}^{n\times n}} \frac{\left|\left(h^1(K) - h^2(K)\right) - \left(h^1(L) - h^2(L)\right)\right|}{\max_{i,j\in[n]}(K_{ij} - L_{ij})^2}$$
$$\text{s.t.} \quad K \neq L, \quad \operatorname{rank}(K - L) \leq 2r.$$

For every constant $\alpha \in [0, 1]$, the distance between two instances (f, M^*) and $(\tilde{f}, \tilde{M}^*)$ is defined as

$$\operatorname{dist}_\alpha\left[(f, M^*), (\tilde{f}, \tilde{M}^*)\right] := \alpha\|f(\cdot; M^*) - \tilde{f}(\cdot; \tilde{M}^*)\|_1 + (1-\alpha)\|U^* - \tilde{U}^*\|_1,$$

where $U^*, \tilde{U}^* \in \mathbb{R}^{n\times r}$ satisfy $U^*(U^*)^T = M^*$ and $\tilde{U}^*(\tilde{U}^*)^T = \tilde{M}^*$. Finally, the complexity metric is given by

$$\mathbb{D}_\alpha(f, M^*) := \left[\inf_{(\tilde{f},\tilde{M}^*)\in\mathcal{D}} \operatorname{dist}_\alpha\left[(f, M^*), (\tilde{f}, \tilde{M}^*)\right]\right]^{-1}. \tag{3.5}$$

We note that the definitions of $\mathcal{D}$ and $\mathbb{D}_\alpha(f, M^*)$ are consistent with those of instance (1.7). The following theorem provides an upper bound on the complexity metric of any instance satisfying the $\text{RIP}_{2r,2r}$ condition.

Theorem 3.1 *Let $\alpha \in [0, 1]$ and $\delta \in [0, 1)$ be two constants. Suppose that the function $f(\cdot; M^*)$ satisfies the δ-$RIP_{2r,2r}$ condition and the normalization constraint (3.3), where r is the rank of M^*. Then, it holds that*

$$\mathbb{D}_\alpha\left(f, M^*\right) \leq \frac{n^2(1+\delta)}{\alpha(1-\delta)}$$

Proof We fix the instance (f, M^*) and assume that $\left(\tilde{f}, \tilde{M}^*\right) \in \mathcal{D}$. Suppose that the matrix $M \neq \tilde{M}^*$ satisfies

$$\tilde{f}\left(M; \tilde{M}^*\right) = \tilde{f}\left(\tilde{M}^*; \tilde{M}^*\right).$$

We first consider the case when $M \neq M^*$. In this case, we can estimate that

$$\begin{aligned}
&\|f\left(\cdot; M^*\right) - \tilde{f}\left(\cdot; \tilde{M}^*\right)\|_1 \\
&\geq \frac{\left|\left[f\left(M; M^*\right) - \tilde{f}\left(M; \tilde{M}^*\right)\right] - \left[f\left(M^*; M^*\right) - \tilde{f}\left(M^*; \tilde{M}^*\right)\right]\right|}{\max_{i,j\in[n]}\left(M_{ij} - M^*_{ij}\right)^2} \\
&= \frac{\left|\left[f\left(M; M^*\right) - f\left(M^*; M^*\right)\right] + \left[\tilde{f}\left(M^*; \tilde{M}^*\right) - \tilde{f}\left(M; \tilde{M}^*\right)\right]\right|}{\max_{i,j\in[n]}\left(M_{ij} - M^*_{ij}\right)^2}
\end{aligned}$$

$$= \frac{\left|\left[f\left(M;M^*\right)-f\left(M^*;M^*\right)\right]+\left[\tilde{f}\left(M^*;\tilde{M}^*\right)-\tilde{f}\left(\tilde{M}^*;\tilde{M}^*\right)\right]\right|}{\max_{i,j\in[n]}\left(M_{ij}-M^*_{ij}\right)^2}$$

$$\geq \frac{f\left(M;M^*\right)-f\left(M^*;M^*\right)}{\max_{i,j\in[n]}\left(M_{ij}-M^*_{ij}\right)^2} \geq \frac{(c_1/2)\cdot\|M-M^*\|_F^2}{\max_{i,j\in[n]}\left(M_{ij}-M^*_{ij}\right)^2} \geq \frac{c_1}{2}, \tag{3.6}$$

where c_1 is the constant in the RIP condition of $f(\cdot; M^*)$. The second inequality is due to

$$f(M;M^*)-f\left(M^*;M^*\right)\geq 0, \quad \tilde{f}\left(M^*;\tilde{M}^*\right)-\tilde{f}\left(\tilde{M}^*;\tilde{M}^*\right)\geq 0.$$

The second last inequality follows from the global optimality of M^* and the second inequality after inequality (12) in [59], namely,

$$f\left(M;M^*\right)\geq f\left(M^*;M^*\right)+\frac{c_1}{2}\|M-M^*\|_F^2, \quad \forall M\in\mathbb{R}^{n\times n},\ \operatorname{rank}(M)\leq r.$$

Now, we provide a lower bound on c_1. Using the normalization constraint (3.3) and the stationarity of M^*, it holds that

$$1=\sum_{i,j\in[n]}\left[f\left(M^*+E_{ij};M^*\right)-f\left(M^*;M^*\right)\right]\leq\frac{c_2}{2}\cdot\sum_{i,j\in[n]}\|E_{ij}\|_F^2=\frac{c_2n^2}{2},$$

which implies that $c_2\geq 2n^{-2}$. Using the relation $c_2/c_1=(1+\delta)/(1-\delta)$, we obtain that

$$c_1\geq\frac{2(1-\delta)}{n^2(1+\delta)}.$$

By substituting into inequality (3.6), it follows that

$$\|f(\cdot;M^*)-\tilde{f}(\cdot;\tilde{M}^*)\|_1\geq\frac{1-\delta}{n^2(1+\delta)}.$$

which leads to $\operatorname{dist}_\alpha\left[(f,M^*),\left(\tilde{f},\tilde{M}^*\right)\right]\geq\alpha(1-\delta)/\left[n^2(1+\delta)\right]$. Now, the desired bound on $\mathbb{D}_\alpha(f,M^*)$ follows from taking the inverse. In the case when $M=M^*$, we can replace M with $\tilde{M}^*$ and the proof can be done in the same way. □

We note that the upper bound on $\mathbb{D}_\alpha(C,u^*)$ is increasing in δ, which is consistent with the intuition that a smaller δ will lead to a better optimization landscape. Moreover, in the case when $\alpha(1-\delta)=\Theta(1)$, the upper bound is on the order of $O\left(n^2\right)$, which is $O(n)$ larger than the minimum possible complexity metric in Theorem 2.4. Now, we provide a remedy to the aforementioned issue for problem (1.7). With the knowledge about the incoherence of the global solution, we can improve the upper bound on the complexity metric.

Theorem 3.2 *Suppose that the instance $\mathcal{MC}(C, u^*)$ satisfies the δ-RIP$_{2,2}$ condition and u^* has incoherence μ. Then, it holds that*

$$\mathbb{D}_\alpha\left(C, u^*\right) \leq \max\left\{\frac{n(1+\delta)}{4\alpha(1-\delta)}, \frac{1}{2(1-\alpha)\mu}\right\} \times \min\left\{\left(\frac{1}{\mu}-\frac{1}{n}\right)^{-1}, 3\mu\right\}.$$

The proof of Theorem 3.2 can be found in Appendix B.1. From the above theorem, we can use the weight α to control the balance between the RIP constant δ and the incoherence μ. If we choose $1-\alpha = \Theta(n^{-1})$, then the complexity can be upper-bounded by

$$\mathbb{D}_\alpha\left(C, u^*\right) = \mu n \cdot \max\left\{O\left(\frac{1+\delta}{1-\delta}\right), O\left(\frac{1}{\mu}\right)\right\} = O\left(\mu n \cdot \frac{1+\delta}{1-\delta}\right).$$

In addition, if it holds that $\mu = O(1)$ and $(1-\delta)^{-1} = O(1)$, then the complexity is upper-bounded by $O(n)$, which matches the minimum possible complexity in Theorem 2.4 up to a constant. Although the complexity metric may have a large value for extreme instances (i.e., instances with a large incoherence), the complexity of regular instances achieves the optimal value up to a constant. Furthermore, we conjecture in Sect. 4 that the condition $\mathbb{D}_\alpha(C, u^*) = O(n\mu/\alpha)$ is sufficient to guarantee the success of local search methods. Assuming that this conjecture is true, then the condition $(1-\delta)^{-1} = O(1)$ alone is sufficient to guarantee that the optimization landscapes are benign regardless of the value of the incoherence μ. This is consistent with the existing results on the RIP condition. We conclude the discussion of instances with the RIP condition by showing that the dependence of δ in Theorem 3.2 is tight up to a constant.

Theorem 3.3 *Suppose that $n \geq 4$, $\alpha \in [0, 1]$, $\mu \in [1, n]$ and $\delta \in [0, 1)$. Let $\ell := \lceil n/\mu \rceil$. Then, there exists an instance $\mathcal{MC}(C, u^*)$ such that $\mathcal{MC}(C, u^*)$ satisfies the δ-RIP$_{2,2}$ condition, u^* has incoherence μ and*

$$\mathbb{D}_\alpha\left(C, u^*\right) \geq \frac{n(1+\delta)}{4\alpha(1-\delta)} \cdot \min\left\{\frac{n\mu}{\mu\ell-\mu}, \mu\right\}.$$

The proof of Theorem 3.3 is similar to that of Theorem 3.2 and can be found in the online version [60].

3.2 Matrix completion problem: Bernoulli model and incoherence condition

Next, we consider instances $\mathcal{MC}(C, u^*)$ of problem (1.7) where the global solution u^* is μ-incoherent and the random weight matrix C obeys the Bernoulli model. Similar to the RIP condition, we need to generalize the definition of the Bernoulli model under the normalization constraint.

Definition 3.2 Given the sampling rate $p \in (0, 1]$, a random matrix $C \in \mathbb{S}_{+,1}^{n^2-1}$ is said to obey the **Bernoulli model** if

$$C_{ij} = \frac{\delta_{ij}}{\sum_{k,\ell \in [n]} \delta_{k\ell}}, \quad \forall i, j \in [n],$$

where $\{\delta_{k\ell} | k, \ell \in [n]\}$ are independent Bernoulli random variables with the parameter p.

We note that the above model is well defined only when $\sum_{i,j} \delta_{ij} > 0$, which happens with probability $1 - (1-p)^{n^2} \geq 1 - \exp(-n^2 p)$. This probability is sufficiently large if $n^2 p \gg 1$. In [13], the authors showed that $p \geq \Theta(\mu \log n / n)$ is necessary and under this condition, the success probability is at least $1 - O(n^{-\mu n})$. Therefore, we only focus on the case when the event $\sum_{i,j} \delta_{ij} > 0$ happens. In the existing literature [12, 17, 26], the instances obeying the Bernoulli model are proven to have no spurious local minima. We show that our complexity metric is able to characterize this property by proving an upper bound on the complexity metric.

Theorem 3.4 *Given $\mu \in [1, n]$ and $p \in (0, 1]$, suppose that the weight matrix C obeys the Bernoulli model with the parameter p and that u^* has incoherence μ. If $\eta > 2$ is a constant and the sampling rate satisfies*

$$p \geq \min\left\{1, \frac{16(1 + \eta\mu)\log n + 16}{n}\right\},$$

it holds with probability at least $1 - 3n^{-\eta/2+1}$ that

$$\mathbb{D}_\alpha\left(C, u^*\right) \leq \max\left\{\frac{3n}{4\alpha}, \frac{1}{2(1-\alpha)\mu}\right\} \times \min\left\{\left(\frac{1}{\mu} - \frac{1}{n}\right)^{-1}, 3\mu\right\}.$$

The proof of Theorem 3.4 is similar to that of Theorem 3.2. By Theorem 3.4, if $1 - \alpha = \Theta\left(n^{-1}\mu^{-1}\right)$, then the complexity of instances obeying the Bernoulli model is on the order of $\Theta\left[n^2\mu/(n-\mu)\right]$. If the incoherence $\mu = O(1)$, the complexity is on the order of $O(n)$, which matches the minimum possible complexity up to a constant. Therefore, the proposed metric can also serve as a good indicator for the matrix completion problem with the Bernoulli model. Finally, we note that the bound $p \geq \Theta(\mu \log n/n)$ is optimal up to a constant [13]; see also the discussions in Appendix E of [24].

Finally, we note that problem (1.7) may still have spurious local minima when the sampling probability p and the incoherence μ satisfy the condition in Theorem 3.4. In the existing literature, the global convergence of randomly initialized local search methods is established for problem (1.7) only under an extra regularizer or an extra constraint on the incoherence of u. That being said, our proposed complexity metric correctly reflects the common sense that the matrix completion problem is generally easier to solve when the incoherence is small or when the sampling rate p is large. When the complexity is small, it is possible to apply local search methods to find the ground truth. The local search methods may be different for different classes of low-rank matrix optimization problems. In addition, the new complexity metric has the advantage that it is able to simultaneously capture the RIP condition, the incoherence condition, and potentially other existing complexity metrics.

3.3 One-parameter class of instances

In Sects. 3.1 and 3.2, we provided several upper bounds on the complexity metric. In this part, we consider a class of instances that are parameterized by a single parameter $\epsilon \in [0, 1]$. Intuitively, when the parameter grows from 0 to 1, the optimization landscape of the instance becomes more benign. Unlike the previous results in this section, the analysis of the small parameter case provides necessary conditions for the existence of spurious local minima. More specifically, we fix $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ to be an unweighted undirected graph without self-loops, where the node set is $\mathbb{V} = [n]$. We consider the maximal independent set of $\mathbb{G}$, which is defined as follows:

Definition 3.3 For an undirected graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, a set $\mathcal{S} \subset \mathbb{V}$ is called an *independent set* if no two nodes in $\mathcal{S}$ are adjacent. The set $\mathcal{S}$ is called a *maximal independent set* if it is an independent set with the maximum number of nodes.[2]

Suppose that $\mathcal{S} \subset [n]$ is a maximal independent set of $\mathbb{G}$. For every $\epsilon \in [0, 1]$, the instance $\mathcal{MC}(C^\epsilon, u^*)$ is defined by

$$\begin{aligned} &C^\epsilon_{ij} := \epsilon/Z_\epsilon, \quad \forall i, j \in \mathcal{S} \ \text{ s.t. } \ i \neq j; \quad C^\epsilon_{ij} := 1/Z_\epsilon, \quad \text{if } \{i, j\} \in \mathbb{E}; \\ &C^\epsilon_{ii} := 1/Z_\epsilon, \quad \forall i \in [n], \quad C^\epsilon_{ij} := 0, \quad \text{otherwise}, \\ &u^*_i := 1/m, \quad \forall i \in \mathcal{S}; \quad u^*_i := 0, \quad \forall i \notin \mathcal{S}, \end{aligned} \tag{3.7}$$

where $m := |\mathcal{S}|$ and $Z_\epsilon := 2|\mathbb{E}| + n + m(m-1)\epsilon$ is the normalization constant. In the remainder of this subsection, we assume without loss of generality that $\mathcal{S} = [m]$.

First, we study for what values of ϵ the instance $\mathcal{MC}(C^\epsilon, u^*)$ has benign landscape or has spurious local minima. The following theorem guarantees that the threshold $\epsilon = \Theta\left(m^{-1}\right) = \Theta(\mu/n)$ separates the regimes where the instance possesses and does not possess spurious local minima, where $\mu := n/m$ denotes the incoherence of u^*.

Theorem 3.5 *If $\epsilon \geq \Theta(m^{-1})$, the instance $\mathcal{MC}(C^\epsilon, u^*)$ does not have spurious second-order critical points[3] (SSCPs), namely, all second-order critical points are global minima associated with the ground truth solution M^*. If $\epsilon = O\left(m^{-1}\right)$, the instance $\mathcal{MC}(C^\epsilon, u^*)$ has at least $O\left(2^{m/2}\right)$ spurious local minima.*

The proof of Theorem 3.5 can be found in Appendix B.3. In the case when $m = 2$, the proof of Theorem 3.5 (more specifically, Theorem B.2) states that the instance

---

[2] We note that this definition is different from the common definition of a maximum independent set, which only requires that a maximum independent set is not a proper subset of an independent set.

[3] A point $u \in \mathbb{R}^n$ is called a spurious second-order critical point if it satisfies the first-order and the second-order necessary optimality conditions and $uu^T \neq u^*(u^*)^T$.

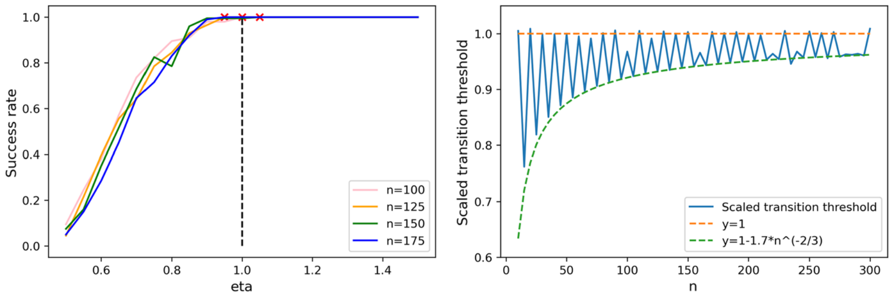

Fig. 2 The left plot shows the transitions of the success rate of the gradient descent algorithm when $n = 100, 125, 150, 175$. The red "×" sign refers to the *transition threshold*, i.e., the smallest value of η that attains 100% success rate. In the right plot, the transition thresholds of η are compared with the curves $y = 1$ and $y = 1 - 1.7(n+1)^{-2/3}$

$\mathcal{MC}(C^\epsilon, u^*)$ has spurious local minima if $\epsilon < 1/3$. The condition $\epsilon = 1/3$ corresponds to the δ-$\text{RIP}_{2,2}$ condition holding with $\delta = 1/2$. Therefore, the RIP constant $\delta < 1/2$ is necessary for the instance $\mathcal{MC}(C^\epsilon, u^*)$ to have no spurious local minima. Combined with the results in [7, 59], we can see that the one-parameter group $\mathcal{MC}(C^\epsilon, u^*)$ also contains difficult instances of the general problem (1.2).

Furthermore, we note that the constants in the proof of Theorem 3.5 are not optimal. We conjecture that the instance $\mathcal{MC}(C^\epsilon, u^*)$ has spurious solutions if $\epsilon < (m+1)^{-1} + o(m^{-1})$ and does not have spurious solutions if $\epsilon > (m+1)^{-1} + o(m^{-1})$. We numerically verify this conjecture in the special case when $m = n$. In numerical examples, we consider the scaled parameter $\eta := (n+1)\epsilon$. For each instance, we implement the randomly initialized gradient descent algorithm for 200 times and check the number of implements for which the distance between the last iterate and $\pm u^*$ has Frobenius norm at most 10^{-5}. The results are plotted in Fig. 2.

In the left plot, we can see that in most cases, the success rate grows with the parameter η, which is proportional to ϵ. This indicates that the optimization landscape becomes more benign when ϵ is larger. In addition, the transition thresholds of η are very close to 1 (to be more accurate, the thresholds of η are between 0.95 and 1.05). This observation is consistent with our conjecture. In the right plot, we compare the transition thresholds of η against the constant number 1. We observe that the thresholds are approximately located between 1 and $1 - 1.7(n+1)^{-2/3}$, which implies that the original thresholds of ϵ are between $(n+1)^{-1}$ and $(n+1)^{-1} - 1.7(n+1)^{-5/3}$. Hence, the thresholds become close to $(n+1)^{-1}$ when n is large, which is also consistent with our conjecture. Moreover, we can see that the threshold of η is not monotone in n and is slightly smaller when n is odd.

Finally, we transform the estimates on the parameter ϵ to the complexity metric.

Theorem 3.6 *Suppose that $n \ge m \ge 36$, $\alpha \in [0, 1]$ and $\epsilon \in [0, 1]$. Then, the following statements hold true:*

1. If

$$\mathbb{D}_\alpha(C^\epsilon, u^*) \le \left[\frac{36\alpha}{n^2} + \min\left\{72\alpha \cdot \frac{m}{n^2}, 2(1-\alpha)\right\}\right]^{-1},$$

then the instance $\mathcal{MC}(C^\epsilon, u^*)$ has no spurious local minima;

2. *If*

$$\mathbb{D}_\alpha(C^\epsilon, u^*) \geq \frac{18}{17} \max\left\{\frac{13n^2}{2\alpha}, \frac{1}{2(1-\alpha)}\right\},$$

then the instance $\mathcal{MC}(C^\epsilon, u^*)$ has spurious local minima.

The proof of Theorem 3.6 can be found in Appendix B.4. In the case when $1 - \alpha \geq \Theta(m/n^2)$, the upper bound on $\mathbb{D}_\alpha(C^\epsilon, u^*)$ is on the order of $O(n\mu/\alpha)$, where $\mu := n/m$ is the incoherence of u^*. This result is consistent with the upper bounds in Sects. 3.1 and 3.2. In addition, the RIP constant is $1 - O(1/m)$ if $\epsilon = O(1/m)$, which shows that the proposed complexity metric can provide better guarantees on the optimization complexities than the RIP constant. On the other hand, the lower bound in Theorem 3.6 is on the order of $O(n^2/\alpha)$ in the case when $1 - \alpha \geq \Theta(n^{-2})$.

In summary, we have provided a consistent upper bound on the complexity metric that is on the order of $\Theta(n\mu/\alpha)$ for all three examples ($\Theta[n\mu/\alpha \cdot (1+\delta)/(1-\delta)]$ for the RIP case) if we choose $1 - \alpha = O(n^{-1})$. These theoretical results provide strong evidence that our proposed complexity metric is able to capture the properties of the optimization landscape for several different models, even when other existing conditions fail to provide theoretical guarantees; see the comparison of the condition and our complexity metric in Sect. 3.3. In Sect. 4, we make some conjectures based on these observations and provide a partial theoretical explanation.

4 Theoretical results for general instances

In this section, we provide a theoretical analysis for the proposed complexity metric (2.2) on the general problem (1.7). Intuitively, we expect the problem (1.7) to have a benign landscape when the complexity metric is small and vice versa. We first prove that the proposed complexity metric is able to provide a sufficient condition on the absence of SSCPs of problem (1.7). Then, we construct another complexity metric that lower-bounds the metric (2.1) and show that the alternative complexity metric is able to provide necessary conditions on the absence of SSCPs.

Recalling the analysis in Sect. 3, one might have the following questions: Suppose that $1 - \alpha \geq \Theta(n^{-1})$ and the solution u^* is μ-incoherent. Can we find two constants $\delta, \Delta > 0$ such that

1. If $\mathbb{D}_\alpha(C, u^*) \leq \delta\mu n/\alpha$, the instance $\mathcal{MC}(C, u^*)$ has no SSCPs;
2. If $\mathbb{D}_\alpha(C, u^*) \geq \Delta n^2/\alpha$, the instance $\mathcal{MC}(C, u^*)$ has SSCPs?

Suppose that the first property in the above question holds. The results in Sect. 3.1 imply that the proposed complexity metric guarantees the absence of SSCPs when the RIP constant is $O[(\delta - 1)/(\delta + 1)]$, which is independent of μ. In addition, the matrix completion problem under the Bernoulli model does not have SSCPs when $p \geq O(\mu \log n/n)$, which matches the lower bound in [13]. In Sect. 4.1, we prove a weaker version of the first property in the case when α is equal to α^* or $\alpha^\diamond$, which are defined in Sect. 2. We note that both α^* and $\alpha^\diamond$ satisfy the condition that

$1 - \alpha = \Theta(n^{-1})$. On the other hand, in Sect. 4.2, we refute the second property in the above question by constructing counterexamples. This observation implies that similar to the RIP constant and the incoherence, the proposed complexity metric cannot provide necessary conditions on the absence of spurious local solutions. However, if we substitute the degenerate set $\mathcal{D}$ with a slightly smaller set, we prove that the complexity metric is able to provide a necessary condition.

4.1 Small complexity case

We first consider instances with a small complexity metric. In the case when α is equal to α^* or $\alpha^\diamond$, we prove that $\mathbb{D}_\alpha(C, u^*) \le \delta n/\alpha$ serves as a sufficient condition for the absence of SSCPs, where $\delta > 0$ is an absolute constant. Since the incoherence μ is at least 1, the aforementioned condition is weaker than the first property in the aforementioned question. By Theorem 2.4, the minimum possible value of the complexity metric is on the order of $O(n/\alpha)$. In this subsection, we show that the constant δ can be chosen such that $\delta n/\alpha$ is strictly larger than the minimum possible complexity. The following theorem deals with the case when $\alpha = \alpha^*$.

Theorem 4.1 *Suppose that $n \ge 5$ and $\alpha = \alpha^*$. Then, there exists a constant $\delta > 1/4$ such that for every instance $\mathcal{MC}(C, u^*)$ satisfying*

$$\mathbb{D}_{\alpha^*}\left(C, u^*\right) \le \delta n/\alpha^*,$$

the instance $\mathcal{MC}(C, u^*)$ does not have any SSCPs.

Since the minimum possible complexity metric is $n/(4\alpha^*)$, the upper bound in Theorem 4.1 is *non-trivial* in the sense that there exist instances satisfying the inequality. By Theorem 2.5, the minimum complexity metric $n/(4\alpha^*)$ is only attained by instances in $\mathcal{M}$, where

$$\mathcal{M} := \left\{ \left(C, u^*\right) \,\middle|\, |u_i^*| = \frac{1}{n},\ C_{ii} = 0,\ \forall i \in [n],\ C_{ij} = \frac{1}{n(n-1)},\ \forall i, j \in [n],\ i \ne j \right\}.$$

In the next lemma, we prove the strict-saddle property [48] of the ℓ_1-norm for instances in $\mathcal{M}$, which can be viewed as a robust version of the absence of SSCPs.

Lemma 4.1 *Suppose that $n \ge 2$ and $\left(C^0, u^0\right) \in \mathcal{M}$. Then, there exist a positive constant η_0 and two positive-valued functions $\beta(\eta)$ and $\gamma(\eta)$ such that for all $\eta \in (0, \eta_0]$ and $u \in \mathbb{R}^n$, at least one of the following properties holds:*

1. $\min\{\|u - u^*\|_1, \|u + u^*\|_1\} \le \eta$;
2. $\|\nabla g(u; C, u^*)\|_\infty \ge \beta(\eta)$;
3. $\lambda_{min}[\nabla^2 g(u; C, u^*)] \le -\gamma(\eta)$.

We then show that after a sufficiently small perturbation to any point $\left(C^0, u^0\right) \in \mathcal{M}$, the new instance does not have any SSCPs.

Lemma 4.2 *Suppose that $n \geq 3$. There exists a small positive constant ϵ such that for every pair $(C^0, u^0) \in \mathcal{M}$ and $(\tilde{C}, \tilde{u}^*)$ satisfying*

$$\alpha^* \|\tilde{C} - C^0\|_1 + (1 - \alpha^*)\|\tilde{u}^* - u^0\|_1 < \epsilon,$$

the instance $\mathcal{MC}(\tilde{C}, \tilde{u}^*)$ does not have SSCPs.

The proofs of the last two lemmas involve several standard calculations and can be found in the online version [60]. Now, we prove the existence of a non-trivial upper bound on the metric.

Proof of Theorem 4.1 Let ϵ be the constant in Lemma 4.2. We consider the compact set

$$\mathcal{C} := \left\{ \left(C, u^*\right) \middle| \|C\|_1 = \|u^*\|_1 = 1, \right.$$
$$\left. \alpha^* \|C - C^0\|_1 + (1 - \alpha^*)\|u^* - u^0\|_1 \geq \epsilon, \quad \forall \left(C^0, u^0\right) \in \mathcal{M} \right\}.$$

Since the minimum possible complexity metric $n/(4\alpha^*)$ is only attained by points in $\mathcal{M}$, it holds that

$$\mathbb{D}_{\alpha^*}(\mathcal{C}) := \max_{(C,u^*) \in \mathcal{C}} \mathbb{D}_{\alpha^*}\left(C, u^*\right) > n/(4\alpha^*).$$

Therefore, choosing

$$\delta := (\alpha^*/n) \cdot \mathbb{D}_{\alpha^*}(\mathcal{C}) > 1/4,$$

we have

$$\mathbb{D}_{\alpha^*}\left(C, u^*\right) \leq \delta n/\alpha^* \implies \left(C, u^*\right) \notin \mathcal{C} \implies \text{the instance } \mathcal{MC}\left(C, u^*\right) \text{ has no SSCPs.}$$

This completes the proof. □

The case when $\alpha = \alpha^\diamond$ can be analyzed in a similar way. We note that the strict-saddle property of the instances in Theorem 2.6 has been established in [32]. Hence, we present the results in the following theorem and omit the proof.

Theorem 4.2 *Suppose that $n \geq 5$ and $\alpha = \alpha^\diamond$. Then, there exists a constant $\delta > 1/4$ such that for every pair (C, u^*) satisfying*

$$\mathbb{D}_{\alpha^\diamond}\left(C, u^*\right) \leq \delta n(n+2)/(n+1),$$

the instance $\mathcal{MC}\,(C, u^*)$ does not have any SSCPs.

Similar to Theorem 4.1, since the minimum possible complexity metric is attained with $\delta = 1/4$, the upper bound in Theorem 4.2 is non-trivial.

4.2 Large complexity case

In this subsection, we first refute the second property in the question that we asked at the beginning of Sect. 4 and then refine its statement to make it hold true. We note that the RIP condition and the incoherence condition cannot provide necessary conditions for the absence of SSCPs either. Namely, there exist instances that satisfy the δ-RIP$_{2,2}$ condition with δ as high as 1 which do not have SSCPs. Similarly, in the case when the incoherence of the global solution is n, it is still possible to have an instance of the matrix completion problem without any SSCPs. In other words, although small values for the RIP constant and incoherence guarantee the absence of spurious solutions, these notions cannot capture the complexity of the problem since there are low-complexity problems with large values for these parameters. We first show that our new metric suffers from the same shortcoming, but we then propose a simple refinement to address this issue.

Example 1 Suppose that the weight matrix and the ground truth are

$$C^\delta := \frac{1}{1+3\delta}\begin{bmatrix}1 & \delta\\ \delta & \delta\end{bmatrix}, \quad u^* := \begin{bmatrix}1\\0\end{bmatrix},$$

where $\delta \geq 0$ is a constant. One can verify that $\pm u^*$ are the only local minima to the instance $\mathcal{MC}(C^\delta, u^*)$ for all $\delta > 0$. However, in the case when $\delta = 0$, the instance $\mathcal{MC}(C^0, u^*)$ has the set of global solutions

$$\pm\begin{bmatrix}1\\c\end{bmatrix}, \quad \forall c \in \mathbb{R}.$$

Moreover, we consider the case when both components of u^* are measured, where the instance $\mathcal{MC}(\tilde{C}^\epsilon, \tilde{u}^\epsilon)$ is defined by

$$\tilde{C}^\epsilon := \frac{1}{1+\epsilon}\begin{bmatrix}1 & 0\\ 0 & \epsilon\end{bmatrix}, \quad \tilde{u}^\epsilon := \frac{1}{1+\epsilon}\begin{bmatrix}1\\ \epsilon\end{bmatrix},$$

where ϵ is a positive constant. One can verify that the pair $(\tilde{C}^\epsilon, \tilde{u}^\epsilon)$ belongs to $\mathcal{D}$ for all $\epsilon > 0$. Setting δ and ϵ to be small enough, the instances $\mathcal{MC}(C^\delta, u^*)$ and $\mathcal{MC}(\tilde{C}^\epsilon, \tilde{u}^\epsilon)$ can be arbitrarily close to each other in the sense that

$$\alpha\|C^\delta - \tilde{C}^\epsilon\|_1 + (1-\alpha)\|u^* - \tilde{u}^\epsilon\|_1 = O(\alpha\delta + \epsilon).$$

Therefore, the complexity metric of $\mathcal{MC}(C^\delta, u^*)$ can be arbitrarily large. This example shows that instances without SSCPs can be arbitrarily close to those in $\mathcal{D}$, which have non-unique global solutions.

Nevertheless, we derive a lower bound on the complexity metric (2.2) by constructing a subset of $\mathcal{D}$, which allows obtaining a necessary condition. Intuitively, if an instance has multiple global minima, these global minima are still locally optimal after a sufficiently small perturbation to the instance. To ensure the "robustness" of the local optimality, we require the positive-definiteness of the Hessian matrix. For each instance $\mathcal{MC}(C, u^*)$, let $\mathbb{G}_{1k}(C, u^*)$ for all $k \in [n_1]$ be the connected components of $\mathbb{G}_1(C, u^*)$, where n_1 is the number of connected components. Moreover, we use $\mathcal{I}_{1k}(C, u^*)$ to denote the node set of $\mathbb{G}_{1k}(C, u^*)$ for all $k \in [n_1]$. We define the following subset of $\mathcal{D}$:

$$\mathcal{SD} := \{\left(C, u^*\right) \in \mathcal{D} \mid \mathbb{G}_{1k}\left(C, u^*\right) \text{ is not bipartite for all } k \in [n_1], \\ \mathbb{G}_1\left(C, u^*\right) \text{ is disconnected}, \quad \mathcal{I}_{00}\left(C, u^*\right) = \emptyset\}.$$

The following theorem provides a characterization of the Hessian matrix at global solutions for pairs in $\mathcal{SD}$.

Theorem 4.3 *Suppose that $(C, u^*) \in \mathcal{SD}$. Then, the Hessian matrix is positive definite at all global solutions of the instance $\mathcal{MC}(C, u^*)$.*

The proof of Theorem 4.3 can be found in the appendix. Using the positive-definiteness of the Hessian matrix, we are able to apply the implicit function theorem to guarantee the existence of spurious local minima in a neighborhood of each instance in $\mathcal{SD}$; see Appendix C.2 for more details. The global guarantee can be established by considering closed subsets of $\mathcal{SD}$. For every constant $\epsilon \geq 0$, we consider the closed subset $\mathcal{SD}_\epsilon$, which is defined as

$$\mathcal{SD}_\epsilon := \left\{ \left(C, u^*\right) \in \mathcal{SD} \mid C_{ij} \in \{0\} \cup [\epsilon, 1], \quad \forall i, j \in [n], \\ |u_i^*| \in \{0\} \cup [\epsilon, 1], \quad \forall i \in [n]\right\}.$$

Basically, the extra condition in the definition of $\mathcal{SD}_\epsilon$ requires that the nonzero components of C and u^* be at least ϵ. We can verify that the set $\mathcal{SD}_\epsilon$ is a compact set and for every $\epsilon_n \to 0$, it holds that

$$\lim_{n\to\infty} \cup_{i=1}^n \mathcal{SD}_{\epsilon_i} = \mathcal{SD}_0 = \mathcal{SD}.$$

Now, we define the alternative complexity metric

$$\mathbb{D}_{\alpha,\epsilon}\left(C, u^*\right) := \left[\min_{(\tilde{C},\tilde{u}^*)\in\mathcal{SD}_\epsilon} \alpha\|C - \tilde{C}\|_1 + (1-\alpha)\|u^* - \tilde{u}^*\|_1\right]^{-1}. \tag{4.1}$$

Since $\mathcal{SD}_\epsilon$ is a subset of $\mathcal{D}$, it holds that

$$\mathbb{D}_{\alpha,\epsilon}\left(C, u^*\right) \leq \mathbb{D}_\alpha\left(C, u^*\right).$$

Similar to Theorem 2.3, we can prove the following relation:

$$\overline{\mathcal{SD}} = \left\{\left(C, u^*\right) \mid C \in \mathbb{S}_{+,1}^{n^2-1}, u^* \in \mathbb{S}_1^{n-1}, \mathbb{G}_1\left(C, u^*\right) \text{ is disconnected}\right\} \\ \cup \left\{\left(C, u^*\right) \mid C \in \mathbb{S}_{+,1}^{n^2-1}, u^* \in \mathbb{S}_1^{n-1}, \mathcal{I}_{00}\left(C, u^*\right) \text{ is not empty}\right\}.$$

Hence, the closure of $\mathcal{SD}$ is a proper subset of $\overline{\mathcal{D}}$. Combining with the fact that $\mathcal{SD}_\epsilon$ is a subset of $\mathcal{SD}$, the metric $\mathbb{D}_{\alpha,\epsilon}(C, u^*)$ is not equivalent to $\mathbb{D}_\alpha(C, u^*)$. Using the compactness of $\mathcal{SD}_\epsilon$, the following theorem provides a necessary condition for the existence of spurious local minima.

Theorem 4.4 *Suppose that $\epsilon > 0$ is a constant. Then, there exists a large constant $\Delta(\epsilon) > 0$ such that for every instance $\mathcal{MC}(C, u^*)$ satisfying*

$$\mathbb{D}_{\alpha,\epsilon}\left(C, u^*\right) \geq \Delta(\epsilon),$$

the instance $\mathcal{MC}(C, u^*)$ has spurious local minima.

Proof For every pair $(C, u^*) \in \mathcal{SD}_\epsilon$, Lemma C.3 implies that there exists an open neighborhood of (C, u^*) such that the desired properties hold. Now, we consider the union of such open neighborhoods over all points $(C, u^*) \in \mathcal{SD}_\epsilon$, which is an open cover of $\mathcal{SD}_\epsilon$. Using the Heine-Borel covering theorem, there exists an open sub-cover of $\mathcal{SD}_\epsilon$. Therefore, we obtain the existence of $\Delta(\epsilon)$. □

We note that the maximum possible value of $\mathbb{D}_{\alpha,\epsilon}(C, u^*)$ is $+\infty$, which is attained by instances in $\mathcal{SD}_\epsilon$. Therefore, there exist instances satisfying the condition of Theorem 4.4 and the lower bound is *non-trivial*. Using Theorem 4.4, the slightly modified complexity metric is able to provide a necessary condition on the absence of SSCPs. This result implies that our complexity metric is able to provide conditions that are much better than the RIP condition and the incoherence condition that fail to provide necessary conditions.

Finally, we conjecture that the second property in the question we asked at the beginning of the section holds for any fixed weight matrix. More specifically, we define

$$\mathbb{D}_C(u^*) := \left(\min_{(C,\tilde{u}^*)\in\overline{\mathcal{D}}} \|u^* - \tilde{u}^*\|_1 \right)^{-1}. \tag{4.2}$$

We have the following conjecture:

Conjecture 1 Suppose that $\epsilon \in [0, 1]$. Then, there exists a large constant $\Gamma(\epsilon) > 0$ such that for every instance $\mathcal{MC}(C, u^*)$ satisfying

$$C_{ij} \in \{0\} \cup [\epsilon, 1], \quad \mathbb{D}_C(u^*) \geq \Gamma(\epsilon),$$

the instance $\mathcal{MC}(C, u^*)$ has spurious local minima.
We note that the metric $\mathbb{D}_C(u^*)$ is equal to 0 if $\mathcal{MC}(C, u^*)$ satisfies the δ-RIP$_{2,2}$ condition with $\delta \in [0, 1)$.

5 Conclusions

In this work, we propose a new complexity metric for an important class of the low-rank matrix optimization problems, which has the potential to generalize major existing recovery guarantees and is applicable to a much broader set of problems. The proposed complexity metric aims to measure the complexity of the non-convex optimization landscape of each problem and quantifies the likelihood of local search methods in successfully solving each instance of the problem under a random initialization. We focus on the rank-1 generalized matrix completion problem (1.7) to mathematically prove the usefulness of the new metric from three aspects. Namely, we show that the complexity metric has a small value if the instance satisfies the RIP condition or the incoherence condition. The results in these two scenarios are consistent with the existing results on the RIP condition and the incoherence condition. In addition, we analyze a one-parameter class of instances to illustrate that the proposed metric captures the true complexity of this class as the parameter varies and has consistent behavior with the aforementioned two scenarios. This consistency implies that our proposed complexity metric is able to characterize the optimization landscapes of different applications, which the RIP condition and the incoherence condition fail to capture. Finally, we provide strong theoretical results on the generalized matrix completion problem by showing that a small value for the proposed complexity metric guarantees the absence of spurious solutions, whereas a large value for a slightly modified complexity metric guarantees the existence of spurious solutions. This also shows the superiority of this metric over the RIP condition and the incoherence condition since those notions cannot offer any necessary conditions for having spurious solutions.

Acknowledgements This work was supported by grants from ARO, AFOSR, ONR and NSF.

A Proofs in Sect. 2

A.1 Proof of Theorem 2.4

The proof of Theorem 2.4 relies on the following two lemmas, which transform the computation of $\mathbb{D}_\alpha^{min}$ into a one-dimensional optimization problem. The first lemma upper-bounds the maximum possible distance.

Lemma A.1 *Suppose that $n \geq 2$. It holds that*

$$\left(\mathbb{D}_\alpha^{min}\right)^{-1} \leq \max_{c \in \left[0, \frac{1}{n(n-1)}\right]} g(\alpha, c),$$

where the function $g(\alpha, c)$ is defined by

$$\begin{aligned} g(\alpha, c) := \min\Bigg\{ & 2(1-\alpha) \cdot \frac{n-2}{n} + 4\alpha c, \quad 4\alpha(n-1)c, \\ & 2(1-\alpha) \cdot \frac{n-4}{n} + 2\alpha\left(\frac{4}{n} - 4(n-2)c\right), \\ & 2(1-\alpha) \cdot \frac{n-3}{n} + 2\alpha\left(\frac{3}{n} - (3n-5)c\right), \end{aligned}$$

$$2(1-\alpha)\cdot\frac{n-2}{n}+2\alpha\left(\frac{2}{n}-2(n-1)c\right),$$
$$2(1-\alpha)\cdot\frac{n-1}{n}+2\alpha\left(\frac{1}{n}-(n-1)c\right)\Bigg\}.$$

We denote $g_i(\alpha, c)$ be the i-th term in the above minimization for all $i \in \{1, \ldots, 6\}$. The next lemma proves the other direction.

Lemma A.2 *Suppose that $n \geq 2$. It holds that*

$$\left(\mathbb{D}_\alpha^{min}\right)^{-1} \geq \max_{c\in\left[0,\frac{1}{n(n-1)}\right]} g(\alpha, c),$$

where the function $g(\alpha, c)$ is defined in Lemma A.1.

The proof of Lemmas A.1 and A.2 can be found in the online version [60].

Proof of Theorem 2.4 By the results of Lemmas A.1 and A.2, we only need to compute $\max_{c\in\left[0,\frac{1}{n(n-1)}\right]} g(\alpha, c)$. Let $\kappa := (1-\alpha)/\alpha \in [0, +\infty]$. We study three cases below.

Case 1. We first consider the case when $\kappa \geq 2(n-3)/[(n-4)(n-1)]$. We prove that $g(\alpha, c) = g_2(\alpha, c)$. Since $g_2(\alpha, c)$ has a larger gradient than $g_1(\alpha, c)$ and the function $g_i(\alpha, c)$ is decreasing in c for $i = 3, 4, 5, 6$, we only need to show that

$$g_i\left(\alpha, \frac{1}{n(n-1)}\right) \geq g_2\left(\alpha, \frac{1}{n(n-1)}\right), \quad \forall i \in \{1, 3, 4, 5, 6\}. \tag{A.1}$$

The above inequality with $i = 1$ is equivalent to $\kappa \geq 2/(n-1)$, which is guaranteed by the assumption that $\kappa \geq 2(n-3)/[(n-4)(n-1)]$. For $i \in \{3, 4, 5, 6\}$, the inequality (A.1) is equivalent to

$$\kappa \geq \max\left\{\frac{2(n-3)}{(n-1)(n-4)}, \frac{2(n-2)}{(n-1)(n-3)}, \frac{2}{n-2}, \frac{2}{n-1}\right\} = \frac{2(n-3)}{(n-1)(n-4)}.$$

Therefore, it holds that

$$g(\alpha, c) = g_2(\alpha, c) = 4\alpha(n-1)c.$$

whose maximum is attained at $c = [n(n-1)]^{-1}$ and

$$\max_{C,u^*} \mathbb{T}_\alpha\left(C, u^*\right) = g_2\left(\alpha, \frac{1}{n(n-1)}\right) = \frac{4\alpha}{n}.$$

The other two cases can be proved in the same way and the proof can be found in the online version [60]. □

B Proofs in Sect. 3

B.1 Proof of Theorem 3.2

Before proving the estimation of the complexity metric, we prove two properties of μ-incoherent vectors.

Lemma B.1 *Given any constant $\mu \in [1, n]$, suppose that u^* has incoherence μ and $\|u^*\|_1 = 1$. Then, the following properties hold:*

1. *u^* has at least n/μ nonzero components;*
2. *$|u_i^*| \leq \mu/n$ for all $i \in [n]$.*

The following lemma lower-bounds the perturbation of the weight matrix C.

Lemma B.2 *Suppose that the instance $\mathcal{MC}(C, u^*)$ satisfies the δ-RIP$_{2,2}$ condition and the weight matrix $\tilde{C} \in \mathbb{S}^{n^2-1}_{+,1}$ has N zero entries, where $\delta \in [0, 1)$ and $\mathcal{N} \subset [n] \times [n]$. Then, it holds that*

$$\|C - \tilde{C}\|_1 \geq 2 \sum_{(i,j)\in\mathcal{N}} C_{ij} \geq \frac{2(1-\delta)N}{(1+\delta)n^2 - 2\delta N},$$

where $\mathcal{N}$ is the set of indices of zero entries of $\tilde{C}$.

The proofs of Lemmas B.1 and B.2 are direct calculations and can be found in the online version [60]. Now, we prove the main theorem.

Proof of Theorem 3.2 Suppose that $\mathcal{MC}(\tilde{C}, \tilde{u}^*) \in \overline{\mathcal{D}}$ is the instance such that

$$\left[\mathbb{D}_\alpha\left(C, u^*\right)\right]^{-1} = \alpha\|C - \tilde{C}\|_1 + (1-\alpha)\|u^* - \tilde{u}^*\|_1.$$

In the following, we split the proof into two steps.

Step I. We first fix $\tilde{u}^*$ and consider the closest matrix $\tilde{C}$ to C such that $(\tilde{C}, \tilde{u}^*) \in \overline{\mathcal{D}}$. Let $k := |\mathcal{I}_1(\tilde{C}, \tilde{u}^*)|$. Without loss of generality, we assume that

$$\mathcal{I}_1(\tilde{C}, \tilde{u}^*) = \{1, \ldots, k\}, \quad \mathcal{I}_0(\tilde{C}, \tilde{u}^*) = \{k+1, \ldots, n\}.$$

We first consider the case when $k \geq 2$. If $\mathbb{G}_1(\tilde{C}, \tilde{u}^*)$ is disconnected, at least $2(k-1)$ entries of $\tilde{C}$ are 0. If $\mathbb{G}_1(\tilde{C}, \tilde{u}^*)$ are bipartite, at least $k^2/2 \geq 2(k-1)$ entries of $\tilde{C}$ are 0. If $\mathcal{I}_{00}(\tilde{C}, \tilde{u}^*)$ is non-empty, at least $2k$ entries of $\tilde{C}$ are 0. Otherwise if $k = 1$, at least one entry of $\tilde{C}$ should be 0 to make $\mathbb{G}_1(\tilde{C}, \tilde{u}^*)$ bipartite. In summary, at least $N(k)$ entries of $\tilde{C}$ are 0, where

$$N(k) := \max\{2(k-1), 1\}.$$

Using the results in Lemma B.2, the distance between C and $\tilde{C}$ is at least

$$\|C - \tilde{C}\|_1 \geq \frac{2(1-\delta)N(k)}{(1+\delta)n^2 - 2\delta N(k)}. \tag{B.1}$$

We note that the distance is monotonously increasing as a function of k.

Step II. Now, we consider the optimal choice of $\tilde{u}^*$ based on the lower bound in (B.1). Let

$$\ell := |\mathcal{I}_1\left(C, u^*\right)|, \quad k := |\mathcal{I}_1(\tilde{C}, \tilde{u}^*)|.$$

Since the distance between C and $\tilde{C}$ is a monotonously increasing function of k, the minimum distance between (C, u^*) and $(\tilde{C}, \tilde{u}^*)$ cannot be attained by $k > \ell$. Therefore, we focus on the case when $k \leq \ell$. Without loss of generality, we assume that

$$|u_1^*| \geq |u_2^*| \geq \cdots \geq |u_\ell^*| > 0; \quad |u_i^*| = 0, \quad \forall i \geq \ell + 1.$$

Then, the distance between u^* and $\tilde{u}^*$ satisfies

$$\|u^* - \tilde{u}^*\|_1 \geq 2 \sum_{i=k+1}^{\ell} |u_i^*|. \tag{B.2}$$

Denote the distance between (C, u^*) and $(\tilde{C}, \tilde{u}^*)$ by

$$d_\alpha := \alpha\|C - \tilde{C}\|_1 + (1 - \alpha)\|u^* - \tilde{u}^*\|_1.$$

Step II-1. We first consider the case when $\mu \leq 2n/3$. Combining inequalities (B.1) and (B.2), we obtain a lower bound on d_α:

$$d_\alpha \geq \min_{k \in [\ell]} \left[\frac{2\alpha(1-\delta)N(k)}{n^2(1+\delta) - 2\delta N(k)} + 2(1-\alpha) \sum_{i=k+1}^{\ell} |u_i^*| \right].$$

For every $k \in [\ell]$, the term inside the above minimization can be lower-bounded by

$$\begin{aligned} &\frac{2\alpha(1-\delta)N(k)}{n^2(1+\delta) - 2\delta N(k)} + 2(1-\alpha) \sum_{i=k+1}^{\ell} |u_i^*| \\ &\qquad \geq \frac{2\alpha(1-\delta) \cdot 2(k-1)}{n^2(1+\delta)} + 2(1-\alpha) \sum_{i=k+1}^{\ell} |u_i^*| \\ &\qquad = \frac{4\alpha(1-\delta)}{n^2(1+\delta)} \cdot (k-1) + 2(1-\alpha) \sum_{i=k+1}^{\ell} |u_i^*|. \end{aligned}$$

The minimum of the right-hand side over $k \in [\ell]$ can be solved in closed form and is equal to

$$\sum_{i=2}^{\ell} \min\left\{ \frac{4\alpha(1-\delta)}{n^2(1+\delta)}, 2(1-\alpha)|u_i^*| \right\}.$$

Using the second property in Lemma B.1, we have

$$\min\left\{\frac{4\alpha(1-\delta)}{n^2(1+\delta)}, 2(1-\alpha)|u_i^*|\right\} \geq \min\left\{\frac{4\alpha(1-\delta)}{n^2(1+\delta)} \cdot \frac{n|u_i^*|}{\mu}, 2(1-\alpha)|u_i^*|\right\}$$
$$= \min\left\{\frac{4\alpha(1-\delta)}{\mu n(1+\delta)}, 2(1-\alpha)\right\} \cdot |u_i^*|.$$

Taking the summation over $k \in \{2, \ldots, \ell\}$, we can conclude that

$$\begin{aligned} d_\alpha &\geq \sum_{k=2}^{\ell} \min\left\{\tfrac{4\alpha(1-\delta)}{\mu n(1+\delta)}, 2(1-\alpha)\right\} \cdot |u_i^*| \\ &= \min\left\{\tfrac{4\alpha(1-\delta)}{\mu n(1+\delta)}, 2(1-\alpha)\right\} \cdot \sum_{k=2}^{\ell} |u_i^*|. \end{aligned} \tag{B.3}$$

Using the second property in Lemma B.1 and $\|u^*\|_1 = 1$, it follows that

$$\sum_{k=2}^{\ell} |u_i^*| \geq 1 - \frac{\mu}{n}.$$

Substituting back into inequality (B.3), we have

$$d_\alpha \geq \min\left\{\frac{4\alpha(1-\delta)}{\mu n(1+\delta)}, 2(1-\alpha)\right\} \cdot \left(1 - \frac{\mu}{n}\right).$$

Step II-2. Next, we consider the case when $\mu \geq 2n/3$. By Theorem 3.1, the distance is at least

$$d_\alpha \geq \frac{2\alpha(1-\delta)}{n^2(1+\delta) - 2\delta} \geq \frac{2\alpha(1-\delta)}{(3/2)\mu \cdot n(1+\delta)} \geq \min\left\{\frac{4\alpha(1-\delta)}{\mu n(1+\delta)}, 2(1-\alpha)\right\} \cdot \frac{1}{3},$$

where the second inequality is due to the assumption that $\mu \geq 2n/3$.

By combining *Steps II-1* and *II-2*, the distance is lower-bounded by

$$\begin{aligned} d_\alpha &\geq \min\left\{\frac{4\alpha(1-\delta)}{\mu n(1+\delta)}, 2(1-\alpha)\right\} \times \max\left\{1 - \frac{\mu}{n}, \frac{1}{3}\right\} \\ &= \min\left\{\frac{4\alpha(1-\delta)}{n(1+\delta)}, 2(1-\alpha)\mu\right\} \times \max\left\{\frac{1}{\mu} - \frac{1}{n}, \frac{1}{3\mu}\right\} \end{aligned}$$

The proof is completed by using the relation between d_α and $\mathbb{T}_\alpha(C, u^*)$. □

B.2 Reduction of problem (3.7)

Before discussing the properties of problem instances in Sect. 3.3, we prove that the SSCPs of the instance $\mathcal{MC}(C^\epsilon, u^*)$ are closely related to those of the m-dimensional problem

$$\min_{x\in\mathbb{R}^m} \quad \sum_{i\in[m]} \left(x_i^2 - 1\right)^2 + \epsilon \sum_{i,j\in[m],i\neq j} (x_i x_j - 1)^2. \tag{B.4}$$

Lemma B.3 *If problem (B.4) has no SSCPs, then the instance $\mathcal{MC}(C^\epsilon, u^*)$ has no SSCPs. In addition, given a number $N \in \mathbb{N}$, suppose that problem (B.4) has N SSCPs with nonzero components at which the objective function has a positive definite Hessian matrix. Then, the instance $\mathcal{MC}(C^\epsilon, u^*)$ has at least N spurious local minima.*

The proof of Lemma B.3 is a direct calculation and can be found in the online version [60].

B.3 Proof of Theorem 3.5

To simplify the notations in the following proofs, we denote the gradient and the Hessian matrix of the objective function of problem (B.4) by

$$\begin{aligned} g_i(x;\epsilon) &:= 4\left[x_i^3 - x_i + \epsilon \sum_{j\neq i} x_j(x_i x_j - 1)\right], \quad \forall i \in [m]; \\ H_{ii}(x;\epsilon) &:= 4\left[3x_i^2 - 1 + \epsilon \sum_{j\neq i} x_j^2\right], \quad \forall i \in [m]; \\ H_{ij}(x;\epsilon) &:= 4\epsilon(2x_i x_j - 1), \quad \forall i,j \in [m] \quad \text{s.t.} \quad i \neq j. \end{aligned}$$

The following theorem guarantees that the instance $\mathcal{MC}(C^\epsilon, u^*)$ does not have spurious local minima when $\epsilon \geq O(m^{-1})$.

Theorem B.1 *If $\epsilon > 18/m$, the instance $\mathcal{MC}(C^\epsilon, u^*)$ does not have SSCPs, namely, all second-order critical points are global minima associated with the ground truth solution M^*.*

The proof of Theorem B.1 can be found in the online version [60]. Then, we consider the regime of ϵ where the instance $\mathcal{MC}(C^\epsilon, u^*)$ has spurious solutions. The following theorem studies the case when m is an even number.

Theorem B.2 *Suppose that m is an even number. If $\epsilon < 1/(m+1)$, then the instance $\mathcal{MC}(C^\epsilon, u^*)$ has at least $2^{m/2}$ spurious local minima.*

Proof By Lemma B.3, we only need to show that problem (B.4) has at least $\binom{m}{m/2}$ SSCPs whose associated Hessian matrices are positive definite and whose components are nonzero. We consider a point $x^0 \in \mathbb{R}^m$ such that

$$\left(x_i^0\right)^2 = \frac{1-\epsilon}{1+(m-1)\epsilon} > 0, \quad \forall i \in [m]; \quad \sum_{i\in[m]} x_i^0 = 0.$$

The above equations have a solution since m is an even number. By a direct calculation, we can verify that the gradient $g(x^0; \epsilon)$ is equal to 0. We only need to show that the Hessian matrix $H(x^0; \epsilon)$ is positive definite, namely

$$c^T H(x^0; \epsilon)c > 0, \quad \forall c \in \mathbb{R}^m \backslash \{0\}.$$

The above condition is equivalent to

$$\left[(3 + (m-3)\epsilon)\left(x_1^0\right)^2 - 1 + \epsilon\right] \sum_{i\in[m]} c_i^2 - \epsilon \left(\sum_{i\in[m]} c_i\right)^2 + 2\epsilon \left(x_1^0\right)^2 \left(\sum_{i\in[m]} \operatorname{sign}\left(x_1^0\right) c_i\right)^2 > 0, \quad \forall c \in \mathbb{R}^n \backslash \{0\}.$$

Under the normalization constraint $\|c\|_2 = 1$, the Cauchy inequality implies that the minimum of the left-hand side is attained by

$$c_1 = \cdots = c_m = 1/\sqrt{m}.$$

Therefore, the Hessian is positive definite if and only if

$$(3 + (m-3)\epsilon)\left(x_1^0\right)^2 - 1 + \epsilon > m\epsilon.$$

By substituting $\left(x_1^0\right)^2 = (1-\epsilon)/[1 + (m-1)\epsilon]$, the above condition is equivalent to

$$2 - (m+4)\epsilon - (m-2)(m+1)\epsilon^2 > 0.$$

Using the condition that $(m+1)\epsilon < 1$, we obtain that

$$2 - (m+4)\epsilon - (m-2)(m+1)\epsilon^2 > 1 - 3\epsilon - (m-2)\epsilon = 1 - (m+1)\epsilon > 0,$$

where the first inequality is from the fact that $m \geq 2$, which follows from the assumption that $m > 0$ is an even number.

To estimate the number of SSCPs, we observe that $m/2$ components of x^0 have a positive sign and the other $m/2$ components have a negative sign. Hence, there are at least

$$\binom{m}{m/2}$$

spurious SSCPs. The estimate on the combinatorial number is in light of the inequality $\binom{n}{k} \geq (n/k)^k$. □

The estimation of the odd number case is similar and we present the result in the following theorem.

Theorem B.3 *Suppose that m is an odd number. If* $\epsilon < 1/[13(m+1)]$*, then the instance* $\mathcal{MC}(C^\epsilon, u^*)$ *has at least* $[2m/(m+1)]^{(m+1)/2}$ *spurious local minima.*

The proof of Theorem B.3 is similar to that of Theorem B.2 and can be found in the online version [60]. By combining Theorems B.1–B.3, we complete the proof of Theorem 3.5.

B.4 Proof of Theorem 3.6

The proof of Theorem 3.6 relies on the following lemma, which calculates the complexity metric of the instance $\mathcal{MC}(C^\epsilon, u^*)$. The proof of Lemma B.4 is similar to that of Theorem 2.4.

Lemma B.4 *Suppose that* $n \geq m \geq 5$, $\alpha \in [0, 1]$ *and* $\epsilon \in [0, 1]$. *The complexity metric* $\mathbb{D}_\alpha(C^\epsilon, u^*)$ *has the closed form*

$$[\mathbb{D}_\alpha(C^\epsilon, u^*)]^{-1} = \min\left\{ \frac{2\alpha}{Z_\epsilon} + \frac{2(1-\alpha)(m-1)}{m}, \frac{4\alpha\epsilon}{Z_\epsilon} + \frac{2(1-\alpha)(m-2)}{m}, \frac{4\alpha(m-1)\epsilon}{Z_\epsilon} \right\}.$$

Moreover, $\mathbb{D}_\alpha(C^\epsilon, u^*)$ *is strictly decreasing in* ϵ *on* $[0, 1/2]$.

The proof of Lemma B.4 can be found in the online version [60]. Combining Theorem 3.5 and Lemma B.4, we are able to estimate the range of the complexity metric.

Proof of Theorem 3.6 By defining constants $\delta := 1/26$ and $\Delta := 18$, Theorem 3.5 implies that

1. If $\epsilon < \delta/m$, the instance $\mathcal{MC}(C^\epsilon, u^*)$ has spurious local minima;
2. If $\epsilon > \Delta/m$, the instance $\mathcal{MC}(C^\epsilon, u^*)$ has no spurious local minima.

Then, we study two different cases.

Case I. We first consider the case when $m\epsilon$ is large. Since $\epsilon < \Delta/m \leq 1/2$, the threshold is located in the regime where $\mathbb{D}_\alpha(C^\epsilon, u^*)$ is strictly decreasing. Hence, it suffices to show that

$$\left[\frac{2\alpha\Delta}{n^2} + \min\left\{4\alpha\Delta \cdot \frac{m}{n^2}, 2(1-\alpha)\right\}\right]^{-1}$$

is a lower bound on $\mathbb{D}_\alpha(C^\epsilon, u^*)$ when $\epsilon = \Delta/m$. By Lemma B.4, it holds that

$$\left[\mathbb{D}_\alpha(C^\epsilon, u^*)\right]^{-1}$$

$$
\begin{aligned}
&=\min\left\{\frac{2\alpha}{Z_\epsilon}+\frac{2(1-\alpha)(m-1)}{m},\frac{4\alpha\epsilon}{Z_\epsilon}+\frac{2(1-\alpha)(m-2)}{m},\frac{4\alpha(m-1)\epsilon}{Z_\epsilon}\right\}\\
&\leq\min\left\{\frac{4\alpha\epsilon}{Z_\epsilon}+\frac{2(1-\alpha)(m-2)}{m},\frac{4\alpha(m-1)\epsilon}{Z_\epsilon}\right\}\\
&=\frac{4\alpha\epsilon}{Z_\epsilon}+(m-2)\min\left\{\frac{4\alpha\epsilon}{Z_\epsilon},\frac{2(1-\alpha)}{m}\right\}\leq\frac{4\alpha\epsilon}{Z_\epsilon}+m\min\left\{\frac{4\alpha\epsilon}{Z_\epsilon},\frac{2(1-\alpha)}{m}\right\}.
\end{aligned}
$$

Since the graph $\mathbb{G}$ does not contain any independence set with $m+1$ nodes, Turán's theorem [3] implies that the graph $\mathbb{G}$ has at least $n^2/(2m)$ edges, namely,

$$|\mathbb{E}|\geq n^2/(2m).$$

We note that the above bound is asymptotically tight and is attained by the Turán graph. Hence, we obtain that

$$Z_\epsilon=2|\mathbb{E}|+n+m(m-1)\epsilon\geq 2|\mathbb{E}|\geq n^2/m.$$

By substituting into the estimate of $\mathbb{D}_\alpha(C^\epsilon,u^*)$, it follows that

$$
\begin{aligned}
\left[\mathbb{D}_\alpha(C^\epsilon,u^*)\right]^{-1}&\leq\frac{4\alpha\epsilon\cdot m}{n^2}+m\min\left\{\frac{4\alpha\epsilon\cdot m}{n^2},\frac{2(1-\alpha)}{m}\right\}\\
&=\frac{2\alpha\Delta}{n^2}+\min\left\{4\alpha\Delta\cdot\frac{m}{n^2},2(1-\alpha)\right\}.
\end{aligned}
$$

Case II. Next, we consider the case when ϵm is small. Similar to *Case I*, it suffices to show that

$$\frac{18}{17}\max\left\{\frac{n^2}{4\alpha\delta},\frac{1}{2(1-\alpha)}\right\}$$

is an upper bound for $\mathbb{D}_\alpha(C^\epsilon,u^*)$ when $\epsilon=\delta/m$. Since $\delta<1/2$, we have

$$2\alpha/Z_\epsilon>4\alpha\epsilon/Z_\epsilon.$$

By Lemma B.4, it holds that

$$
\begin{aligned}
&\left[\mathbb{D}_\alpha(C^\epsilon,u^*)\right]^{-1}\\
&=\min\left\{\frac{2\alpha}{Z_\epsilon}+\frac{2(1-\alpha)(m-1)}{m},\frac{4\alpha\epsilon}{Z_\epsilon}+\frac{2(1-\alpha)(m-2)}{m},\frac{4\alpha(m-1)\epsilon}{Z_\epsilon}\right\}\\
&=\min\left\{\frac{4\alpha\epsilon}{Z_\epsilon}+\frac{2(1-\alpha)(m-2)}{m},\frac{4\alpha(m-1)\epsilon}{Z_\epsilon}\right\}\\
&=\frac{4\alpha\epsilon}{Z_\epsilon}+(m-2)\min\left\{\frac{4\alpha\epsilon}{Z_\epsilon},\frac{2(1-\alpha)}{m}\right\}\geq\frac{17}{18}\min\left\{\frac{4\alpha\epsilon m}{Z_\epsilon},2(1-\alpha)\right\},
\end{aligned}
$$

where the last inequality is from $m \geq 36$. Since $\epsilon \leq 1$, the definition of Z_ϵ implies that $Z_\epsilon \leq n^2$. By substituting into the estimate of $\mathbb{D}_\alpha(C^\epsilon, u^*)$, it follows that

$$\left[\mathbb{D}_\alpha(C^\epsilon, u^*)\right]^{-1} \geq \frac{17}{18}\min\left\{\frac{4\alpha\epsilon m}{n^2}, 2(1-\alpha)\right\} = \frac{17}{18}\min\left\{\frac{4\alpha\delta}{n^2}, 2(1-\alpha)\right\}.$$

By combining *Cases I* and *II*, we complete the proof. □

C Proofs in Sect. 4

C.1 Proof of Theorem 4.3

The proof of Theorem 4.3 directly follows from the next two lemmas.

Lemma C.1 *Suppose that* $(C, u^*) \in \mathcal{SD}$ *and that* u^0 *is a global solution to* $\mathcal{MC}(C, u^*)$. *Then, for all* $k \in [n_1]$, *it holds that* $u_i^0 u_j^0 = u_i^* u_j^*$ *for all* $i, j \in \mathcal{I}_{1k}$. *In addition,* $u_i^0 = 0$ *for all* $i \in \mathcal{I}_0(C, u^*)$.

Proof Denote $M^* := u^*(u^*)^T$. We first consider nodes in $\mathcal{G}_{1k}$ for some $k \in [n_1]$. Since the subgraph is not bipartite, there exists a cycle with an odd length $2\ell + 1$, which we denote as

$$\{i_1, \ldots, i_{2\ell+1}\}.$$

Then, we have

$$\begin{aligned}\left(u_{i_1}^0\right)^2 &= \prod_{s=1}^{2\ell+1}\left(u_{i_s}^0 u_{i_{s+1}}^0\right)^{(-1)^{s-1}} = \prod_{s=1}^{2\ell+1}\left(M_{i_s i_{s+1}}^*\right)^{(-1)^{s-1}} \\ &= \prod_{s=1}^{2\ell+1}\left(u_{i_s}^* u_{i_{s+1}}^*\right)^{(-1)^{s-1}} = \left(u_{i_1}^*\right)^2,\end{aligned}$$

which implies that the conclusion holds for $i = j = i_1$. Using the connectivity of $\mathbb{G}_{1k}(C, u^*)$, we know

$$u_i^0 u_j^0 = u_i^* u_k^*, \quad \forall i, j \in \mathcal{I}_{1k}\left(C, u^*\right).$$

Then, we consider nodes in $\mathcal{I}_0(C, u^*)$. Since $\mathcal{I}_{00}(C, u^*)$ is empty, for every node $i \in \mathcal{I}_0(C, u^*)$, there exists another node $j \in \mathcal{I}_1(C, u^*)$ such that $C_{ij} > 0$. Hence, we have

$$u_i^0 = M_{ij}^*/u_j^0 = 0.$$

This completes the proof. □

The following lemma provides a necessary and sufficient condition for instances with a positive definite Hessian matrix at global solutions, which is stronger than what Theorem 4.3 requires.

Lemma C.2 *Suppose that $u^0 \in \mathbb{R}^n$ is a global minimizer of the instance $\mathcal{MC}(C, u^*)$ such that the conditions in Lemma C.1 hold. Then, the Hessian matrix is positive definite at u^0 if and only if*

1. *$\mathbb{G}_{1i}(C, u^*)$ is not bipartite for all $i \in [n_1]$;*
2. *$\mathcal{I}_{00}(C, u^*) = \emptyset$.*

Proof We prove the positive definiteness of the Hessian matrix for the sufficiency part since it is the part used in this manuscript. The necessity part is proved in the online version [60].

Sufficiency. Now, we consider the sufficiency part, namely, we prove that the Hessian matrix is positive definite under the two conditions stated in the theorem. Suppose that there exists a nonzero vector $q \in \mathbb{R}^n$ such that

$$\left[\nabla^2 g\left(u^0; C, u^*\right)\right](q, q) = 0.$$

Then, after straightforward calculations, we arrive at

$$\begin{aligned} u_i^0 q_j + u_j^0 q_i &= 0, \quad \forall i, j \quad \text{s.t.} \quad C_{ij} > 0,\ i \neq j; \\ \left[2\left(u_i^0\right)^2 - \left(u_i^*\right)^2\right] q_i^2 = \left(u_i^0 q_i\right)^2 &= 0, \quad \forall i \quad \text{s.t.} \quad C_{ii} > 0. \end{aligned}$$

The two conditions can be written compactly as

$$u_i^0 q_j + u_j^0 q_i = 0, \quad \forall i, j \quad \text{s.t.} \quad C_{ij} > 0. \tag{C.1}$$

Consider the index set $\mathcal{I}_{1k}(C, u^*)$ for some $k \in [n_1]$. The equality (C.1) implies that

$$q_i/u_i^0 + q_j/u_j^0 = 0, \quad \forall i, j \in \mathcal{I}_{1k}\left(C, u^*\right). \tag{C.2}$$

Since the graph $\mathbb{G}_{1k}(C, u^*)$ is not bipartite, there exists a cycle with an odd length $2\ell + 1$, which we denote as

$$\{i_1, i_2, \ldots, i_{2\ell+1}\}.$$

Denoting $i_{2\ell+2} := i_1$, we can calculate that

$$2\frac{q_{i_1}}{u_{i_1}^0} = \sum_{s=1}^{2\ell+1} (-1)^{s-1} \left(\frac{q_{i_s}}{u_{i_s}^0} + \frac{q_{i_{s+1}}}{u_{i_{s+1}}^0}\right) = 0,$$

which leads to $q_{i_1} = 0$. Using the connectivity of $\mathcal{G}_{1k}$ and the relation (C.2), it follows that

$$q_i = 0, \quad \forall i \in \mathcal{I}_{1k}\left(C, u^*\right).$$

Moreover, the same conclusion holds for all $k \in [n_1]$ and, thus, we conclude that

$$q_i = 0, \quad \forall i \in \mathcal{I}_1 \left(C, u^*\right).$$

Since $\mathcal{I}_{00}(C, u^*) = \emptyset$, for every node $i \in \mathcal{I}_0(C, u^*)$, there exists another node $j \in \mathcal{I}_1(C, u^*)$ such that $C_{ij} > 0$. Considering the relation (C.2), we obtain that

$$q_j = -u_j^0 q_i / u_i^0 = 0.$$

In summary, we have proved that $q_i = 0$ for all $i \in [n]$, which contradicts the assumption that $q \neq 0$. Hence, the Hessian matrix at u^0 is positive definite. □

C.2 Application of the implicit function theorem

Using the positive-definiteness of the Hessian matrix, we are able to apply the implicit function theorem to certify the existence of spurious local minima.

Lemma C.3 *Suppose that $\alpha \in [0, 1]$ and consider a pair $(C, u^*) \in \mathcal{SD}$. Then, there exists a small constant $\delta(C, u^*) > 0$ such that for every instance $\mathcal{MC}(\tilde{C}, \tilde{u}^*)$ satisfying*

$$\alpha \|\tilde{C} - C\|_1 + (1 - \alpha)\|\tilde{u}^* - u^*\|_1 < \delta \left(C, u^*\right),$$

the instance $\mathcal{MC}(\tilde{C}, \tilde{u}^*)$ has spurious local minima.

Proof By Theorem 4.3, there exists a global solution u^0 to the instance $\mathcal{MC}(C, u^*)$ such that

$$u^0(u^0)^T \neq u^*(u^*)^T, \quad \nabla^2 g(u^0; C, u^*) \succ 0.$$

Consider the system of equations:

$$\nabla g(u; C, u^*) = 0.$$

Since the Jacobi matrix of $\nabla g(u; C, u^*)$ with respect to u is the Hessian matrix $\nabla^2 g(u; C, u^*)$ and (u^0, C, u^*) is a solution, the implicit function theorem guarantees that there exists a small constant $\delta(C, u^*) > 0$ such that in the neighborhood

$$\mathcal{N} := \left\{(\tilde{C}, \tilde{u}^*) \mid \alpha\|\tilde{C} - C\|_1 + (1 - \alpha)\|\tilde{u}^* - u^*\|_1 < \delta \left(C, u^*\right)\right\},$$

there exists a function $u(\tilde{C}, \tilde{u}^*) : \mathcal{N} \mapsto \mathbb{R}^n$ such that

1. $u(C, u^*) = u^0$;
2. $u(\cdot, \cdot)$ is a continuous function in $\mathcal{N}$;
3. $\nabla g[u(\tilde{C}, \tilde{u}^*); \tilde{C}, \tilde{u}^*] = 0$.

Using the continuity of the Hessian matrix and $u(\cdot, \cdot)$, we can choose $\delta(C, u^*)$ to be small enough such that

$$u(\tilde{C}, \tilde{u}^*)[u(\tilde{C}, \tilde{u}^*)]^T \neq \tilde{u}^*(\tilde{u}^*)^T, \quad \nabla^2 g\left[u(\tilde{C}, \tilde{u}^*); \tilde{C}, \tilde{u}^*\right] \succ 0, \quad \forall (\tilde{C}, \tilde{u}^*) \in \mathcal{N}.$$

Therefore, the point $u(\tilde{C}, \tilde{u}^*)$ is a spurious local minimum of the instance $\mathcal{MC}(\tilde{C}, \tilde{u}^*)$. □

References

1. Agarwal, A., Anandkumar, A., Jain, P., Netrapalli, P.: Learning sparsely used overcomplete dictionaries via alternating minimization. SIAM J. Optim. **26**(4), 2775–2799 (2016)
2. Ahn, K., Suarez, F.: Riemannian perspective on matrix factorization. arXiv preprint arXiv:2102.00937 (2021)
3. Aigner, M.: Turán's graph theorem. Am. Math. Mon. **102**(9), 808–816 (1995)
4. Ajayi, T., Mildebrath, D., Kyrillidis, A., Ubaru, S., Kollias, G., Bouchard, K.: Provably convergent acceleration in factored gradient descent with applications in matrix sensing. arXiv preprint arXiv:1806.00534 (2018)
5. Allen-Zhu, Z., Li, Y.: Neon2: Finding local minima via first-order oracles. Adv. Neural Inform. Process. Syst. **31** (2018)
6. Bi, Y., Lavaei, J.: On the absence of spurious local minima in nonlinear low-rank matrix recovery problems. In: International Conference on Artificial Intelligence and Statistics, pp. 379–387. PMLR (2021)
7. Bi, Y., Zhang, H., Lavaei, J.: Local and global linear convergence of general low-rank matrix recovery problems. In: Proceedings of 36th AAAI Conference on Artificial Intelligence (AAAI), pp. 1–9. Vancouver, Canada (2022)
8. Burer, S., Monteiro, R.D.: A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. Math. Program. **95**(2), 329–357 (2003)
9. Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? J. ACM (JACM) **58**(3), 1–37 (2011)
10. Candes, E.J., Li, X., Soltanolkotabi, M.: Phase retrieval via wirtinger flow: theory and algorithms. IEEE Trans. Inform. Theory **61**(4), 1985–2007 (2015)
11. Candes, E.J., Plan, Y.: Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. IEEE Trans. Inform. Theory **57**(4), 2342–2359 (2011)
12. Candès, E.J., Recht, B.: Exact matrix completion via convex optimization. Found. Comput. Math. **9**(6), 717–772 (2009)
13. Candès, E.J., Tao, T.: The power of convex relaxation: near-optimal matrix completion. IEEE Trans. Inform. Theory **56**(5), 2053–2080 (2010)
14. Cartis, C., Gould, N.I., Toint, P.L.: Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results. Math. Program. **127**(2), 245–295 (2011)
15. Charisopoulos, V., Chen, Y., Davis, D., Díaz, M., Ding, L., Drusvyatskiy, D.: Low-rank matrix recovery with composite optimization: good conditioning and rapid convergence. Found. Comput. Math. **21**(6), 1505–1593 (2021)
16. Chen, J., Li, X.: Model-free nonconvex matrix completion: local minima analysis and applications in memory-efficient kernel PCA. J. Mach. Learn. Res. **20**(142), 1–39 (2019)
17. Chen, J., Liu, D., Li, X.: Nonconvex rectangular matrix completion via gradient descent without $\ell_{2,\infty}$ regularization. IEEE Trans. Inform. Theory **66**(9), 5806–5841 (2020)
18. Chen, Y., Chi, Y.: Harnessing structures in big data via guaranteed low-rank matrix estimation: recent theory and fast algorithms via convex and nonconvex optimization. IEEE Signal Process. Mag. **35**(4), 14–31 (2018)
19. Chen, Y., Chi, Y., Fan, J., Ma, C.: Gradient descent with random initialization: fast global convergence for nonconvex phase retrieval. Math. Program. **176**(1), 5–37 (2019)

20. Chen, Y., Chi, Y., Fan, J., Ma, C., Yan, Y.: Noisy matrix completion: understanding statistical guarantees for convex relaxation via nonconvex optimization. SIAM J. Optim. **30**(4), 3098–3121 (2020)
21. Chen, Y., Fan, J., Ma, C., Yan, Y.: Bridging convex and nonconvex optimization in robust PCA: noise, outliers and missing data. Ann. Stat. **49**(5), 2948–2971 (2021)
22. Chi, Y., Lu, Y.M., Chen, Y.: Nonconvex optimization meets low-rank matrix factorization: an overview. IEEE Trans. Signal Process. **67**(20), 5239–5269 (2019)
23. Chou, H.H., Gieshoff, C., Maly, J., Rauhut, H.: Gradient descent for deep matrix factorization: dynamics and implicit bias towards low rank. arXiv preprint arXiv:2011.13772 (2020)
24. Fattahi, S., Sojoudi, S.: Exact guarantees on the absence of spurious local minima for non-negative rank-1 robust principal component analysis. J. Mach. Learn. Res. (2020)
25. Ge, R., Jin, C., Zheng, Y.: No spurious local minima in nonconvex low rank problems: a unified geometric analysis. In: International Conference on Machine Learning, pp. 1233–1242. PMLR (2017)
26. Ge, R., Lee, J.D., Ma, T.: Matrix completion has no spurious local minimum. Adv. Neural Inform. Process. Syst. 2981–2989 (2016)
27. Hardt, M.: Understanding alternating minimization for matrix completion. In: 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, pp. 651–660. IEEE (2014)
28. Hardt, M., Wootters, M.: Fast matrix completion without the condition number. In: Conference on Learning Theory, pp. 638–678. PMLR (2014)
29. Hou, T.Y., Li, Z., Zhang, Z.: Fast global convergence for low-rank matrix recovery via riemannian gradient descent with random initialization. arXiv preprint arXiv:2012.15467 (2020)
30. Jain, P., Meka, R., Dhillon, I.: Guaranteed rank minimization via singular value projection. Adv. Neural Inform. Process. Syst. **23** (2010)
31. Jain, P., Netrapalli, P., Sanghavi, S.: Low-rank matrix completion using alternating minimization. In: Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, pp. 665–674. (2013)
32. Jin, C., Ge, R., Netrapalli, P., Kakade, S.M., Jordan, M.I.: How to escape saddle points efficiently. In: International Conference on Machine Learning, pp. 1724–1732. PMLR (2017)
33. Jin, C., Netrapalli, P., Jordan, M.I.: Accelerated gradient descent escapes saddle points faster than gradient descent. In: Conference On Learning Theory, pp. 1042–1085. PMLR (2018)
34. Lee, J.D., Simchowitz, M., Jordan, M.I., Recht, B.: Gradient descent only converges to minimizers. In: Conference on Learning Theory, pp. 1246–1257. PMLR (2016)
35. Levin, E., Kileel, J., Boumal, N.: The effect of smooth parametrizations on nonconvex optimization landscapes. arXiv preprint arXiv:2207.03512 (2022)
36. Li, X., Zhu, Z., Man-Cho So, A., Vidal, R.: Nonconvex robust low-rank matrix recovery. SIAM J. Optim. **30**(1), 660–686 (2020)
37. Li, Y., Ma, T., Zhang, H.: Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In: Conference On Learning Theory, pp. 2–47. PMLR (2018)
38. Luo, Y., Li, X., Zhang, A.R.: Nonconvex factorization and manifold formulations are almost equivalent in low-rank matrix optimization. arXiv preprint arXiv:2108.01772 (2021)
39. Ma, J., Fattahi, S.: Sign-rip: a robust restricted isometry property for low-rank matrix recovery. arXiv preprint arXiv:2102.02969 (2021)
40. Ma, Z., Bi, Y., Lavaei, J., Sojoudi, S.: Sharp restricted isometry property bounds for low-rank matrix recovery problems with corrupted measurements. arXiv preprint arXiv:2105.08232 (2021)
41. Netrapalli, P., Jain, P., Sanghavi, S.: Phase retrieval using alternating minimization. Adv. Neural Inform. Process. Syst. **26** (2013)
42. Netrapalli, P., Un, N., Sanghavi, S., Anandkumar, A., Jain, P.: Non-convex robust PCA. Adv. Neural Inform. Process. Syst. **27** (2014)
43. Recht, B., Fazel, M., Parrilo, P.A.: Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. SIAM Rev. **52**(3), 471–501 (2010)
44. Renegar, J.: Linear programming, complexity theory and elementary functional analysis. Math. Program. **70**(1), 279–351 (1995)
45. Renegar, J.: Condition numbers, the barrier method, and the conjugate-gradient method. SIAM J. Optim. **6**(4), 879–912 (1996)
46. Stöger, D., Soltanolkotabi, M.: Small random initialization is akin to spectral learning: optimization and generalization guarantees for overparameterized low-rank matrix reconstruction. Adv. Neural Inform. Process. Syst. **34** (2021)
47. Sun, J., Qu, Q., Wright, J.: Complete dictionary recovery over the sphere I: overview and the geometric picture. IEEE Trans. Inform. Theory **63**(2), 853–884 (2016)

48. Sun, J., Qu, Q., Wright, J.: A geometric analysis of phase retrieval. Found. Comput. Math. **18**(5), 1131–1198 (2018)
49. Sun, R., Luo, Z.Q.: Guaranteed matrix completion via non-convex factorization. IEEE Trans. Inform. Theory **62**(11), 6535–6579 (2016)
50. Tong, T., Ma, C., Chi, Y.: Accelerating ill-conditioned low-rank matrix estimation via scaled gradient descent. J. Mach. Learn. Res. **22**(150), 1–63 (2021)
51. Tong, T., Ma, C., Chi, Y.: Low-rank matrix recovery with scaled subgradient methods: fast and robust convergence without the condition number. IEEE Trans. Signal Process. **69**, 2396–2409 (2021)
52. Tong, T., Ma, C., Prater-Bennette, A., Tripp, E., Chi, Y.: Scaling and scalability: provable nonconvex low-rank tensor estimation from incomplete measurements. arXiv preprint arXiv:2104.14526 (2021)
53. Tu, S., Boczar, R., Simchowitz, M., Soltanolkotabi, M., Recht, B.: Low-rank solutions of linear matrix equations via procrustes flow. In: International Conference on Machine Learning, pp. 964–973. PMLR (2016)
54. Wei, K., Cai, J.F., Chan, T.F., Leung, S.: Guarantees of riemannian optimization for low rank matrix recovery. SIAM J. Matrix Anal. Appl. **37**(3), 1198–1222 (2016)
55. Wei, K., Cai, J.F., Chan, T.F., Leung, S.: Guarantees of riemannian optimization for low rank matrix completion. Inverse Probl. Ima. **14**(2) (2020)
56. Yalcin, B., Zhang, H., Lavaei, J., Sojoudi, S.: Factorization approach for low-complexity matrix completion problems: exponential number of spurious solutions and failure of gradient methods. In: International Conference on Artificial Intelligence and Statistics, pp. 1–9. PMLR (2022)
57. Yi, X., Park, D., Chen, Y., Caramanis, C.: Fast algorithms for robust pca via gradient descent. Adv. Neural Inform. Process. Syst. **29** (2016)
58. Yurtsever, A., Tropp, J.A., Fercoq, O., Udell, M., Cevher, V.: Scalable semidefinite programming. SIAM J. Math. Data Sci. **3**(1), 171–200 (2021)
59. Zhang, H., Bi, Y., Lavaei, J.: General low-rank matrix optimization: geometric analysis and sharper bounds. Adv. Neural Inform. Process. Syst. **34** (2021)
60. Zhang, H., Yalcin, B., Lavaei, J., Sojoudi, S.: A new complexity metric for nonconvex rank-one generalized matrix completion. arXiv preprint arXiv:2204.02364 (2022)
61. Zhang, J., Fattahi, S., Zhang, R.: Preconditioned gradient descent for over-parameterized nonconvex matrix factorization. Adv. Neural Inform. Process. Syst. **34** (2021)
62. Zhang, R.Y., Sojoudi, S., Lavaei, J.: Sharp restricted isometry bounds for the inexistence of spurious local minima in nonconvex matrix recovery. J. Mach. Learn. Res. **20**(114), 1–34 (2019)
63. Zheng, Q., Lafferty, J.: A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. In: Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1, pp. 109–117 (2015)
64. Zhu, Z., Li, Q., Tang, G., Wakin, M.B.: Global optimality in low-rank matrix optimization. IEEE Trans. Signal Process. **66**(13), 3614–3628 (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.