
Factorization Approach for Low-complexity Matrix Completion Problems: Exponential Number of Spurious Solutions and Failure of Gradient Methods

Baturalp Yalcin
UC Berkeley

Haixiang Zhang
UC Berkeley

Javad Lavaei
UC Berkeley

Somayeh Sojoudi
UC Berkeley

Abstract

Burer-Monteiro (B-M) factorization approach can efficiently solve low-rank matrix optimization problems under the Restricted Isometry Property (RIP) condition. It is natural to ask whether B-M factorization-based methods can succeed on any low-rank matrix optimization problems with a low information-theoretic complexity, i.e., polynomial-time solvable problems that have a unique solution. We provide negative answer to this question. We investigate the landscape of B-M factorized polynomial-time solvable matrix completion (MC) problems, which are the most popular subclass of low-rank matrix optimization problems without the RIP condition. We construct an instance of polynomial-time solvable MC problems with exponentially many spurious local minima, which leads to the failure of most gradient-based methods. We define a new complexity metric that measures the solvability of low-rank matrix optimization problems based on B-M factorization approach. In addition, we show that more measurements can deteriorate the landscape, which further reveals the unfavorable behavior of B-M factorization.

1 INTRODUCTION

The low-rank matrix optimization problem aims to recover a low-rank ground truth matrix $\mathbf{M}^*$ through

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

some measurements modeled as $\mathcal{A}(\mathbf{M}^*)$, where the measurement operator $\mathcal{A}$ is a function from $\mathbb{R}^{n \times n}$ to $\mathbb{R}^d$. The operator $\mathcal{A}$ can be either linear as in the linear matrix sensing problem and the matrix completion problem (Candès & Recht, 2009; Recht et al., 2010), or nonlinear as in the one-bit matrix sensing problem (Davenport et al., 2014) and the phase retrieval problem (Shechtman et al., 2015). There are two variants of the problem, known as symmetric and asymmetric problems. The first one assumes that $\mathbf{M}^*$ is a positive semi-definite (PSD) matrix, whereas the second one makes no such assumption and allows $\mathbf{M}^*$ to be non-symmetric or sign indefinite. Since the asymmetric problem can be equivalently transformed into a symmetric problem (Zhang et al., 2021), we focus on the latter one.

There are in general two different approaches to overcome the non-convex low-rank constraint. The first approach is to design a convex penalty function that prefers low-rank matrices and then optimize the penalty function under the measurement constraint (Candès & Recht, 2009; Recht et al., 2010; Candès & Tao, 2010). However, this approach works in the matrix space $\mathbb{R}^{n \times n}$ and has a high computational complexity. The other widely accepted technique is the Burer-Monteiro (B-M) factorization approach (Burer & Monteiro, 2003), which converts the original problem into an unconstrained one by replacing the original PSD matrix variable $\mathbf{M} \in \mathbb{R}^{n \times n}$ with the product of a low-dimensional variable $\mathbf{X} \in \mathbb{R}^{n \times r}$ and its transpose. The optimization problem based on the B-M factorization approach can be written as

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times r}} g \left[\mathcal{A}(\mathbf{X}\mathbf{X}^T) - \mathcal{A}(\mathbf{M}^*) \right],$$

where $g(\cdot)$ is a loss function that penalizes the mismatch between $\mathbf{X}\mathbf{X}^T$ and $\mathbf{M}^*$. Using the B-M factorization, the objective function is generally non-convex even if the loss function $g(\cdot)$ is convex. Nonetheless, it has been proved that under certain

strong conditions, such as the Restricted Isometry Property (RIP) condition, saddle-escaping methods can converge to the ground truth solution with a random initialization (Zhang et al., 2021; Bi et al., 2021) and first-order methods with spectral initialization converge locally (Tu et al., 2016; Bhojanapalli et al., 2016); see Chi et al. (2019) for an overview.

The B-M factorization is practical due to its low computational complexity. When it is successful, the problem is considered to be tractable. Then, it is natural to ask whether optimization methods based on the B-M factorization approach can succeed on general low-rank matrix optimization problems with a low information-theoretic complexity (i.e., problems that have a unique global solution and can be solved in polynomial time), especially when the RIP condition does not hold. In this work, we focus on a common class of problems without the RIP condition, namely the matrix completion (MC) problem. For the MC problem, the measurement operator $\mathcal{A}_\Omega : \mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$ is given by

$$\mathcal{A}_\Omega(\mathbf{M})_{ij} := \begin{cases} \mathbf{M}_{ij} & \text{if } (i, j) \in \Omega \\ 0 & \text{otherwise,} \end{cases}$$

where Ω is the set of indices of observed entries. We denote the measurement operator as $\mathbf{M}_\Omega := \mathcal{A}_\Omega(\mathbf{M})$ for simplicity. An instance of the MC problem, denoted as $\mathcal{P}_{\mathbf{M}^*, \Omega, n, r}$, can be formulated as

$$\begin{aligned} \text{find } & \mathbf{M} \in \mathbb{R}^{n \times n} & (\mathcal{P}_{\mathbf{M}^*, \Omega, n, r}) \\ \text{s.t. } & \text{rank}(\mathbf{M}) \leq r, \quad \mathbf{M} \succeq 0, \quad \mathbf{M}_\Omega = \mathbf{M}_\Omega^*. \end{aligned}$$

If $\mathbf{M}^*$ is the only solution of this problem, we will say that $\mathcal{P}_{\mathbf{M}^*, \Omega, n, r}$ has a unique solution. Using the B-M factorization approach, the MC problem can be solved via the optimization problem

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times r}} f(\mathbf{X}), \quad (1)$$

where $f(\mathbf{X}) := g[(\mathbf{X}\mathbf{X}^T - \mathbf{M}^*)_\Omega]$. For example, if the ℓ_2 -loss function is used, the problem becomes

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times r}} \left\| (\mathbf{X}\mathbf{X}^T - \mathbf{M}^*)_\Omega \right\|_F^2. \quad (2)$$

Contributions. We provide a negative answer to the preceding question by constructing MC problem instances for which the *optimization complexity* of local search methods using the B-M factorization does not align with the *information-theoretic complexity* of the underlying MC problem instance. The information-theoretic complexity refers to the

minimum number of operations that the best possible algorithm takes to find the ground truth matrix, while the optimization complexity refers to the minimum number of operations that a given optimization method takes to find the ground truth matrix. In general, the optimization complexity of local search methods depends on the properties of spurious solutions of the optimization problem, e.g., the number, the sharpness and the regions of attraction of spurious solutions. The optimization complexity predicts the performance of an algorithm and provides a hint on which algorithm to use for a given problem. Therefore, the results in this work imply that the popular B-M factorization approach is not able to capture the benign properties of the low-rank problem when the RIP condition does not hold. We summarize our contributions as follows:

- i) Given natural numbers n and r with $n \geq 2r$, we construct a class of MC problem instances with a deterministic measurement set in $\mathcal{L}(\mathcal{G}, n, r)$, whose ground matrix $\mathbf{M}^* \in \mathbb{S}_+^n$ has rank r . For every instance in this class, there exists a unique global solution and the solution can be found in polynomial time via graph-theoretic algorithms.
- ii) Next, we show the existence of an instance in $\mathcal{L}(\mathcal{G}, n, r)$ whose B-M factorization formulation (2) has at least $\mathcal{O}(2^{n-2r})$ equivalence classes of spurious¹ local solutions. This claim holds for general loss functions under a weak assumption.
- iii) For the rank-1 case, we prove that most gradient-based methods with a random initialization converge to a spurious local minimum with probability at least $1 - \mathcal{O}(2^{-n/2})$. Numerical studies verify that the failure of the gradient-based methods happens for general rank.
- iv) We present an instance that has no spurious solution under the B-M factorization formulation (2), but introducing additional observations of the ground truth matrix leads to at least exponentially many spurious solutions. This example further reveals the unfavorable behavior of the B-M factorization approach on general low-rank matrix optimization problems.
- v) Based on these results, we define a new complexity metric that potentially captures the optimization complexity of optimization methods based on the B-M factorization.

Related Work. The low-rank optimization problem has been well studied under the RIP condition (Recht et al., 2010). Several recent works (Zhang

¹A solution is called spurious if it is a local minimum, but has a larger objective value than the optimal.

et al., 2019; Bi & Lavei, 2021; Zhang et al., 2021) showed that the non-convex formulation has no spurious local minima with a small RIP parameter. To understand how conservative the RIP condition is, we consider a class of polynomial-time solvable problems without the RIP condition and study the behavior of optimization methods on this class. Specifically, we consider the polynomial-time solvable MC problems with deterministic measurement pattern. Most existing literature on the MC problem is based on the assumption that the measurement set is randomly constructed and the global solution is coherent (Candès & Recht, 2009; Candès & Tao, 2010; Ge et al., 2016, 2017; Ma et al., 2019; Chen et al., 2020). In comparison, there is a small range of works that have focused on the deterministic MC problem (Bhojanapalli & Jain, 2014; Király et al., 2015; Pimentel-Alarcón et al., 2016; Li et al., 2016). Furthermore, efficient graph-theoretic algorithms utilizing the special structures of a deterministic measurement set can be designed (Ma et al., 2018). Existing works on a deterministic measurement set case have focused on the completability problem and the convex relaxation approach, while the B-M factorization approach has not been analyzed. Moreover, there are several existing works that also provided negative results on the low-rank matrix optimization problem. The counterexamples in Candès & Tao (2010); Bhojanapalli & Jain (2014) have non-unique global solutions, which make the recovery of the ground truth matrix impossible. The counterexamples in Waldspurger & Waters (2020) have a unique global solution but the objective function must be a linear function. We refer the reader to Chi et al. (2019) for a review of the low-rank matrix optimization problem. Our work is the first one in the literature that studies the optimization complexity in the case when the information-theoretic complexity is low.

Notations. The set $[n]$ represents the set of integers from 1 to n . We use lower case bold letters $\mathbf{x}$ to represent vectors and capital bold letters $\mathbf{X}$ to represent matrices. $\|\mathbf{X}\|$ and $\|\mathbf{X}\|_F$ are the 2-norm and the Frobenius norm of the matrix $\mathbf{X}$, respectively. Let $\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr}(\mathbf{A}^T \mathbf{B})$ be the inner product between matrices. The notations $\mathbf{X} \succeq 0$ and $\mathbf{X} \succ 0$ mean that the matrix $\mathbf{X}$ is PSD and positive definite, respectively. The set of $n \times n$ PSD matrices is denoted as $\mathbb{S}_+^n$. For a function $f : \mathbb{R}^{m \times n} \mapsto \mathbb{R}$, we denote the gradient and the Hessian as $\nabla f(\cdot)$ and $\nabla^2 f(\cdot)$, respectively. The Hessian is a four-dimensional tensor with $[\nabla^2 f(\mathbf{X})]_{i,j,k,l} = \frac{\partial^2 f(\mathbf{X})}{\partial \mathbf{x}_{i,j} \partial \mathbf{x}_{k,l}}$ for all $i, k \in [m]$ and $j, l \in [n]$. The quadratic form of the Hessian in the direction $\Delta \in \mathbb{R}^{m \times n}$ is defined

as $\Delta : \nabla^2 f(\mathbf{X}) : \Delta = \sum_{i,j,k,l} [\nabla^2 f(\mathbf{X})]_{i,j,k,l} \Delta_{ij} \Delta_{kl}$. We use $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ to denote the ceiling and flooring functions, respectively. The cardinality of a set $\mathcal{S}$ is shown as $|\mathcal{S}|$.

2 EXPONENTIAL NUMBER OF SPURIOUS LOCAL MINIMA

In this section, we show that MC problem instances with a low information-theoretic complexity may have exponentially many spurious local minima if the B-M factorization is employed. We first construct a class of MC problem instances with a low information-theoretic complexity and then identify the problematic instances.

2.1 Low-complexity Class of MC Problems

Suppose that $r \geq 1$ and $n \geq 2r$ are two given integers. We construct a class of MC problem instances whose ground truth matrix $\mathbf{M}^* \in \mathbb{S}_+^n$ is rank- r . For every instance in this class, the global solution is unique and can be found in polynomial time in terms of n and r . Let $m := \lfloor n/r \rfloor \geq 2$. We divide the first mr rows and the first mr columns of the matrix $\mathbf{M}^*$ into $m \times m$ block matrices, where each block has dimension $r \times r$. For every $i, j \in [m]$, we denote the block matrix at position (i, j) as $\mathbf{M}_{i,j}^*$. We now define the block measurement patterns induced by a given graph.

Definition 1 (Induced measurement set). *Let $\mathcal{G} = (\mathcal{G}_1, \mathcal{G}_2) = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ be a pair of undirected graphs with the node set $\mathcal{V} = [m]$ and the disjoint edge sets $\mathcal{E}_1, \mathcal{E}_2 \subset [m] \times [m]$. The induced measurement set $\Omega(\mathcal{G})$ is defined as: if $(i, j) \in \mathcal{E}_1$, then the entire block $\mathbf{M}_{i,j}^*$ are observed; if $(i, j) \in \mathcal{E}_2$, then all nondiagonal entries of the block $\mathbf{M}_{i,j}^*$ are observed; otherwise, none of the entries of the block is observed. In addition, the last $n - mr$ rows and the last $n - mr$ columns of the matrix $\mathbf{M}^*$ are fully observed. We refer to the graph $\mathcal{G}$ as the block sparsity graph.*

The following definition introduces a low-complexity class of MC problem instances.

Definition 2 (Low-complexity class of MC problems). *Define $\mathcal{L}(\mathcal{G}, n, r)$ to be the class of low-complexity MC problems $\mathcal{P}_{\mathbf{M}^*, \Omega, n, r}$ with the following properties:*

- i) *The ground truth $\mathbf{M}^* \in \mathbb{S}_+^n$ is rank- r .*
- ii) *Matrix $\mathbf{M}_{i,j}^* \in \mathbb{R}^{r \times r}$ is rank- r for all $i, j \in [m]$.*
- iii) *The measurement set $\Omega = \Omega(\mathcal{G})$ is induced by $\mathcal{G} = (\mathcal{G}_1, \mathcal{G}_2)$, where $\mathcal{G}_1$ is connected and non-bipartite.*

The next proposition states that every MC problem instance in $\mathcal{L}(\mathcal{G}, n, r)$ is polynomial-time solvable. A graph-theoretical algorithm is used to find the solution within polynomial time by only using the partially observed rank- r matrix.

Proposition 1. *The ground truth $\mathbf{M}^*$ is the unique solution of this problem and can be found in $\mathcal{O}(n^2/r^2 + nr^2)$ time for every problem in $\mathcal{L}(\mathcal{G}, n, r)$.*

2.2 Intuition for Rank-1 Case with ℓ_2 -loss Function

We start with the case when the rank r is equal to 1 and the loss function $g(\cdot)$ is the ℓ_2 -loss. We study two instances in the class $\mathcal{L}(\mathcal{G}, n, 1)$ with $\mathcal{O}(n)$ and $\mathcal{O}(n^2)$ observations, respectively. The B-M formulation (2) of both instances contains exponentially many spurious local minima. Since the decomposition variable $\mathbf{X}$ is a column vector in the rank-1 case, we write it as $\mathbf{x}$.

Example 1. *We first provide an instance with $\mathcal{O}(n)$ observations. Note that the number of blocks, namely m , is equal to n in the rank-1 case. Let the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ be chosen as $\mathcal{V} := [n]$ and*

$$\mathcal{E}_1 := \{(i, j) \mid i, j \in [n], |i - j| \leq 1\}, \quad \mathcal{E}_2 := \emptyset.$$

The measurement set is the induced set $\Omega := \Omega(\mathcal{G})$. Namely, we observe the diagonal, sub-diagonal and super-diagonal entries of the ground truth matrix. One can verify that the subgraph $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ is connected and non-bipartite. Now, we construct a specific ground truth matrix. We define the vector $\mathbf{x}^ \in \mathbb{R}^n$ by*

$$\mathbf{x}_{2k+1}^* := 1, \quad \forall k \in \llbracket n/2 \rrbracket, \quad \mathbf{x}_{2k}^* := 0, \quad \forall k \in \llbracket n/2 \rrbracket,$$

and let $\mathbf{M}^ := \mathbf{x}^*(\mathbf{x}^*)^T$. For the B-M factorization formulation (2), the set of global minima is given by*

$$\mathcal{X}^* := \left\{ \mathbf{x} \in \mathbb{R}^n \mid \begin{aligned} &\mathbf{x}_{2k+1}^2 = 1, \quad \forall k \in \llbracket n/2 \rrbracket, \\ &\mathbf{x}_{2k} = 0, \quad \forall k \in \llbracket n/2 \rrbracket \end{aligned} \right\},$$

which has cardinality $2^{\lceil \frac{n}{2} \rceil}$. For every global solution $\hat{\mathbf{x}} \in \mathcal{X}^$ and every $\Delta \in \mathbb{R}^n \setminus \{0\}$, the Hessian satisfies*

$$\begin{aligned} \Delta : \nabla^2 f(\hat{\mathbf{x}}) : \Delta &= 2 \left\| (\hat{\mathbf{x}} \Delta^T + \Delta \hat{\mathbf{x}}^T)_{\Omega} \right\|_F^2 \\ &= 8 \|\Delta\|^2 - \mathbb{1}[n \text{ is even}] 4\Delta_n^2 > 0, \end{aligned}$$

where $\mathbb{1}[\cdot]$ is the indicator function. Therefore, the Hessian is positive definite at every global minimum. Then, we perturb the ground truth solution $\mathbf{M}^$ to*

$$\mathbf{M}^*(\epsilon) := \mathbf{x}^*(\epsilon) [\mathbf{x}^*(\epsilon)]^T = (\mathbf{x}^* + \epsilon)(\mathbf{x}^* + \epsilon)^T,$$

where $\mathbf{x}^(\epsilon) := \mathbf{x}^* + \epsilon$ and $\epsilon \in \mathbb{R}^n$ is a small perturbation. We denote the associated problem (2) as*

$$\min_{\mathbf{x} \in \mathbb{R}^n} \tilde{f}(\mathbf{x}; \epsilon), \quad (3)$$

where $\tilde{f}(\mathbf{x}; \epsilon) := \|(\mathbf{x}\mathbf{x}^T - \mathbf{M}(\epsilon)^)_{\Omega}\|_F^2$. For a generic perturbation ϵ , all components of ϵ are nonzero and problem $\mathcal{P}_{\mathbf{M}^*(\epsilon), \Omega, n, 1}$ belongs to the class $\mathcal{L}(\mathcal{G}, n, 1)$. This implies that the global solution of problem (3) is unique up to a sign flip.*

We analyze the relation between the local minima of the original problem and those of the perturbed problem. Consider the equation $\nabla_{\mathbf{x}} \tilde{f}(\mathbf{x}; \epsilon) = 0$ near an unperturbed global minimum $\hat{\mathbf{x}} \in \mathcal{X}^$. Since $(\hat{\mathbf{x}}; 0)$ is a solution to the gradient equation and the Jacobian matrix with respect to $\mathbf{x}$ is equal to the positive definite Hessian $\nabla^2 f(\hat{\mathbf{x}})$, the Implicit Function Theorem (IFT) states that there exists a unique solution $\hat{\mathbf{x}}(\epsilon)$ in a neighbourhood of $\hat{\mathbf{x}}$ for all values of ϵ with a small norm. In addition, the continuity of Hessian implies that $\nabla_{\mathbf{x}\mathbf{x}} \tilde{f}(\hat{\mathbf{x}}^*(\epsilon); \epsilon) \succ 0$. Thus, $\hat{\mathbf{x}}(\epsilon)$ is a local minimum of the perturbed problem (3). As a result, we have proved the existence of a local minimum for the perturbed problem corresponding to each of the $2^{\lceil n/2 \rceil}$ global minima of the unperturbed problem. Hence, the problem (3) has at least $2^{\lceil n/2 \rceil}$ local minima, while only two of them are global minima. In summary, an instance in $\mathcal{L}(\mathcal{G}, n, 1)$ has exponentially many spurious local solutions.*

Example 2. *Next, we construct an MC problem instance with exponentially many spurious local minima and $\mathcal{O}(n^2)$ observations. We choose the same ground truth matrix $\mathbf{M}^*(\epsilon)$ as in the last example, but assume that the measurement set Ω is induced by the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ with $\mathcal{V} := [n]$, $\mathcal{E}_2 = \emptyset$, and $\mathcal{E}_1 := \{(i, i), (i, 2k), (2k, i) \mid \forall i \in [n], k \in \llbracket n/2 \rrbracket\}$. Since the subgraph $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ is connected and non-bipartite, the perturbed problem $\mathcal{P}_{\mathbf{M}^*(\epsilon), \Omega, n, 1}$ belongs to the class $\mathcal{L}(\mathcal{G}, n, 1)$. Moreover, one can verify that the set of global minima of this problem is still $\mathcal{X}^*$ and the Hessian at every global solution is positive definite. By the same argument, IFT implies that problem (3) has at least $2^{\lceil n/2 \rceil} - 2$ spurious local minima for a generic and small perturbation ϵ .*

Note that the instances analyzed in this section and those in the remainder of this paper, satisfy the incoherence condition (Candès & Recht, 2009) with the parameter $\mu = \mathcal{O}(1)$.

2.3 Rank-1 Case with General Measurement Sets

In this subsection, we estimate the largest lower bound on the number of spurious local minima for

the given parameters n and r . We address the problem by first finding a lower bound on the number of spurious local minima given a general measurement set Ω , and then maximizing the lower bound over Ω . The following theorem utilizes the topology of $\mathcal{G}$ to quantify a lower bound on the number of spurious solutions for the measurement set $\Omega(\mathcal{G})$.

Theorem 1. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \emptyset)$ such that $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ is connected and non-bipartite with n vertices. Assume that there exists a maximal independent set² $\mathcal{S}(\mathcal{G}_1)$ of $\mathcal{G}_1$ such that every vertex in the set has a self-loop. There exists an instance in $\mathcal{L}(\mathcal{G}, n, 1)$ for which the problem (2) has at least $2^{|\mathcal{S}(\mathcal{G}_1)|} - 2$ spurious local minima.*

In both Examples 1 and 2, a maximal independent set is $\mathcal{S} = \{2k + 1 \mid \forall k \in [\lceil n/2 \rceil]\}$. Hence, Theorem 1 implies that there are $2^{\lceil n/2 \rceil} - 2$ spurious local minima, which is consistent with our analysis. Since a maximal independent set of a connected and non-bipartite graph can have up to $n - 1$ vertices, the number of spurious local minima can be as large as $2^{n-1} - 2$.

Corollary 1. *There exist a graph $\mathcal{G}$ and an instance in $\mathcal{L}(\mathcal{G}, n, 1)$ such that problem (2) has $2^{n-1} - 2$ spurious solutions. In addition, there exist a graph $\mathcal{G}$ and an instance in $\mathcal{L}(\mathcal{G}, n, 1)$ with $|\Omega| = n^2 - 2$ such that the problem (2) has spurious solutions.*

Corollary 1 implies that the B-M factorization may not be an efficient approach to the MC problem, since it has a spurious solution even in the highly ideal case when almost all entries of the matrix are measured. Since these results work for almost all graphs in the rank-1 case, extensions to the random graph case can be achieved via estimating the size of the maximum independent set of a random graph. Generally, the proof of Theorem 1 implies that, as a necessary condition for not having a spurious solution in formulation (2), the elements of $\mathbf{x}^*$ associated with the nodes outside of the maximal independent set $\mathcal{S}$ should not be much smaller than those associated with the nodes in $\mathcal{S}$.

Corollary 2. *Under the setting of Theorem 1, there exists a function $h_{\mathcal{S}}(\cdot) : (0, \infty) \mapsto (0, \infty)$ such that $\mathcal{P}_{\mathbf{x}^*(\mathbf{x}^*)^T, \Omega(\mathcal{G}), n, 1}$ has at least $2^{|\mathcal{S}|} - 2$ spurious local minima in formulation (2) for every generic $\mathbf{x}^* \in \mathbb{R}^n$ satisfying*

$$\|\mathbf{x}_{\mathcal{S}^c}^*\| \leq h_{\mathcal{S}}(\min_{i \in \mathcal{S}} |x_i^*|) \cdot \|\mathbf{x}_{\mathcal{S}}^*\|,$$

²For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the set $\mathcal{S} \subset \mathcal{V}$ 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 and is not a strict subset of another independent set.

where $\mathcal{S}^c := [n] \setminus \mathcal{S}$ and $\mathbf{x}_{\mathcal{S}^c} := (x_i : i \notin \mathcal{S})$.

Because the maximal independent set of a graph $\mathcal{G}$ is not necessarily unique, the set of functions $h_{\mathcal{S}, M}(\cdot)$ over all maximal independent sets $\mathcal{S}$ designates a necessary condition for the nonexistence of spurious local minima given a measurement set $\Omega(\mathcal{G})$.

2.4 Extension to General Rank- r Case

We generalize the results to the case when the ground truth matrix has a general rank. Eisenberg-Nagy et al. (2013) showed that the rank- r MC problem is $\mathcal{NP}$ -hard in the worst case for every $r \geq 2$. However, we focus on instances in the low-complexity class $\mathcal{L}(\mathcal{G}, n, r)$ and show that there are instances in this class whose B-M factorization formulation (2) has a highly undesirable landscape. We cannot simply extend the proof of the rank-1 case to the rank- r case since there exist an infinite number of matrices $\mathbf{X}^*$ such that $\mathbf{M}^* = \mathbf{X}^*(\mathbf{X}^*)^T$ when $r > 1$. The global optimality of a solution $\mathbf{X}^*$ is not lost under any orthogonal transformation. This implies that the Hessian at the global solutions of problem (2) cannot be positive definite, which fails the applicability of IFT. Therefore, we consider the quotient manifold $\mathcal{M}_{n,r}^r(\mathbb{R})/O_r$, where O_r is the lie group of $r \times r$ orthogonal matrices and $\mathcal{M}_{n,r}^r(\mathbb{R})/O_r$ is the open-submanifold of rank- r matrices of size $n \times r$. To simplify the analysis, we instead consider the following lower-diagonal subspace

$$\mathcal{W}^{n \times r} := \{\mathbf{X} \in \mathbb{R}^{n \times r} \mid \mathbf{X}_{ij} = 0, \forall i \in [n], j \in [r] \text{ s.t. } i < j\}.$$

We define an embedding of the manifold $\mathcal{M}_{n,r}^r(\mathbb{R})/O_r$ into $\mathcal{W}^{n \times r}$ and composite it with the quotient map.

Definition 3 (Restriction map). *Given a matrix $\mathbf{X} \in \mathbb{R}^{n \times r}$, we define the embedding $\phi^{emb}([\mathbf{X}]) := \mathbf{R} \in \mathcal{W}^{n \times r}$, where $\mathbf{X} = \mathbf{R}\mathbf{Q}$ is the RQ decomposition with $\mathbf{Q}$ being an orthogonal matrix and $\mathbf{R}$ having non-negative diagonal elements. The restriction map is defined as $\phi(\mathbf{X}) := \phi^{emb}([\mathbf{X}])$.*

The properties of the RQ decomposition ensure that $\phi^{emb}(\cdot)$ is a bijection in a small neighborhood of each matrix $\mathbf{X}$ whose first r rows are linearly independent. Consider the restricted version of problem (2):

$$\min_{\mathbf{X} \in \mathcal{W}^{n \times r}} f(\mathbf{X}), \quad (4)$$

Results in Section 2.3 can be extended to the problem (4), and be translated back to the problem (2).

Lemma 1. *Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ with $m = \lfloor n/r \rfloor$ vertices for which the subgraph $\mathcal{G}_1 =$*

$(\mathcal{V}, \mathcal{E}_1)$ is connected and non-bipartite. Assume that there exists a maximal independent set $\mathcal{S}(\mathcal{G}_1)$ of $\mathcal{G}_1$ whose vertices each have a self-loop. If the induced subgraph³ $\mathcal{G}_2[\mathcal{S}]$ is connected, then there exists an instance in $\mathcal{L}(\mathcal{G}, n, r)$ for which the problem (4) has at least $2^{r|\mathcal{S}(\mathcal{G}_1)|} - 2^r$ spurious local minima. In addition, the first r rows of each local minimum are linearly independent.

The Hessian at each local minimum of the unperturbed problem is positive definite along the tangent space of $\mathcal{W}^{n \times r}$, where the off-diagonal observations of $\Omega(\mathcal{G})$ play a key role. If the first r rows of a local minimum $\hat{\mathbf{X}}$ are linearly independent, the diagonal elements of $\hat{\mathbf{X}}$ are nonzero. Therefore, by flipping the signs of columns, we can find an equivalent local minimum $\tilde{\mathbf{X}}$ with positive diagonal elements, i.e., $\tilde{\mathbf{X}}$ lies in the range of $\phi^{emb}(\cdot)$. By symmetry, the total number of such local minima is $2^{r(|\mathcal{S}(\mathcal{G}_1)|-1)} - 1$. Since the restriction map $\phi^{emb}(\cdot)$ is a bijection in a neighborhood of $\tilde{\mathbf{X}}$, the equivalence class $[\tilde{\mathbf{X}}] \in \mathcal{M}_{n,r}^r(\mathbb{R})/O_r$ is a local minimum of problem (2) on the quotient manifold and thus $\tilde{\mathbf{X}}$ is a local minimum of problem (2). The above argument leads to Theorem 2.

Theorem 2. Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ satisfying the conditions of Lemma 1. There exists an instance in $\mathcal{L}(\mathcal{G}, n, r)$ for which problem (2) has at least $2^{r(|\mathcal{S}(\mathcal{G}_1)|-1)} - 1$ equivalence classes of spurious solutions.

Finally, we give an estimate on the largest lower bound for the number of spurious local minima.

Corollary 3. There exists an instance in $\mathcal{L}(\mathcal{G}, n, r)$ for which the problem (2) has at least $2^{n-2r} - 1$ equivalence classes of spurious solutions. In addition, there exists an instance in $\mathcal{L}(\mathcal{G}, n, r)$ with $|\Omega| = n^2 - 2r$ for which the problem (2) has spurious solutions.

Low-rank matrix optimization problems are easier to solve as the rank- r approaches the dimension n . However, the above generalization demonstrates that the undesired behavior of the B-M factorization is still valid for arbitrary ranks and they never become easy to solve using this method.

2.5 General Loss Functions

In this part, we generalize the preceding results to the problem (1). To extend the constructions to a general loss function $g(\cdot)$, we require a few weak assumptions on the loss function $g(\cdot)$.

³See Harary (2018) for the definition.

Assumption 1. The following conditions hold for the function $g(\cdot)$:

- i) $g(\cdot)$ is twice continuously differentiable;
- ii) the matrix $\mathbf{0}_{n \times r}$ is the unique minimizer of $g(\cdot)$;
- iii) the Hessian of $g(\cdot)$ at $\mathbf{0}_{n \times r}$ is positive definite.

We can extend the results in Section 2.4 to the general loss function case under the above assumption.

Theorem 3. Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ satisfying the conditions of Lemma 1 and suppose that Assumption 1 holds. There exists an instance in $\mathcal{L}(\mathcal{G}, n, r)$ for which the problem (1) has at least $2^{r(|\mathcal{S}(\mathcal{G}_1)|-1)} - 1$ equivalence classes of spurious local minima.

We note that the ℓ_2 -loss function $g(\cdot) = \|\cdot\|_F^2$ satisfies the conditions in Assumption 1. As another example, regularizers are ubiquitously used in the low-rank matrix optimization literature (Ge et al., 2016, 2017; Fattahi & Sojoudi, 2020). As a corollary to Theorem 3, the regularized version of the problem (2) also suffers from the same issue. In this case, the loss function is equal to

$$g(\mathbf{X}) := \left\| (\mathbf{X}\mathbf{X}^T - \mathbf{M}^*)_{\Omega} \right\|_F^2 + Q(\mathbf{X}),$$

where $Q(\mathbf{X}) := \lambda (\|\mathbf{X}_i\| - \alpha)_+^4$ is from Ge et al. (2016), $(x)_+ := \max\{x, 0\}$ and $\alpha, \lambda > 0$ are constants. Since the regularizer does not change the landscape around global solutions with a large α , Assumption 1 is satisfied and Theorem 3 is applicable to the regularized problem.

3 MORE OBSERVATIONS LEAD TO SPURIOUS LOCAL MINIMA

In Section 2, we showed that the B-M factorization formulation (2) has an exponential number of spurious local minima on low-complexity MC problem instances. In this section, we exhibit another unfavorable behaviour of the B-M factorization. We identify an MC problem instance in $\mathcal{L}(\mathcal{G}, n, r)$ with some pattern Ω that has no spurious solution while adding observations to Ω leads to spurious solutions. Let m and r be natural numbers with $m \geq 2r$. We define $n := mr$ and let the graph be $\mathcal{G}_k := (\mathcal{V}, \mathcal{E}_k)$ where $k \in [m]$ is an arbitrary index and

$$\mathcal{V} := [m], \quad \mathcal{E}_k := \{(k, j), (j, k) \mid \forall j \in [m]\}.$$

In the measurement set Ω_k , we observe the blocks $\mathbf{M}_{i,j}$ for all $(i, j) \in \mathcal{E}_k$; see Figure 1. Ω_k contains only full block observations induced by $\mathcal{G}_k$. The next proposition states that if the ground truth matrix

$$\mathbf{M}_{\Omega_k} = \left[\begin{array}{c|c|c} \mathbf{0} & \mathbf{M}_{1,k} & \mathbf{0} \\ \hline \mathbf{M}_{k,1} & \ddots & \mathbf{M}_{k,m} \\ \hline \mathbf{0} & \mathbf{M}_{m,k} & \mathbf{0} \end{array} \right]$$

 Figure 1: Measurement operator $\mathcal{A}_{\Omega_k}$ on matrix $\mathbf{M}$.

is generic, every second-order critical point⁴ of the problem (2) is a global minimum.

Proposition 2. *Given an index $k \in [m]$, let the measurement set Ω be equal to Ω_k . Assume that the block $\mathbf{M}_{i,j}^*$ of the ground truth matrix $\mathbf{M}^*$ has rank r for all $i, j \in [m]$. Then, every SOCP of problem (2) is a global minimum.*

Next, we construct a graph $\tilde{\mathcal{G}}_k := (\mathcal{V}, \tilde{\mathcal{E}}_k, \mathcal{E}_2)$, where

$$\begin{aligned} \tilde{\mathcal{E}}_k &:= \mathcal{E}_k \cup \{(i, i) \mid \forall i \in [m]\}, \\ \mathcal{E}_2 &:= \{(i, j) \mid \forall i, j \in [m], i \neq j\}. \end{aligned}$$

Namely, we have included all self-loops and nondiagonal observations of each block in the new graph. A maximal independent set for the subgraph $\tilde{\mathcal{G}}_{k,1} := (\mathcal{V}, \tilde{\mathcal{E}}_k)$ is $\mathcal{S} := [m] \setminus \{k\}$. We define a new measurement set $\tilde{\Omega}_k := \Omega(\tilde{\mathcal{G}}_k)$. Since $\tilde{\mathcal{E}}_k$ is a superset of $\mathcal{E}_k$, the measurement set $\tilde{\Omega}_k$ is larger than Ω_k . Using Theorem 2, we obtain the following result.

Corollary 4. *Every instance of the MC problem with the measurement set $\tilde{\Omega}_k$ and full rank blocks $\mathbf{M}_{i,j}^*$ for all $i, j \in [m]$ belongs to $\mathcal{L}(\mathcal{G}, n, r)$. The formulation (2) of an instance of the problem has at least $2^{r(m-2)} - 1$ equivalence classes of spurious local minima, while all spurious solutions disappear when using the smaller set Ω_k .*

Results of Proposition 2 and Corollary 4 conclude that the landscape of the problem (2) deteriorates when the number of observations is increased. This phenomenon further reveals the unfavorable behavior of the B-M factorization on low-rank matrix optimization problems, even if the information-theoretic complexity is low.

4 MEASURE OF COMPLEXITY

In Section 2, we showed that if there is an MC problem with a non-unique completion, a slightly per-

⁴Second-order critical points are defined as those points that satisfy the first-order and the second-order necessary optimality conditions.

turbed problem will have exponentially many spurious local minima in the B-M factorization formulation (1). Hence, bifurcation behaviors appear around measurement matrices $\mathbf{M}_{\Omega}^*$ that are associated with multiple global solutions. For a given measurement operator $\mathcal{A}$, a measurement matrix $\mathcal{A}(\mathbf{M})$ that allows multiple global solutions designates an unacceptable region in the space of ground truth solutions. Based on that, we define a metric to capture the extent of the bifurcation behavior. The set of measurements that allow d non-unique solutions:

$$\begin{aligned} \mathcal{T}_{\mathcal{A}}^d &:= \left\{ \mathcal{A}(\mathbf{M}) : \exists \{\mathbf{X}_i\}_{i=1}^d \in \mathbb{R}^{n \times r} \right. \\ &\quad \text{s.t. } \mathbf{X}_i \mathbf{X}_i^T \neq \mathbf{X}_j \mathbf{X}_j^T, \forall i, j \in [d], i \neq j, \\ &\quad \left. \mathcal{A}(\mathbf{M}) = \mathcal{A}(\mathbf{X}_i \mathbf{X}_i^T), \forall i \in [d] \right\}. \end{aligned}$$

Then, we define a complexity metric below.

Definition 4 (Complexity metric). *The complexity metric for operator $\mathcal{A}$ and ground truth $\mathbf{M}^*$ is defined as $\tau_{\mathcal{A}, \mathbf{M}^*}^d = 1/\text{dist}(\mathcal{A}(\mathbf{M}^*), \mathcal{T}_{\mathcal{A}}^d)$ where*

$$\text{dist}(\mathcal{A}(\mathbf{M}^*), \mathcal{T}_{\mathcal{A}}^d) := \min_{\mathcal{A}(\mathbf{M}) \in \mathcal{T}_{\mathcal{A}}^d} \|\mathcal{A}(\mathbf{M}^*) - \mathcal{A}(\mathbf{M})\|_F.$$

Analysis in Section 2 implies if the complexity metric $\tau_{\mathcal{A}, \mathbf{M}^*}^d$ is large, there exist $\mathcal{O}(d)$ spurious solutions. Hence, the optimization complexity of algorithms based on the B-M factorization approach may be aligned with the corresponding information-theoretic complexity for instances with a small complexity metric $\tau_{\mathcal{A}, \mathbf{M}^*}^2$. For example, if the RIP condition is satisfied, the set $\mathcal{T}$ is empty and the complexity metric is always 0. Ge et al. (2017) show when the incoherence condition is small and the measurement graph is random, the B-M factorization is successful with high probability. Since a small incoherence condition requires balanced columns and rows while the non-unique completion occurs when the columns and rows of the matrix are less balanced, we have a small complexity metric for matrices with small incoherence parameters. This implies that our complexity metric supports the existing results on the success of the B-M factorization approach under incoherence conditions. Oppositely, the instances studied in Section 2 with spurious solutions all have very large complexity metrics. Consequently, the complexity metric is a possible measure of the optimization complexity for the MC problem with the B-M factorization: the optimization problem should be more difficult if the complexity metric is higher.

5 GRADIENT METHODS FAIL

We show that the exponential number of spurious local minima in preceding instances will make

most randomly initialized gradient-based methods fail with a high probability. The existence of spurious local minima does not necessarily imply the failure of gradient-based methods; see Ma et al. (2018); Chen et al. (2020). The analysis in this section is based on the gradient flow

$$\dot{\mathbf{X}}(t) = -\nabla_{\mathbf{X}} f(\mathbf{X}(t)), \quad \mathbf{X}(0) = \mathbf{X}_0. \quad (5)$$

It is known that the trajectories of gradient-based methods with a small enough step size are close to those of the gradient flow. We can view a variety of gradient-based methods as ordinary differential equation (ODE) solvers applied to the gradient flow (5). Then, the convergence of the discrete trajectories when the step size goes to 0 can be guaranteed by the consistency and the stability of the ODE solver. Scieur et al. (2017) proved that the consistency and stability conditions are satisfied by several commonly used gradient-based methods, such as the gradient descent, the proximal point and the accelerated gradient descent methods. Although Scieur et al. (2017) considered minimizing a strongly convex function, the consistency and stability conditions only depend on the ODE solver and the Lipschitz continuity of the underlying gradient flow. We need the following assumption on the loss function to characterize the global landscape.

Assumption 2. *The loss function $g(\cdot)$ satisfies the sparse (δ, r) -RIP condition in the Ω -norm for some constant $\delta \in [0, 1)$ and integer $r \geq 1$. Namely, the inequality*

$$(1 - \delta) \|\mathbf{N}_\Omega\|_F^2 \leq \mathbf{N} : \nabla^2 g(\mathbf{M}_\Omega) : \mathbf{N} \leq (1 + \delta) \|\mathbf{N}_\Omega\|_F^2$$

holds for all matrices $\mathbf{M}$ and $\mathbf{N}$ with rank at most $2r$.

The sparse RIP condition is remarkably different from the conservative RIP condition. For example, the ℓ_2 -loss function satisfies the sparse $(0, r)$ -RIP condition for every r , while the RIP condition does not hold if Ω is not a complete graph. Under the above assumption, we show that for the unperturbed example constructed in Section 2, the gradient flow will converge to each global minimum with equal probability in the rank-1 case. The main difficulty is to show that all saddle points of the objective function $f(\mathbf{X})$ are strict and therefore their region of attractions (ROAs) have measure zero (Lee et al., 2016).

Lemma 2. *Suppose that Assumption 2 holds for $r = 1$ and $\delta = \mathcal{O}(1/n)$, where $n \geq 3$ is the size of the ground truth matrix. There exists an MC problem instance such that the following statements hold for the problem (1):*

- *there are $2^{\lceil n/2 \rceil}$ equivalent global minima;*
- *if the gradient flow (5) is initialized with an absolutely continuous radial probability distribution, it converges to each global minimum with the equal probability $2^{-\lceil n/2 \rceil}$,*

where a probability distribution with respect to the Lebesgue measure is called radial if its density function at point x only depends on $\|x\|$.

Examples of absolutely continuous radial probability distributions include zero-mean Gaussian distributions and uniform distributions over a ball centered at the origin. Note that the ℓ_2 -loss function satisfies the assumption of Lemma 2. Next, we show that with a sufficiently small perturbation to the previous instance of the problem, the ROA of each local minimum will not shrink significantly. Therefore, the gradient flow will converge to each global minimum or spurious solution with approximately the same probability.

Theorem 4. *Under the setting of Lemma 2, consider an absolutely continuous radial probability distribution. There exists an instance in $\mathcal{L}(\mathcal{G}, n, 1)$ for which the problem (1) satisfies the following properties:*

- *the global minima are unique up to a sign flip;*
- *if the gradient flow (5) is initialized with the given distribution, it converges to a global minimum associated with the ground truth solution with probability at most $\mathcal{O}(2^{-\lceil n/2 \rceil})$.*

The results of Theorem 4 imply that, in the rank-1 case, most gradient-based methods with a small enough step size and a suitable random initialization will converge to a spurious solution with an overwhelming probability. The proof works for the general rank case if it can be shown that there is no degenerate saddle points for the above-mentioned unperturbed instances of the problem.

Remark. *We remark that the trajectories of stochastic gradient descent (SGD) methods cannot be approximated by those of the gradient flow. Hence, our analysis cannot automatically imply the failure of SGD. However, the proof of Lemma 2 can be adopted to conclude that SGD methods with a random initialization will converge to each global minimum with equal probability in the unperturbed case. Since the trajectories of SGD methods will not vary dramatically with a sufficiently small perturbation, they still converge to each solution with approximately the same probability. Therefore, we also expect the SGD methods to fail with high probability.*

6 EXPERIMENTS

Numerical results are presented to support the failure of the gradient descent algorithm. Each MC problem with the B-M factorization formulation (2) is solved by the gradient descent algorithm with a constant step size, where the step size is chosen to be small enough to guarantee that the algorithm converges to a stationary point. Regarding the measurement set, the graph $\mathcal{G}_1 := (\mathcal{V}, \mathcal{E}_1)$ is generated randomly by the Erdos-Renyi model $G(m, p)$, where $\mathcal{V} := [m]$ and each edge of the graph is included independently with probability p . If $\mathcal{G}_1$ is not connected or a node in the maximal independent set $\mathcal{S}$ does not have a self-loop, the missed edges are added to satisfy these conditions. In addition, a connected subtree $\mathcal{G}_2 = (\mathcal{S}, \mathcal{E}_2(\mathcal{S}))$ is generated for nondiagonal observations. We define $\mathcal{G} := (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ and subsequently the measurement set $\Omega(\mathcal{G})$. In addition, the unperturbed ground truth matrix $\mathbf{M}^* = \mathbf{X}^*(\mathbf{X}^*)^T$ is defined as $\mathbf{M}_{i,j}^* := \mathbf{I}_r$ for all $i, j \in \mathcal{S}$, and $\mathbf{M}_{i,j}^* := \mathbf{0}_{r \times r}$ otherwise. Lastly, a Gaussian random perturbation matrix $\epsilon \in \mathbb{R}^{n \times r}$ is generated and normalized, e.g. $\|\epsilon\|_F = 1$. Then, the perturbed ground truth matrix $\mathbf{M}^*(\epsilon) := (\mathbf{X}^* + \gamma\epsilon)(\mathbf{X}^* + \gamma\epsilon)^T$ is generated, where $\gamma > 0$ is the perturbation size. We evaluate the success rate of the algorithm at 100 equally distributed values of $\gamma \in (0, 0.5)$ with 300 random initializations of the gradient algorithm for each instance and each γ .

The top figure in Figure 2 illustrates the success rate of the gradient descent algorithm with the rank $r = 1$, dimension $n = 20$ and various maximum independent set sizes $|\mathcal{S}|$. The results conform to Theorem 4 implying that the success rate is less than $1/(2^{|\mathcal{S}|-1})$ when the perturbation size γ is small. The bottom figure in Figure 2 makes similar observations with different ranks and maximum independent set sizes $|\mathcal{S}|$ when m is equal to 10. These observations imply that Theorem 4 can be extended to the general rank case, since the success rate is less than $1/(2^{r(|\mathcal{S}|-1)})$ when γ is small. We note that the behavior of the algorithm may change when γ is large. Specifically, significant improvements in success rate are observed when $\gamma > 0.2$ for most problem instances. This is in accordance with our notion of complexity metric.

7 CONCLUSION AND DISCUSSION

In this paper, we provided a negative answer to the question of whether the B-M factorization approach




Figure 2: Success rate of gradient descent method for (top) rank-1 and (bottom) rank-r MC problem instances with randomly generated measurement sets.

can capture the benign properties of low-complexity MC problem instances. More specifically, we defined a class of MC problem instances that could be solved in polynomial time. We showed that there exist MC problem instances in this class that have exponentially many spurious local minima in the B-M factorization formulation (1). The results hold for a general class of loss functions, including the commonly used regularized formulation. The incoherence condition determines the number of random samples required to guarantee a high probability of convergence. Nevertheless, there are ground truth matrices that satisfy the incoherence condition and belong to the low-complexity class of problems $\mathcal{L}$ simultaneously. Moreover, there exist convex optimization methods that solve those problems in polynomial time if the incoherence condition is sufficiently small whereas our results indicate that the B-M factorization still fails even with very few missing observations. Furthermore, for the rank-1 case, we proved that gradient-based methods fail with high probability for such instances. Numerical results verify that similar behaviors also hold for higher rank. These results imply that the optimization complexity of methods based on the factorized problem (1) are not aligned with the information-theoretic complexity of the MC problem. Furthermore, we derived a complexity metric that potentially captures the complexity of the B-M factorization formulation (1).

Acknowledgements

This work was supported by grants from AFOSR, ARO, ONR, NSF and C3.ai Digital Transformation Institute.

References

- Bhojanapalli, S. and Jain, P. Universal matrix completion. In *International Conference on Machine Learning*, pp. 1881–1889. PMLR, 2014.
- Bhojanapalli, S., Kyrillidis, A., and Sanghavi, S. Dropping convexity for faster semi-definite optimization. In *Conference on Learning Theory*, pp. 530–582. PMLR, 2016.
- Bi, Y. and 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.
- Bi, Y., Zhang, H., and Lavaei, J. Local and global linear convergence of general low-rank matrix recovery problems. *arXiv preprint arXiv:2104.13348*, 2021.
- Burer, S. and Monteiro, R. D. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- Candès, E. J. and Recht, B. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- Candès, E. J. and Tao, T. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- Chen, J., Liu, D., and Li, X. Nonconvex rectangular matrix completion via gradient descent without $\ell_{2,\infty}$ regularization. *IEEE Transactions on Information Theory*, 66(9):5806–5841, 2020.
- Chi, Y., Lu, Y. M., and Chen, Y. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.
- Davenport, M. A., Plan, Y., Van Den Berg, E., and Wootters, M. 1-bit matrix completion. *Information and Inference: A Journal of the IMA*, 3(3): 189–223, 2014.
- Eisenberg-Nagy, M., Laurent, M., and Varvitsiotis, A. Complexity of the positive semidefinite matrix completion problem with a rank constraint. *Springer, Fields Institute Communications*, vol. 69, 2013.
- Fattahi, S. and Sojoudi, S. Exact guarantees on the absence of spurious local minima for non-negative rank-1 robust principal component analysis. *Journal of machine learning research*, 2020.
- Ge, R., Lee, J. D., and Ma, T. Matrix completion has no spurious local minimum. *Advances in Neural Information Processing Systems*, pp. 2981–2989, 2016.
- Ge, R., Jin, C., and 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.
- Harary, F. Graph theory. *CRC Press*, 2018.
- Khalil, H. K. *Nonlinear systems; 3rd ed.* Prentice-Hall, Upper Saddle River, NJ, 2002.
- Király, F. J., Theran, L., and Tomioka, R. The algebraic combinatorial approach for low-rank matrix completion. *J. Mach. Learn. Res.*, 16(1):1391–1436, 2015.
- Lee, J. D., Simchowitz, M., Jordan, M. I., and Recht, B. Gradient descent only converges to minimizers. In *Conference on learning theory*, pp. 1246–1257. PMLR, 2016.
- Li, Y., Liang, Y., and Risteski, A. Recovery guarantee of weighted low-rank approximation via alternating minimization. In *International Conference on Machine Learning*, pp. 2358–2367. PMLR, 2016.
- Ma, C., Wang, K., Chi, Y., and Chen, Y. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion, and blind deconvolution. *Foundations of Computational Mathematics*, 2019.
- Ma, Y., Olshevsky, A., Szepesvari, C., and Saligrama, V. Gradient descent for sparse rank-one matrix completion for crowd-sourced aggregation of sparsely interacting workers. In *International Conference on Machine Learning*, pp. 3335–3344. PMLR, 2018.
- Pimentel-Alarcón, D. L., Boston, N., and Nowak, R. D. A characterization of deterministic sampling patterns for low-rank matrix completion. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):623–636, 2016.
- Recht, B., Fazel, M., and Parrilo, P. A. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.

- Scieur, D., Roulet, V., Bach, F., and d'Aspremont, A. Integration methods and optimization algorithms. *Advances in Neural Information Processing Systems*, 30, 2017.
- Shechtman, Y., Eldar, Y. C., Cohen, O., Chapman, H. N., Miao, J., and Segev, M. Phase retrieval with application to optical imaging: a contemporary overview. *IEEE signal processing magazine*, 32(3):87–109, 2015.
- Tu, S., Boczar, R., Simchowitz, M., Soltanolkotabi, M., and Recht, B. Low-rank solutions of linear matrix equations via procrustes flow. In *International Conference on Machine Learning*, pp. 964–973. PMLR, 2016.
- Waldspurger, I. and Waters, A. Rank optimality for the burer–monteiro factorization. *SIAM journal on Optimization*, 30(3):2577–2602, 2020.
- Zhang, H., Bi, Y., and Lavaei, J. General low-rank matrix optimization: Geometric analysis and sharper bounds. *arXiv preprint arXiv:2104.10356*, 2021.
- Zhang, R. Y., Sojoudi, S., and 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.

A PROOFS IN SECTION 2

A.1 Proof of Proposition 1

Proof. The condition $\mathcal{P}_{\mathbf{M}^*, \Omega, n, r} \in \mathcal{L}(\mathcal{G}, n, r)$ implies that $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ is connected and non-bipartite. Since the graph is non-bipartite, there exists a cycle with an odd number of vertices $\mathcal{C}_{odd} = (\mathcal{V}_{odd}, \mathcal{E}_{odd})$ in $\mathcal{G}_1$ in which $\mathcal{E}_{odd} \subset \mathcal{E}_1$. To numerically find an odd cycle, the breadth first search method requires $\mathcal{O}(|\mathcal{V}| + |\mathcal{E}_1|) = \mathcal{O}(m^2)$ operations. Without loss of generality, we assume that the set of vertices of the cycle is $\mathcal{V}_{odd} = \{1, 2, \dots, 2k+1\}$ and the set of edges is $\mathcal{E}_{odd} = \{(1, 2), (2, 3), \dots, (2k+1, 1)\}$, where k is a non-negative integer. Suppose that the matrix $\mathbf{X}^* \in \mathbb{R}^{n \times r}$ satisfies $\mathbf{M}^* = \mathbf{X}^*(\mathbf{X}^*)^T$. We denote the i -th $r \times r$ block of $\mathbf{X}^*$ as $\mathbf{X}_i^*$ for all $i \in [m]$, i.e.,

$$\mathbf{M}^* = \mathbf{X}^*(\mathbf{X}^*)^T = \begin{bmatrix} \mathbf{X}_1^* \\ \vdots \\ \mathbf{X}_m^* \end{bmatrix} \begin{bmatrix} (\mathbf{X}_1^*)^T & \dots & (\mathbf{X}_m^*)^T \end{bmatrix}.$$

Since $\mathcal{P}_{\mathbf{M}^*, \Omega, n, r} \in \mathcal{L}(\mathcal{G}, n, r)$, the block $\mathbf{X}_i^*$ is nonsingular for every $i \in [m]$, which further implies that the block $\mathbf{M}_{i,j}^*$ is nonsingular for all $i, j \in [m]$. Using the relation that $\mathbf{M}_{i,j}^* = \mathbf{X}_i^*(\mathbf{X}_j^*)^T$, we can calculate that

$$\left[\prod_{i=1}^k \left(\mathbf{M}_{2i-1, 2i}^* (\mathbf{M}_{2i, 2i+1}^*)^{-T} \right) \right] \mathbf{M}_{2k+1, 1}^* = \mathbf{X}_1^*(\mathbf{X}_1^*)^T,$$

Since the left-hand side only contains observed blocks, the matrix $\mathbf{X}_1^*(\mathbf{X}_1^*)^T$ can be computed via observed blocks. Since computing the inverse of an $r \times r$ matrix and computing the product of two $r \times r$ matrices both require $\mathcal{O}(r^3)$ operations, the total number of operations required for computing $\mathbf{X}_1^*(\mathbf{X}_1^*)^T$ is $\mathcal{O}[(2k+1)r^3] = \mathcal{O}(mr^3)$. In addition, computing the Cholesky decomposition of $\mathbf{X}_1^*(\mathbf{X}_1^*)^T$ requires $\mathcal{O}(r^3)$ operations, which produces a matrix $\mathbf{X}_1^* \mathbf{R}$ for some orthogonal matrix $\mathbf{R} \in \mathbb{R}^{r \times r}$.

With the knowledge of $\mathbf{X}_1^* \mathbf{R}$, we can recursively compute the block $\mathbf{X}_i^*$ using the connectivity of $\mathcal{G}_1$. More specifically, we use $\mathcal{P} \subset [m]$ to denote the set of vertices i or which we have computed $\mathbf{X}_i^* \mathbf{R}$. We start with $\mathcal{P} = \{1\}$. At each iteration, we choose indices $i \in \mathcal{P}$ and $j \notin \mathcal{P}$ such that $(i, j) \in \mathcal{E}_1$. Such a pair of indices always exists unless $\mathcal{P} = [m]$, since the graph $\mathcal{G}_1$ is connected. Then, using the observation

$$\mathbf{M}_{j,i}^* = \mathbf{X}_j^*(\mathbf{X}_i^*)^T = (\mathbf{X}_j^* \mathbf{R})(\mathbf{X}_i^* \mathbf{R})^T,$$

we first compute the matrix $\mathbf{X}_j^* \mathbf{R}$ with $\mathcal{O}(r^3)$ operations and then add j to the set $\mathcal{P}$. We stop the iteration when $\mathcal{P} = [m]$. After this process, we can concatenate $\mathbf{X}_i^* \mathbf{R}$ for all $i \in [m]$ to obtain the matrix $\mathbf{X}^* \mathbf{R}$. The number of iterations is $m-1$ and thus the total number of operations is $\mathcal{O}(mr^3)$.

Summarizing the two parts, the total number of operations to compute the matrix $\mathbf{X}^* \mathbf{U}$ is $\mathcal{O}(m^2 + mr^3) = \mathcal{O}[n^2/r^2 + nr^2]$.

□

A.2 Gradient and Hessian of the Problem (2)

Before proceeding with the analysis for the proofs in the remainder of this paper, we first derive the gradient and the Hessian of the objective function of the problem (2). We omit the proof since the calculation can be done via basic calculus. The gradient of the objective function can be written as

$$\nabla f(\mathbf{X}) = 2(\mathbf{X}\mathbf{X}^T - \mathbf{M}^*)_{\Omega} \mathbf{X}. \quad (6)$$

Similarly, the quadratic variant of the Hessian can be written as

$$\Delta : \nabla^2 f(\mathbf{X}) : \Delta = 4\langle (\mathbf{X}\mathbf{X}^T - \mathbf{M}^*)_{\Omega}, \Delta \Delta^T \rangle + 2\|(\mathbf{X}\Delta^T + \Delta \mathbf{X}^T)_{\Omega}\|_F^2. \quad (7)$$

A.3 Proof of Theorem 1

Proof. Let $\mathcal{S}(\mathcal{G}_1)$ be a maximal independent set of $\mathcal{G}_1$ such that every vertex in the set has a self-loop. We define the global solution as $\mathbf{M}^* := \mathbf{x}^*(\mathbf{x}^*)^T$, where

$$x_i^* := c_i, \quad \forall i \in \mathcal{S}, \quad x_i^* := 0, \quad \forall i \notin \mathcal{S}$$

and $\{c_i \mid \forall i \in \mathcal{S}\}$ is a set of nonzero constants. We note that in the case when $r = 1$, the factor $\mathbf{X} \in \mathbb{R}^n$ is a vector. Therefore, we represent it using the notation for vectors, i.e., $\mathbf{x}$. Considering the problem instance $\mathcal{P}_{\mathbf{M}^*, \Omega(\mathcal{G}), n, 1}$, the set of global solutions of the problem (2) is given by

$$\mathcal{X}^* := \{\mathbf{x} \in \mathbb{R}^n \mid x_i^2 = c_i^2, \forall i \in \mathcal{S}, x_i = 0, \forall i \notin \mathcal{S}\},$$

which has the cardinality $|\mathcal{X}^*| = 2^{|\mathcal{S}|}$. For every global solution $\hat{\mathbf{x}} \in \mathcal{X}^*$, we have $\hat{\mathbf{x}}\hat{\mathbf{x}}^T = \mathbf{M}^*$. Thus, we know that $\hat{\mathbf{x}}$ is a first-order critical point of the problem (2). For every $\Delta \in \mathbb{R}^n$, the quadratic form of the Hessian (7) can be written as

$$\begin{aligned} \Delta : \nabla^2 f(\hat{\mathbf{x}}) : \Delta &= 2 \left\| (\hat{\mathbf{x}}\Delta^T + \Delta\hat{\mathbf{x}}^T)_{\Omega} \right\|_F^2 = \sum_{i \in \mathcal{S}} 2(\Delta_i \hat{x}_i + \hat{x}_i \Delta_i)^2 + \sum_{j \notin \mathcal{S}} \sum_{\substack{i \in \mathcal{S} \\ (i,j) \in \mathcal{E}_1}} [2(\Delta_j \hat{x}_i)^2 + 2(\hat{x}_i \Delta_j)^2] \\ &= \sum_{i \in \mathcal{S}} 4(\Delta_i \hat{x}_i)^2 + \sum_{j \notin \mathcal{S}} \sum_{\substack{i \in \mathcal{S} \\ (i,j) \in \mathcal{E}_1}} 4(\Delta_j \hat{x}_i)^2 = \sum_{i \in \mathcal{S}} 4\hat{x}_i^2 \cdot \Delta_i^2 + \sum_{j \notin \mathcal{S}} \left(\sum_{\substack{i \in \mathcal{S} \\ (i,j) \in \mathcal{E}_1}} 4\hat{x}_i^2 \right) \cdot \Delta_j^2. \end{aligned}$$

The first term in the above expression corresponds to self-loops in $\mathcal{S}$, while the second term corresponds to the edges between $\mathcal{S}$ and $\mathcal{S}^c := \mathcal{V} \setminus \mathcal{S}$. We note that the edges whose endpoints are both in $\mathcal{S}^c$ do not contribute to the quadratic form. Since $\mathcal{S}$ is a maximal independent, we know that

$$\{i \in \mathcal{S} \mid (i, j) \in \mathcal{E}_1\} \neq \emptyset, \quad \forall j \notin \mathcal{S}.$$

As a result, it holds that

$$\Delta : \nabla^2 f(\hat{\mathbf{x}}) : \Delta > 0, \quad \forall \Delta \in \mathbb{R}^n \setminus \{0\},$$

which implies that the Hessian at the global solution $\hat{\mathbf{x}}$ is positive definite. Then, we perturb the global solution of the above problem to be

$$\mathbf{M}^*(\epsilon) := \mathbf{x}^*(\epsilon) [\mathbf{x}^*(\epsilon)]^T = (\mathbf{x}^* + \epsilon)(\mathbf{x}^* + \epsilon)^T,$$

where $\mathbf{x}^*(\epsilon) := \mathbf{x}^* + \epsilon$ and $\epsilon \in \mathbb{R}^n$ is a small perturbation. We denote the problem (2) after perturbation as

$$\min_{\mathbf{x} \in \mathbb{R}^n} \tilde{f}(\mathbf{x}; \epsilon),$$

where $\tilde{f}(\mathbf{x}; \epsilon) := \|(\mathbf{x}\mathbf{x}^T - \mathbf{M}(\epsilon)^*)_{\Omega}\|_F^2$. For a generic perturbation ϵ , all components of ϵ are nonzero and the problem $\mathcal{P}_{\mathbf{M}^*(\epsilon), \Omega, n, 1}$ belongs to the class $\mathcal{L}(\mathcal{G}, n, 1)$. This implies that the global solution of the problem (3) is unique up to a sign flip.

The earlier argument implies that $\nabla_x \tilde{f}(\hat{\mathbf{x}}; 0) = 0$ and $\nabla_{xx} \tilde{f}(\hat{\mathbf{x}}; 0) \succ 0$. Now, we analyze the relationship between the local minima of the original problem and those of the perturbed problem. We consider the equation $\nabla_{\mathbf{x}} \tilde{f}(\mathbf{x}; \epsilon) = 0$ near an unperturbed global minimum $\hat{\mathbf{x}} \in \mathcal{X}^*$. Since $(\hat{\mathbf{x}}; 0)$ is a solution to the gradient equation and the Jacobian matrix with respect to $\mathbf{x}$ is equal to the positive definite Hessian $\nabla^2 f(\hat{\mathbf{x}})$, the Implicit Function Theorem (IFT) states that there exists a unique solution $\hat{\mathbf{x}}(\epsilon)$ in a neighbourhood of $\hat{\mathbf{x}}$ for every ϵ with a small norm. In addition, the continuity of the Hessian implies that $\nabla_{\mathbf{xx}} \tilde{f}(\hat{\mathbf{x}}(\epsilon); \epsilon) \succ 0$. Otherwise, we can find a sequence of $\epsilon^k \rightarrow 0$ and another sequence $\{y^k\}$ with $\|y^k\| = 1$ such that $(y^k)^T \nabla_{\mathbf{xx}} \tilde{f}(\hat{\mathbf{x}}(\epsilon); \epsilon) y^k < 0$. By taking the limit along a convergent subsequence of $\{y^k\}$, we arrive at a contradiction. Therefore, every point in $\hat{\mathbf{x}}^*(\epsilon)$ still satisfies the second-order sufficient conditions. Since there are $2^{|\mathcal{S}|}$ global minima for the unperturbed problem, an analysis through the IFT implies that there are $2^{|\mathcal{S}|}$ strict local minima for the problem (3). Hence, there are $2^{|\mathcal{S}|} - 2$ spurious local minima for the perturbed problem (3).

□

A.4 Proof of Corollary 1

Proof. We first prove the claim about the largest lower bound. Theorem 1 implies that there are at least $2^{|\mathcal{S}(\mathcal{G}_1)|} - 2$ spurious local minima for a problem instance in $\mathcal{L}(\mathcal{G}, n, 1)$. For an arbitrary connected and non-bipartite graph $\mathcal{G}_1$ with n vertices, the maximal possible size of a maximal independent set with self-loops is $n - 1$. More specifically, the graph $\mathcal{G}_1^*$ that attains this maximal value is the star graph $K_{1,n-1}$ complemented with self-loops for the $n - 1$ independent vertices. Hence, Theorem 1 implies that there exists a problem instance in $\mathcal{L}(\mathcal{G}^*, n, 1)$ with at least $2^{n-1} - 2$ local minima, where we define $\mathcal{G}^* := (\mathcal{G}_1^*, \emptyset)$.

We then consider the second claim of this corollary. Consider the measurement set Ω that observes all entries of the ground truth $\mathbf{M}^*$ except $\mathbf{M}_{12}^*$ and $\mathbf{M}_{21}^*$. In this case, we have $|\Omega| = n^2 - 2$. Choose the set of vertices to be $\mathcal{V} := [n]$ and the set of edges to be $\mathcal{E}_1 := [n] \times [n] \setminus \{(1, 2), (2, 1)\}$. Then, the graph $\mathcal{G} := (\mathcal{V}, \mathcal{E}_1, \emptyset)$ satisfies that $\Omega = \Omega(\mathcal{G})$ and the maximal independent set is $\mathcal{S} := \{i, j\}$. Therefore, Theorem 1 implies that there exists a problem instance in $\mathcal{L}(\mathcal{G}, n, 1)$ that has at least $2^{|\mathcal{S}|} - 2 = 2$ spurious local minima. $\square$

A.5 Proof of Corollary 2

Proof. For a generic vector $\mathbf{x}^* \in \mathbb{R}^n$, all elements of $\mathbf{x}^*$ are nonzero and we can decompose $\mathbf{x}^*$ into $\mathbf{x}^0 + \mathbf{x}^1$, where

$$\begin{aligned} x_i^0 &:= x_i^*, \quad \forall i \in \mathcal{S}, & x_i^0 &:= 0, & \forall i \notin \mathcal{S}, \\ x_i^1 &:= 0, \quad \forall i \in \mathcal{S}, & x_i^1 &:= x_i^*, & \forall i \notin \mathcal{S}. \end{aligned}$$

We first consider the problem $\mathcal{P}_{\mathbf{x}^0(\mathbf{x}^0)^T, \Omega(\mathcal{G}), n, 1}$. Using a similar proof as Theorem 1, we can prove that this problem has $2^{|\mathcal{S}|}$ equivalent global solutions in formulation (2), which are described by the set

$$\mathcal{X}^* := \{\mathbf{x} \in \mathbb{R}^n \mid x_i^2 = (x_i^*)^2, \forall i \in \mathcal{S}, x_i = 0, \forall i \notin \mathcal{S}\}.$$

In addition, the Hessian is positive definite at every global solution $\hat{\mathbf{x}} \in \mathcal{X}^*$. Using a similar argument as the proof of Theorem 1, we conclude that there exists a small constant $r_{\mathbf{x}^0} > 0$ such that the conditions

$$\|\epsilon\| \leq r_{\mathbf{x}^0}, \quad \epsilon_i \neq 0, \quad \forall i \in [m] \tag{8}$$

imply that the problem $\mathcal{P}_{\mathbf{x}^0(\epsilon)[\mathbf{x}^0(\epsilon)]^T, \Omega(\mathcal{G}), n, 1}$ has at least $2^{|\mathcal{S}|} - 2$ spurious local minima in formulation (2), where $\mathbf{x}^0(\epsilon) := \mathbf{x}^0 + \epsilon$. Moreover, the MC problem is “scale-free” in the sense that the formulation (2) of the problem $\mathcal{P}_{\mathbf{x}'(\mathbf{x}')^T, \Omega(\mathcal{G}), n, 1}$ has spurious local minima if and only if that of the problem $\mathcal{P}_{(c\mathbf{x}')(c\mathbf{x}')^T, \Omega(\mathcal{G}), n, 1}$ has spurious local minima, where $\mathbf{x}' \in \mathbb{R}^n$ is an arbitrary vector and $c \neq 0$ is a constant. Therefore, we have the relation

$$r_{c\mathbf{x}^0} = c \cdot r_{\mathbf{x}^0}, \quad \forall c \neq 0. \tag{9}$$

Hence, it suffices to consider vectors $\mathbf{x}^0 \in \mathcal{X}_0$, where

$$\mathcal{X}_0 := \{\mathbf{x} \in \mathbb{R}^n \mid x_i \neq 0, \forall i \in \mathcal{S}, x_i = 0, \forall i \notin \mathcal{S}, \|\mathbf{x}\| = 1\}.$$

Now, we consider a vector $\hat{\mathbf{x}}^0 \in \mathbb{R}^n$ satisfying

$$\|\hat{\mathbf{x}}^0 - \mathbf{x}^0\| < r_{\mathbf{x}^0}/2.$$

Then, as long as the generic perturbation ϵ satisfies $\|\epsilon\| \leq r_{\mathbf{x}^0}/2$, the condition (8) implies that the problem $\mathcal{P}_{\hat{\mathbf{x}}^0(\epsilon)[\hat{\mathbf{x}}^0(\epsilon)]^T, \Omega(\mathcal{G}), n, 1}$ has at least $2^{|\mathcal{S}|} - 2$ spurious local minima in formulation (2), where $\hat{\mathbf{x}}^0(\epsilon) := \hat{\mathbf{x}}^0 + \epsilon$. This verifies the existence of a function $h_{\mathcal{S}}$ on the open set

$$\mathcal{N}(\mathbf{x}^0) := \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x} - \mathbf{x}^0\| < r_{\mathbf{x}^0}/2\}, \quad \forall \mathbf{x}^0 \in \mathcal{X}_0.$$

Now, we consider the compact set

$$\mathcal{X}_k := \left\{ \mathbf{x} \in \mathbb{R}^n \mid \min_{i \in \mathcal{S}} |x_i| \geq 1/k, x_i = 0, \forall i \notin \mathcal{S}, \|\mathbf{x}\| = 1 \right\}, \quad k = 1, 2, \dots$$

The open set family $\{\mathcal{N}(\mathbf{x}^0) \mid \mathbf{x}^0 \in \mathcal{X}_0\}$ comprises an open covering of the compact set $\mathcal{X}_k$. Hence, there exists a finite covering of the compact set $\mathcal{X}_k$ and thus there exists a small constant ϵ_k such that

$$h_{\mathcal{S}}(\mathbf{x}^0) \geq \epsilon_k, \quad \forall \mathbf{x}^0 \in \mathcal{X}_k.$$

Moreover, using the relation

$$\mathcal{X}_{k-1} \subset \mathcal{X}_k, \quad \forall k \geq 2, \quad \mathcal{X}_0 = \bigcup_{k=1}^{\infty} \mathcal{X}_k,$$

we can decrease the value of the function $h_{\mathcal{S}}$ to be

$$h(\mathbf{x}^0) := \epsilon_k, \quad \forall \mathbf{x}^0 \in \mathcal{X}_k \setminus \mathcal{X}_{k-1}, \quad k \geq 2.$$

Using this definition, $h_{\mathcal{S}}$ reduces to $\min_{i \in \mathcal{S}} |x_i|$ and, with a little abuse of notations, we still write the new function as $h_{\mathcal{S}}(\min_{i \in \mathcal{S}} |x_i|)$.

Now, we can view the problem $\mathcal{P}_{\mathbf{x}^*(\mathbf{x}^*)^T, \Omega(\mathcal{G}), n, 1}$ as the perturbed problem, where the generic perturbation is given by $\mathbf{x}^1$. The above analysis implies that the following condition holds

$$\|\mathbf{x}_{\mathcal{S}^c}\| = \|\mathbf{x}^1\| \leq h_{\mathcal{S}}\left(\min_{i \in \mathcal{S}} |x_i|\right) \cdot \|\mathbf{x}^0\| = h_{\mathcal{S}}\left(\min_{i \in \mathcal{S}} |x_i|\right) \cdot \|\mathbf{x}_{\mathcal{S}}^*\|$$

which demonstrates the existence of at least $2^{|\mathcal{S}|} - 2$ spurious local minima in formulation (2). □

A.6 Proof of Lemma 1

Proof. We follow a similar proof construction as in Theorem 1. Let $\mathcal{D}$ be the set of full-rank diagonal matrices and $\mathcal{D}_{\{1, -1\}}$ be the set of diagonal matrices with the diagonal entries being +1 or -1. Suppose that $\mathcal{S}$ is a maximal independent set of $\mathcal{G}_1$ in which every node has a self-loop. Then, we define the ground truth matrix $\mathbf{M}^* := \mathbf{X}^*(\mathbf{X}^*)^T$, where $\mathbf{X}^* \in \mathbb{R}^{n \times r}$ satisfies

$$\mathbf{X}_i^* = 0, \quad \forall i \notin \mathcal{S}, \quad \mathbf{X}_i^* \in \mathcal{D}, \quad \forall i \in \mathcal{S},$$

where $\mathbf{X}_i$ is the i -th block of $\mathbf{X} \in \mathbb{R}^{n \times r}$; see the definition in the proof of Proposition 1. Then, the set of global solutions is given by

$$\mathcal{X}^* := \left\{ \mathbf{X} \in \mathbb{R}^{n \times r} \mid \mathbf{X}_i = \mathbf{X}_i^* \mathbf{D}, \quad \forall i \in \mathcal{S}, \quad \mathbf{D} \in \mathcal{D}_{\{1, -1\}}, \quad \mathbf{X}_i = 0, \quad \forall i \notin \mathcal{S} \right\},$$

For every global solution $\hat{\mathbf{X}} \in \mathcal{X}^*$, we have $\hat{\mathbf{X}}\hat{\mathbf{X}}^T = \mathbf{M}^*$. Thus, every global solution is a first-order critical point of the problem (2). In addition, let $\Delta \in \mathbb{R}^{n \times r}$ be an arbitrary direction matrix with its $r \times r$ block matrices denoted as $\Delta_1, \Delta_2, \dots, \Delta_m$. Then, the quadratic variant of the Hessian (7) in the direction Δ can be written as

$$\begin{aligned} \Delta : \nabla^2 f(\hat{\mathbf{X}}) : \Delta &= 2 \left\| \left(\hat{\mathbf{X}} \Delta^T + \Delta \hat{\mathbf{X}}^T \right)_{\Omega} \right\|_F^2 \\ &= \sum_{i \in \mathcal{S}} \|\Delta_i \hat{\mathbf{X}}_i^T + \hat{\mathbf{X}}_i \Delta_i^T\|_F^2 + 2 \sum_{j \notin \mathcal{S}} \left(\sum_{\substack{i \in \mathcal{S} \\ (i,j) \in \mathcal{E}_1}} \|\Delta_j \hat{\mathbf{X}}_i^T\|_F^2 + \sum_{\substack{i \in \mathcal{S} \\ (i,j) \in \mathcal{E}_2}} \|(\Delta_j \hat{\mathbf{X}}_i^T)_{nd}\|_F^2 \right) \\ &\quad + 2 \sum_{j \in \mathcal{S}} \sum_{\substack{i \in \mathcal{S} \\ (i,j) \in \mathcal{E}_2}} \|(\Delta_j \hat{\mathbf{X}}_i^T + \hat{\mathbf{X}}_j \Delta_i^T)_{nd}\|_F^2 \end{aligned} \tag{10}$$

where $(\cdot)_{nd}$ is the projection onto the matrix space with a zero diagonal. We note that the first term in (10) corresponds to self-loops in $\mathcal{G}_1$, while the second term corresponds to edges between $\mathcal{S}$ and $\mathcal{S}^c$. The

edges whose endpoints are both in $\mathcal{S}^c$ do not contribute to the quadratic form. Moreover, the last term corresponds to partial observations with nondiagonal entries within the independent set $\mathcal{S}$.

Now, we aim to prove that the Hessian at $\hat{\mathbf{X}}$ is positive definite in the tangent space of $\mathcal{W}^{n \times r}$, namely,

$$\Delta : \nabla^2 f(\hat{\mathbf{X}}) : \Delta > 0, \quad \forall \Delta \in \mathbb{R}^{n \times r} \setminus \{0\}, \quad \Delta_1 \text{ is lower triangular.}$$

We assume that $\Delta : \nabla^2 f(\hat{\mathbf{X}}) : \Delta = 0$ for some $\Delta \in \mathbb{R}^{n \times r}$ such that Δ_1 is lower triangular. Under this assumption, all three terms in (10) are equal to zero. Considering the second term, since $\hat{\mathbf{X}}_i$ is full-rank, we have

$$\Delta_j = 0, \quad \forall j \notin \mathcal{S}.$$

For the first term, $\|\Delta_i \hat{\mathbf{X}}_i^T + \hat{\mathbf{X}}_i \Delta_i^T\|_F^2$ is zero only if $\Delta_i \hat{\mathbf{X}}_i^T = -\hat{\mathbf{X}}_i \Delta_i^T$, i.e., $\Delta_i \hat{\mathbf{X}}_i^T$ is skew-symmetric. Since $\hat{\mathbf{X}}_i$ is a diagonal matrix with nonzero diagonal entries, the diagonal entries of Δ_i must be zero for all $i \in \mathcal{S}$. Without loss of generality, we assumed that vertex $1 \in \mathcal{S}$. This is because we can equivalently fix the block $\mathbf{X}_i$ to be lower-diagonal for any $i \in \mathcal{S}$ and consider a similarly constrained optimization problem. Then, since Δ_1 must be lower triangular, we have

$$\Delta_1 = 0,$$

We define the set

$$\mathcal{S}_0 := \{i \in \mathcal{S} \mid \Delta_i = 0\}.$$

We have shown that $i \in \mathcal{S}_0$ and aim to prove that $\mathcal{S}_0 = \mathcal{S}$. Since the induced subgraph $\mathcal{G}_2[\mathcal{S}]$ is connected, there exists a vertex $j \in \mathcal{S}$ such that $(1, j) \in \mathcal{E}_2$. Considering the third term in (10), we have

$$(\Delta_j \hat{\mathbf{X}}_1^T + \hat{\mathbf{X}}_j \Delta_1^T)_{nd} = (\Delta_j \hat{\mathbf{X}}_1^T)_{nd} = 0,$$

which implies that $\Delta_j = 0$ because $\hat{\mathbf{X}}_j$ is a diagonal matrix with nonzero diagonal entries. Hence, we have proved that $j \in \mathcal{S}_0$. By the connectivity of $\mathcal{G}_2[\mathcal{S}]$, we can inductively prove that all elements in $\mathcal{S}$ belong to $\mathcal{S}_0$. Therefore, it holds that $\Delta = 0$ and the quadratic form of Hessian is zero only when $\Delta = 0$. As a result, the Hessian is positive definite at every global solution $\hat{\mathbf{X}} \in \mathcal{X}^*$ of the problem (4).

Then, we perturb the ground truth of the above problem instance to be

$$\mathbf{M}^*(\epsilon) := \mathbf{X}^*(\epsilon) [\mathbf{X}^*(\epsilon)]^T = (\mathbf{X}^* + \epsilon)(\mathbf{X}^* + \epsilon)^T,$$

where $\mathbf{X}^*(\epsilon) := \mathbf{X}^* + \epsilon$ and $\epsilon \in \mathbb{R}^{n \times r}$ is a small perturbation. Similar to Theorem 1, for a generic perturbation ϵ , all block components of ϵ are nonzero and full-rank. Therefore, the problem $\mathcal{P}_{\mathbf{M}^*(\epsilon), \Omega, n, r}$ belongs to the class $\mathcal{L}(\mathcal{G}, n, r)$. This implies that the global solution of the problem (4) is unique up to a right-multiplication with $\mathbf{D} \in \mathcal{D}_{\{1, -1\}}$. Since there are $2^{r|\mathcal{S}|}$ global minima for the unperturbed problem, IFT implies that there are $2^{r|\mathcal{S}|}$ strict local minima for the perturbed problem. Hence, there are $2^{r|\mathcal{S}|} - 2^r$ spurious local minima for the perturbed problem. $\square$

A.7 Proof of Corollary 3

Proof. We first consider the largest lower bound on the number of spurious local minima. Similar to Corollary 1, since an instance in $\mathcal{L}(\mathcal{G}, n, r)$ can have a maximum independent set of size at most $|\mathcal{S}(\mathcal{G}_1)| = m - 1 = n/r - 1$, Theorem 2 implies that the largest lower bound on the number of spurious local solution classes is $2^{r(n/r-2)} - 1 = 2^{n-2r} - 1$.

We prove the second part of this corollary next. We choose $\mathcal{G}_1$ to be a graph with m vertices and $\binom{m}{2} - 1$ edges, where (i, j) is the only missing edge. Then, the maximal independent set is $\mathcal{S} := \{i, j\}$. Since $\mathcal{G}_2[\mathcal{S}]$ must be connected, the nondiagonal entries of the block $\mathbf{M}_{i,j}^*$ are observed. Thus, only $2r$ entries are not observed and $|\Omega| = n^2 - 2r$. Furthermore, Theorem 2 implies that there exists a problem instance in $\mathcal{L}(\mathcal{G}, n, r)$ with at least $2^{r(|\mathcal{S}|-1)} - 1 = 2^r - 1$ equivalence classes of spurious local minima. $\square$

A.8 Proof of Theorem 3

Proof. The proof is similar to those of Theorems 1 and 2. We can prove that the Hessian is positive definite at all global solutions for each loss function $g(\cdot)$ satisfying Assumption 1. $\square$

B PROOFS IN SECTION 3

B.1 Proof of Proposition 2

Proof. For every matrix $\mathbf{X} \in \mathbb{R}^{n \times r}$, we denote $\mathbf{X}_i$ as the i -th $r \times r$ block of $\mathbf{X}$ for all $i \in [m]$. Because each block $\mathbf{M}_{i,j}^*$ is assumed to be full rank, the block $\mathbf{X}_i^*$ is also full rank for all $i \in [m]$, where $\mathbf{X}^*$ satisfies $\mathbf{M}^* = \mathbf{X}^*(\mathbf{X}^*)^T$. It is desirable show that every first-order critical point is either a global solution or a saddle point with a strict descent direction. For every $i \in [m]$, the gradient of the problem (2) with respect to the i -th block $\mathbf{X}_i$ is

$$\nabla_{\mathbf{X}_i} f(\mathbf{X}) = \begin{cases} 2(\mathbf{X}_i \mathbf{X}_k^T - \mathbf{X}_i^*(\mathbf{X}_k^*)^T) \mathbf{X}_k, & \text{if } i \neq k \\ \sum_{j=1}^m 2(\mathbf{X}_k \mathbf{X}_j^T - \mathbf{X}_k^*(\mathbf{X}_j^*)^T) \mathbf{X}_j, & \text{if } i = k. \end{cases}$$

Let $\hat{\mathbf{X}} \in \mathbb{R}^{n \times r}$ be a first-order critical point of problem (2).

We first consider the case when $\hat{\mathbf{X}}_k$ is nonsingular. For every $i \in [m] \setminus \{k\}$, the condition $\nabla_{\mathbf{X}_i} f(\mathbf{X}) = 0$ implies that

$$\hat{\mathbf{X}}_i \hat{\mathbf{X}}_k^T - \mathbf{X}_i^*(\mathbf{X}_k^*)^T = 0.$$

Substituting the above equations into the equation $\nabla_{\mathbf{X}_k} f(\mathbf{X}) = 0$, we obtain

$$(\mathbf{X}_k \mathbf{X}_k^T - \mathbf{X}_k^*(\mathbf{X}_k^*)^T) \mathbf{X}_k = 0,$$

which implies that $\hat{\mathbf{X}}_k \hat{\mathbf{X}}_k^T = \mathbf{X}_k^*(\mathbf{X}_k^*)^T$. Therefore, the matrix $\hat{\mathbf{X}}$ is a global solution of the problem (2) in this case.

Now, we consider the case when $\hat{\mathbf{X}}_k$ is singular. We choose a vector $\mathbf{y}_k \in \mathbb{R}^n$ such that

$$\hat{\mathbf{X}}_k \mathbf{y}_k = 0, \quad \|\mathbf{y}_k\| = 1.$$

Given a small constant $\epsilon > 0$, the i -th block direction $\Delta \in \mathbb{R}^{n \times r}$ is defined as

$$\Delta_i := \begin{cases} \mathbf{z}_i \mathbf{y}_k^T, & \text{if } i \neq k \\ \epsilon \mathbf{y}_k \mathbf{y}_k^T, & \text{if } i = k, \end{cases}$$

where $\mathbf{z}_i \in \mathbb{R}^n$ is arbitrary. The above definition directly implies that $\hat{\mathbf{X}}_k \Delta_i^T = 0$ for all $i \in [m]$. Then, we obtain

$$\begin{aligned} 4 \left\langle (\hat{\mathbf{X}} \hat{\mathbf{X}}^T - \mathbf{X}^*(\mathbf{X}^*)^T)_\Omega, \Delta \Delta^T \right\rangle &= -4 \text{tr} [\mathbf{X}_k^*(\mathbf{X}_k^*)^T \Delta_k \Delta_k^T] - \sum_{j=1, j \neq i}^m 8 \text{tr} (\mathbf{X}_j^*(\mathbf{X}_j^*)^T \Delta_k \Delta_j^T) \\ &= \sum_{j=1, j \neq i}^m -8 \text{tr} [\mathbf{X}_j^*(\mathbf{X}_k^*)^T \mathbf{y}_k \mathbf{z}_j^T] \cdot \epsilon + \mathcal{O}(\epsilon^2), \end{aligned}$$

and

$$2 \|(\mathbf{X} \Delta^T + \Delta \mathbf{X}^T)_\Omega\|_F^2 = 8 \|\mathbf{X}_k \Delta_k^T\|_F^2 + 4 \sum_{j=1, j \neq k}^m \|\mathbf{X}_j \Delta_k^T\|_F^2 = \mathcal{O}(\epsilon^2)$$

Combining two estimates above, the quadratic form of the Hessian (7) can be written as

$$\Delta : \nabla^2 f(\hat{\mathbf{X}}) : \Delta = \sum_{j=1, j \neq i}^m -8 \text{tr} [\mathbf{X}_j^*(\mathbf{X}_k^*)^T \mathbf{y}_k \mathbf{z}_j^T] \cdot \epsilon + \mathcal{O}(\epsilon^2).$$

Since $\mathbf{X}_i^*$ is nonsingular for all $i \in [m]$, it holds that $\mathbf{X}_j^*(\mathbf{X}_k^*)^T \mathbf{y}_k \neq 0$. Choosing

$$\mathbf{z}_j := \mathbf{X}_j^*(\mathbf{X}_k^*)^T \mathbf{y}_k,$$

we obtain

$$\Delta : \nabla^2 f(\hat{\mathbf{X}}) : \Delta = \sum_{j=1, j \neq i}^m -8 \|\mathbf{X}_j^* (\mathbf{X}_k^*)^T \mathbf{y}_k\|^2 \cdot \epsilon + \mathcal{O}(\epsilon^2).$$

Hence, the quadratic form of the Hessian is negative with a sufficiently small ϵ and $\hat{\mathbf{X}}$ is a strict saddle point.

Combining the two cases, we conclude that every second-order critical point of the problem (2) is a global minimum. $\square$

C PROOFS IN SECTION 5

C.1 Proof of Lemma 2

In this proof and the following proofs for Section 5, we consider the instance of the MC problem constructed in Section 2. For completeness, we repeat the instance here. In the unperturbed case, the ground truth matrix is defined as $\mathbf{M}^* := \mathbf{x}^* (\mathbf{x}^*)^T \in \mathbb{R}^{n \times n}$, where vector $\mathbf{x}^* \in \mathbb{R}^n$ satisfies

$$\mathbf{x}_{2k-1}^* = 1, \quad \forall k = 1, \dots, \lceil n/2 \rceil, \quad \mathbf{x}_{2k}^* = 0, \quad \forall k = 1, \dots, \lfloor n/2 \rfloor.$$

The measurement set Ω is given by

$$\Omega := \{(j, j), (2k, j), (j, 2k) \mid j = 1, \dots, n, k = 1, \dots, \lfloor n/2 \rfloor\}.$$

It has been proved in Section 2 that the problem (1) has $2^{\lceil n/2 \rceil}$ global solutions, which are given by the set

$$\mathcal{X}^* := \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}_{2k} = 0, k = 1, \dots, \lfloor n/2 \rfloor, \mathbf{x}_{2k+1}^2 = 1, k = 1, \dots, \lceil n/2 \rceil\}.$$

For each vector $\mathbf{x} \in \mathbb{R}^n$, we denote

$$\mathbf{x}^o := (x_1, x_3, \dots, x_{\lceil n/2 \rceil}), \quad \mathbf{x}^e := (x_2, x_4, \dots, x_{\lfloor n/2 \rfloor}).$$

Before presenting the proof of Lemma 2, we first state three technical lemmas.

Lemma 3. *Suppose that Assumption 1 holds with $(\delta, 1)$, and that $\hat{\mathbf{x}}$ is a first-order critical point of the problem (1). Then, it holds that*

$$\|\hat{\mathbf{x}}^e\| \|\hat{\mathbf{x}}\| \leq 2\sqrt{2}\delta \|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F,$$

where we define $\mathbf{M} := \hat{\mathbf{x}}\hat{\mathbf{x}}^T$.

Proof. Utilizing the first-order optimality condition and the gradient in (6), it holds that

$$\langle \nabla g[(\hat{\mathbf{x}}\hat{\mathbf{x}}^T - \mathbf{x}^* \mathbf{x}^*)_\Omega], \hat{\mathbf{x}}\Delta^T \rangle = \int_0^1 (\mathbf{M} - \mathbf{M}^*) : \nabla^2 g[(\mathbf{M}^*)_\Omega + t(\mathbf{M} - \mathbf{M}^*)_\Omega] : \hat{\mathbf{x}}\Delta^T dt = 0, \quad \forall \Delta \in \mathbb{R}^n,$$

where the first equality is from Taylor's expansion. For every fixed number $t \in [0, 1]$, the proof of Theorem 1 in Bi & Lavaei (2021) implies that

$$\begin{aligned} & (\mathbf{M} - \mathbf{M}^*) : \nabla^2 g[(\mathbf{M}^*)_\Omega + t(\mathbf{M} - \mathbf{M}^*)_\Omega] : \hat{\mathbf{x}}\Delta^T \\ & \geq \langle (\mathbf{M} - \mathbf{M}^*)_\Omega, (\hat{\mathbf{x}}\Delta^T)_\Omega \rangle - 2\sqrt{2}\delta \|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F \|(\hat{\mathbf{x}}\Delta^T)_\Omega\|_F. \end{aligned}$$

Integrating over t , it follows that

$$\langle (\mathbf{M} - \mathbf{M}^*)_\Omega, (\hat{\mathbf{x}}\Delta^T)_\Omega \rangle \leq 2\sqrt{2}\delta \|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F \|(\hat{\mathbf{x}}\Delta^T)_\Omega\|_F. \quad (11)$$

By choosing

$$\Delta_{2k+1} = 0, \quad k = 1, \dots, \lceil n/2 \rceil, \quad \Delta_{2k} = \hat{x}_{2k}, \quad k = 1, \dots, \lfloor n/2 \rfloor,$$

we obtain

$$\langle (\mathbf{M} - \mathbf{M}^*)_{\Omega}, (\hat{\mathbf{x}}\Delta^T)_{\Omega} \rangle = \|\hat{\mathbf{x}}^e\|^2 \|\hat{\mathbf{x}}\|^2, \quad \|(\hat{\mathbf{x}}\Delta^T)_{\Omega}\|_F = \|\hat{\mathbf{x}}^e\| \|\hat{\mathbf{x}}\|.$$

Substituting the above two equalities into (11), we have

$$\|\hat{\mathbf{x}}^e\|^2 \|\hat{\mathbf{x}}\|^2 \leq 2\sqrt{2}\delta \|(\mathbf{M} - \mathbf{M}^*)_{\Omega}\|_F \|\hat{\mathbf{x}}^e\| \|\hat{\mathbf{x}}\|.$$

The above inequality implies that

$$\|\hat{\mathbf{x}}^e\| \|\hat{\mathbf{x}}\| \leq 2\sqrt{2}\delta \|(\mathbf{M} - \mathbf{M}^*)_{\Omega}\|_F \quad \text{or} \quad \hat{\mathbf{x}}^e = 0,$$

since $\|\hat{\mathbf{x}}^e\| \|\hat{\mathbf{x}}\| = 0$ if and only if $\hat{\mathbf{x}}^e = 0$. In both cases, the claim of this lemma holds. $\square$

Lemma 4. Let $\mathcal{D}$ be the set of $r \times r$ diagonal matrices and $\mathcal{D}_{\{1,-1\}}$ be set of $r \times r$ diagonal matrices with the diagonal entries $+1$ or -1 . Under the same setting as Lemma 1, consider the $n \times n$ ground truth matrix $\mathbf{M}^* := \mathbf{X}^*(\mathbf{X}^*)^T$ such that $n = mr$, where

$$\mathbf{X}_i^* \in \mathcal{D}, \quad \forall i \in \mathcal{S}(\mathcal{G}_1), \quad \mathbf{X}_i^* = \mathbf{0}, \quad \forall i \notin \mathcal{S}(\mathcal{G}_1).$$

Let $\mathbf{D} \in \mathbb{R}^{n \times n}$ be an arbitrary matrix with its diagonal blocks denoted as $\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_m$ such that $\mathbf{D}_i \in \mathcal{D}_{\{1,-1\}}$ for all $i \in [m]$. Then, the problem instance $\mathcal{P}_{\mathbf{M}^*, \Omega(\mathcal{G}), n, r}$ in formulation (2) satisfies that

$$\nabla f(\mathbf{D}\mathbf{X}) = \mathbf{D}\nabla f(\mathbf{X}), \quad \forall \mathbf{X} \in \mathbb{R}^{n \times r}.$$

Proof. Since the graph $\mathcal{G}$ satisfies the conditions in Lemma 1, the gradient $\nabla f(\mathbf{X})$ for problem (2) can be written as

$$\nabla_{\mathbf{X}_k} f(\mathbf{D}\mathbf{X}) = 2 \begin{cases} (\mathbf{X}_k \mathbf{X}_k^T - \mathbf{X}_k^* (\mathbf{X}_k^*)^T) \mathbf{X}_k + \sum_{(k,j) \in \mathcal{E}_1} \mathbf{X}_k \mathbf{X}_j^T \mathbf{X}_j + \sum_{(k,j) \in \mathcal{E}_2} (\mathbf{X}_k \mathbf{X}_j^T - \mathbf{X}_k^* (\mathbf{X}_j^*)^T)_{nd} \mathbf{X}_j, & \text{if } k \in \mathcal{S} \\ \sum_{(k,j) \in \mathcal{E}_1} \mathbf{X}_k \mathbf{X}_j^T \mathbf{X}_j + \sum_{(k,j) \in \mathcal{E}_2} (\mathbf{X}_k \mathbf{X}_j^T)_{nd} \mathbf{X}_j, & \text{if } k \notin \mathcal{S}, \end{cases}$$

where $(\cdot)_{nd}$ is the projection onto the matrix space with the zero diagonals. Note that blocks of the transformed variable $\mathbf{D}\mathbf{X}$ are $\mathbf{D}_i \mathbf{X}_i$ for all $i \in [m]$. First, we consider the change in the i -th block of the gradient function for $i \in \mathcal{S}$:

$$\begin{aligned} \nabla_{\mathbf{X}_i} f(\mathbf{D}\mathbf{X}) &= 2 \left((\mathbf{D}_i \mathbf{X}_i \mathbf{X}_i^T \mathbf{D}_i^T - \mathbf{X}_i^* (\mathbf{X}_i^*)^T) \mathbf{D}_i \mathbf{X}_i + \sum_{(i,j) \in \mathcal{E}_1} \mathbf{D}_i \mathbf{X}_i \mathbf{X}_j^T \mathbf{D}_j^T \mathbf{D}_j \mathbf{X}_j + \right. \\ &\quad \left. \sum_{(i,j) \in \mathcal{E}_2} (\mathbf{D}_i \mathbf{X}_i \mathbf{X}_j^T \mathbf{D}_j^T - \mathbf{X}_i^* (\mathbf{X}_j^*)^T)_{nd} \mathbf{D}_j \mathbf{X}_j \right). \end{aligned} \quad (12)$$

Using $\mathbf{D}_i^T \mathbf{D}_i = \mathbf{I}$ and $\mathbf{X}_i^* \in \mathcal{D}$, the first term in (12) can be written as

$$\begin{aligned} (\mathbf{D}_i \mathbf{X}_i \mathbf{X}_i^T \mathbf{D}_i^T - \mathbf{X}_i^* (\mathbf{X}_i^*)^T) \mathbf{D}_i \mathbf{X}_i &= \mathbf{D}_i \mathbf{X}_i \mathbf{X}_i^T \mathbf{X}_i - \mathbf{X}_i^* (\mathbf{X}_i^*)^T \mathbf{D}_i \mathbf{X}_i \\ &= \mathbf{D}_i (\mathbf{X}_i \mathbf{X}_i^T - \mathbf{X}_i^* (\mathbf{X}_i^*)^T) \mathbf{X}_i, \end{aligned}$$

where the second equality is justified by the commutative property of diagonal matrix multiplication. Similarly, the second term in (12) can be written as

$$\sum_{(i,j) \in \mathcal{E}_1} \mathbf{D}_i \mathbf{X}_i \mathbf{X}_j^T \mathbf{D}_j^T \mathbf{D}_j \mathbf{X}_j = \mathbf{D}_i \sum_{(i,j) \in \mathcal{E}_1} \mathbf{X}_i \mathbf{X}_j^T \mathbf{X}_j.$$

For the last term in (12), we use the relation $\mathbf{X}_i^* \in \mathcal{D}$ and the fact that $(\cdot)_{nd}$ is nonzero only at positions associated with the nondiagonal entries to obtain

$$\begin{aligned} \sum_{(i,j) \in \mathcal{E}_2} (\mathbf{D}_i \mathbf{X}_i \mathbf{X}_j^T \mathbf{D}_j^T - \mathbf{X}_i^* (\mathbf{X}_j^*)^T)_{nd} \mathbf{D}_j \mathbf{X}_j &= \sum_{(i,j) \in \mathcal{E}_2} \mathbf{D}_i (\mathbf{X}_i \mathbf{X}_j^T - \mathbf{X}_i^* (\mathbf{X}_j^*)^T)_{nd} \mathbf{D}_j^T \mathbf{D}_j \mathbf{X}_j \\ &= \mathbf{D}_i \sum_{(i,j) \in \mathcal{E}_2} (\mathbf{X}_i \mathbf{X}_j^T - \mathbf{X}_i^* (\mathbf{X}_j^*)^T)_{nd} \mathbf{X}_j. \end{aligned}$$

Thus, we have

$$\nabla_{\mathbf{x}_i} f(\mathbf{D}\mathbf{X}) = \mathbf{D}_i \nabla_{\mathbf{x}_i} f(\mathbf{X}), \quad \forall i \in \mathcal{S}.$$

Now, we consider the change in the i -th block of the gradient function for $i \notin \mathcal{S}$:

$$\begin{aligned} \nabla_{\mathbf{x}_i} f(\mathbf{D}\mathbf{X}) &= 2 \left(\sum_{(i,j) \in \mathcal{E}_1} \mathbf{D}_i \mathbf{x}_i \mathbf{x}_j^T \mathbf{D}_j^T \mathbf{D}_j \mathbf{x}_j + \sum_{(i,j) \in \mathcal{E}_2} (\mathbf{D}_i \mathbf{x}_i \mathbf{x}_j^T \mathbf{D}_j^T)_{nd} \mathbf{D}_j \mathbf{x}_j \right) \\ &= 2 \mathbf{D}_i \left(\sum_{(i,j) \in \mathcal{E}_1} \mathbf{x}_i \mathbf{x}_j^T \mathbf{x}_j + \sum_{(i,j) \in \mathcal{E}_2} (\mathbf{x}_i \mathbf{x}_j^T)_{nd} \mathbf{x}_j \right) = \mathbf{D}_i \nabla_{\mathbf{x}_i} f(\mathbf{X}), \end{aligned}$$

where the second equality holds by a similar argument as in the case of $i \in \mathcal{S}$. Consequently, we have

$$\nabla_{\mathbf{x}_i} f(\mathbf{D}\mathbf{X}) = \mathbf{D}_i \nabla_{\mathbf{x}_i} f(\mathbf{X}), \quad \forall i \notin \mathcal{S}.$$

Combining the two cases, it follows that

$$\nabla f(\mathbf{D}\mathbf{X}) = \mathbf{D} \nabla f(\mathbf{X}).$$

□

Lemma 5. Consider the case $r = 1$. Given an arbitrary point $x_0 \in \mathbb{R}^n$, let $\hat{x} \in \mathbb{R}^n$ denote a point with the property that the gradient flow (5) initialized at $\mathbf{x}_0$ converges to $\hat{\mathbf{x}}$. For every diagonal matrix $\mathbf{D} \in \mathbb{R}^{n \times n}$ that satisfies

$$\mathbf{D}_{ii}^2 = 1, \quad \forall i \in [n],$$

the gradient flow initialized at $\mathbf{D}\mathbf{x}_0$ will converge to $\mathbf{D}\hat{\mathbf{x}}$.

Proof. By the results of Lemma 4, we obtain

$$\nabla f(\mathbf{D}\mathbf{x}) = \mathbf{D} \nabla f(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{R}^n. \quad (13)$$

Hence, we know that the gradient flow initialized with $\mathbf{D}\mathbf{x}_0$ is equal to $\mathbf{D}\mathbf{x}(t)$ at time t , for all $t \geq 0$. This leads to the conclusion that the new gradient flow will converge to $\mathbf{D}\hat{\mathbf{x}}$. □

Proof of Lemma 2. Since it is already known that the problem (1) has exponentially many global solutions, it remains to prove that the gradient flow with a radial random initialization will converge to one of the above global solutions with equal probability, i.e., with probability $2^{-\lceil n/2 \rceil}$. It has been proved in Lee et al. (2016) that the gradient flow will only converge to local minima if the objective function does not have degenerate saddle points, i.e., the Hessian of every saddle point has a negative curvature. Since the global solutions of the problem (1) are symmetric with respect to a radial probability distribution, it follows from Lemma 5 that we only need to prove that the objective function of this problem does not have degenerate saddle points. Equivalently, we prove that all second-order critical points are global minima.

Suppose that $\hat{\mathbf{x}}$ is a second-order critical point of the problem (1) that is not a global minimizer. Let $\mathbf{M} := \mathbf{x}\mathbf{x}^T$. Due to the symmetry of the landscape, we can assume without loss of generality that

$$\hat{x}_k \geq 0, \quad k = 1, \dots, n.$$

We define the direction $\Delta \in \mathbb{R}^n$ as

$$\Delta_{2k+1} = \hat{x}_{2k+1} - 1, \quad k = 1, \dots, \lceil n/2 \rceil, \quad \Delta_{2k} = \hat{x}_{2k}, \quad k = 1, \dots, \lfloor n/2 \rfloor.$$

Then, Lemma 7 in Ge et al. (2017) implies that

$$\begin{aligned} \Delta : \nabla^2 f[\hat{\mathbf{x}}] : \Delta &= \Delta \Delta^T : \nabla^2 g[(\hat{\mathbf{x}}\hat{\mathbf{x}}^T - \mathbf{x}^* \mathbf{x}^*)_{\Omega}] : \Delta \Delta^T - 3(\mathbf{M} - \mathbf{M}^*) : \nabla^2 g[(\hat{\mathbf{x}}\hat{\mathbf{x}}^T - \mathbf{x}^* \mathbf{x}^*)_{\Omega}] : (\mathbf{M} - \mathbf{M}^*) \\ &\leq (1 + \delta) \|(\Delta \Delta^T)_{\Omega}\|_F^2 - 3(1 - \delta) \|(\mathbf{M} - \mathbf{M}^*)_{\Omega}\|_F^2, \end{aligned}$$

where the last inequality is from the assumption that $g(\cdot)$ satisfies the sparse RIP condition with $(\delta, 1)$. Combining with the second-order necessary optimality condition, we obtain

$$(1 + \delta) \|(\Delta \Delta^T)_\Omega\|_F^2 \geq 3(1 - \delta) \|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F^2. \quad (14)$$

Using the expression

$$\|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F^2 - \|(\Delta \Delta^T)_\Omega\|_F^2 = \sum_{k=1}^{\lceil n/2 \rceil} [4\hat{x}_{2k+1} \cdot (\hat{x}_{2k+1} - 1)^2 + \|\hat{\mathbf{x}}^e\|^2 (2\hat{x}_{2k+1} - 1)],$$

the inequality (14) can be written as

$$- \sum_{k=1}^{\lceil n/2 \rceil} [4\hat{x}_{2k+1} \cdot (\hat{x}_{2k+1} - 1)^2 + \|\hat{\mathbf{x}}^e\|^2 (2\hat{x}_{2k+1} - 1)] \geq \frac{2 - 4\delta}{1 + \delta} \|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F^2. \quad (15)$$

The above inequality gives that

$$\begin{aligned} \frac{2 - 4\delta}{1 + \delta} \|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F^2 &\leq - \sum_{k=1}^{\lceil n/2 \rceil} [4\hat{x}_{2k+1} \cdot (\hat{x}_{2k+1} - 1)^2 + \|\hat{\mathbf{x}}^e\|^2 (2\hat{x}_{2k+1} - 1)] \\ &\leq - \sum_{k=1}^{\lceil n/2 \rceil} [0 - \|\hat{\mathbf{x}}^e\|^2] = \lceil n/2 \rceil \cdot \|\hat{\mathbf{x}}^e\|^2 \leq (n + 1)/2 \cdot \|\hat{\mathbf{x}}^e\|^2 \\ &\leq (n + 1)/2 \cdot \|\hat{\mathbf{x}}^e\| \|\hat{\mathbf{x}}\| \leq \sqrt{2}(n + 1)\delta \|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F, \end{aligned}$$

where the second inequality is from the assumption that $\hat{x}_{2k+1} \geq 0$ and the second last inequality is from Lemma 3. The above inequality implies that

$$\|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F \leq \frac{\sqrt{2}(n + 1)\delta(1 + \delta)}{2 - 4\delta}.$$

Recalling the condition

$$n \geq 3, \quad \delta \leq \frac{1}{2n},$$

we obtain

$$\|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F \leq \frac{1}{2}.$$

Checking the diagonal entries of $\mathbf{M} - \mathbf{M}^* = \hat{\mathbf{x}}\hat{\mathbf{x}}^T - \mathbf{x}^*(\mathbf{x}^*)^T$, we have

$$|\hat{x}_{2k+1} - 1| \leq \frac{1}{2}, \quad k = 1, \dots, \lceil n/2 \rceil,$$

which gives

$$\hat{x}_{2k+1} \geq \frac{1}{2}, \quad k = 1, \dots, \lceil n/2 \rceil.$$

Applying this condition to inequality (15), the left-hand side of the inequality is non-positive while the right-hand side is non-negative, which implies that

$$\|(\mathbf{M} - \mathbf{M}^*)_\Omega\|_F = 0.$$

This contradicts the assumption that $\hat{\mathbf{x}}$ is not a global solution. Hence, we have completed the proof that all second-order critical points of the problem (1) are global minima.

Furthermore, using Lemma 5, we know that the region of attraction (ROA) of each global minimum is symmetrical. Since the randomly initialized gradient flow converges to a second-order critical point with probability 1 and all second-order critical points are global minima, the gradient flow with a radial random initialization will converge to each global minimum with equal probability. $\square$

C.2 Proof of Theorem 4

Proof. In the unperturbed case, the curvature of the Hessian at each global minimum is given by

$$\Delta : \nabla^2 f(\mathbf{x}) : \Delta \geq (1 - \delta) \|(\mathbf{x}\Delta^T + \Delta\mathbf{x}^T)_\Omega\|_F^2, \quad \forall \Delta \in \mathbb{R}^n, \mathbf{x} \in \mathcal{X}^*,$$

which has been proved to be positive in Section 2. Therefore, we know that the Hessian at each global solution is positive definite and global minima are asymptotically stable for the gradient flow. We choose $R > 0$ to be a large enough constant such that

$$\mathbb{P}[\|\mathbf{x}_0\| \leq R] \geq 1 - 2^{-\lceil n/2 \rceil},$$

where the probability is chosen with respect to the initialization distribution. We consider the level set

$$\mathcal{L}_R := \{\mathbf{x} \in \mathbb{R}^n \mid f(\mathbf{x}) \leq c_R\},$$

where $c_R := \max\{f(\mathbf{x}) \mid \|\mathbf{x}\| \leq R\}$. Since the function $f(\mathbf{x})$ is continuous and coercive, the level set $\mathcal{L}_R$ is compact and the gradient flow will not leave $\mathcal{L}_R$ if it is initialized inside it. In addition, it holds that

$$\mathbb{P}[\mathbf{x}_0 \in \mathcal{L}_R] \geq 1 - 2^{-\lceil n/2 \rceil}.$$

Conditioning on the event that $\mathbf{x}_0 \in \mathcal{L}_R$, Lemma 2 implies that the gradient flow will converge to each global minimum with the same probability. Let $\hat{\mathbf{x}} \in \mathcal{X}^*$ be an arbitrary global minimum and $\mathcal{R}_{\hat{\mathbf{x}}}$ be its ROA of the gradient flow on $\mathcal{L}_R$. Therefore, Lemma 2 implies that

$$\mathbb{P}[\mathbf{x}_0 \in \mathcal{R}_{\hat{\mathbf{x}}} \mid \mathbf{x}_0 \in \mathcal{L}_R] = 2^{-\lceil n/2 \rceil}.$$

By Theorem 4.17 in Khalil (2002), there exist a smooth positive definite function $V(\mathbf{x})$ and a continuous positive definite function $W(\mathbf{x})$ such that every level set of $V(\mathbf{x})$ is compact and

$$\begin{aligned} V(\mathbf{x}) &\rightarrow +\infty, \quad \forall \mathbf{x} \rightarrow \partial\mathcal{R}_{\hat{\mathbf{x}}}, \\ \left\langle \frac{dV(\mathbf{x})}{d\mathbf{x}}, -\nabla_{\mathbf{x}} f(\mathbf{x}) \right\rangle &\leq -W(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{R}_{\hat{\mathbf{x}}}, \end{aligned}$$

where $\mathbf{x} \rightarrow \partial\mathcal{R}_{\hat{\mathbf{x}}}$ means that the distance between $\mathbf{x}$ and $\partial\mathcal{R}_{\hat{\mathbf{x}}}$ goes to zero, and $\partial\mathcal{R}_{\hat{\mathbf{x}}}$ denotes the boundary of the region of attraction of the solution $\hat{\mathbf{x}}$. We choose a large enough constant M such that

$$\mathbb{P}[\mathbf{x}_0 \in \mathcal{V}_M \mid \mathbf{x}_0 \in \mathcal{R}_{\hat{\mathbf{x}}}] \geq 1 - 2^{-\lceil n/2 \rceil},$$

where we define the level set $\mathcal{V}_M := \{\mathbf{x} \in \mathbb{R}^n \mid V(\mathbf{x}) \leq M\}$. Since the level set $\mathcal{V}_M$ is compact, there exists a small enough constant ϵ_0 such that

$$W(\mathbf{x}) \geq \epsilon_0, \quad \forall \mathbf{x} \in \mathcal{V}_M.$$

Now, we consider the perturbed case. We denote the new objective function as $\tilde{f}(\mathbf{x}; \eta)$, where $\eta \in \mathbb{R}$ is the perturbation to the global solution. More explicitly, the objective function is defined as

$$\tilde{f}(\mathbf{x}; \eta) := \|(\mathbf{x}\mathbf{x}^T - (\mathbf{x}^* + \eta)(\mathbf{x}^* + \eta)^T)_\Omega\|_F^2.$$

It has been proved in Section 2 that the global minimum of the perturbed problem is unique up to a sign flip if the perturbation is sufficiently small and generic, and that there exist $2^{\lceil n/2 \rceil} - 2$ spurious local minima. Since the gradient of $\tilde{f}(\mathbf{x}; \eta)$ is a uniformly continuous function of η on the compact set $\mathcal{V}_M$, there exists a small enough $r > 0$ such that for any generic η_0 satisfying $\|\eta_0\| \leq r$, it holds that

$$\left\langle \frac{dV(\mathbf{x})}{d\mathbf{x}}, -\nabla \tilde{f}(\mathbf{x}; \eta_0) \right\rangle \leq -\epsilon_0/2 < 0, \quad \forall \mathbf{x} \in \mathcal{V}_M.$$

This implies that the gradient flow on the perturbed problem will not leave $\mathcal{V}_M$ and will converge to a local minimum inside $\mathcal{V}_M$ if $\mathbf{x}_0 \in \mathcal{V}_M$. Therefore, if the initial point $\mathbf{x}_0$ is initialized with the given distribution, we have

$$\mathbb{P}\left[\lim_{t \rightarrow +\infty} \mathbf{x}(t) \in \mathcal{V}_M\right] \geq 2^{-\lceil n/2 \rceil} \left(1 - 2^{-\lceil n/2 \rceil}\right)^2 \geq 2^{-\lceil n/2 \rceil} \left(1 - 2^{-\lceil n/2 \rceil + 1}\right).$$

By choosing r to be the minimum over all points $\hat{\mathbf{x}} \in \mathcal{X}^*$, the gradient flow on the perturbed problem will converge to a spurious minimum with probability at least

$$\left(2^{\lceil n/2 \rceil} - 2\right) \cdot 2^{-\lceil n/2 \rceil} \left(1 - 2^{-\lceil n/2 \rceil + 1}\right) = 1 - \mathcal{O}\left(2^{-\lceil n/2 \rceil}\right).$$

Thus, we can conclude that the gradient flow on the perturbed problem will fail with probability at least $1 - \mathcal{O}(2^{-\lceil n/2 \rceil})$.

□