

EXACT LINEAR CONVERGENCE RATE ANALYSIS FOR LOW-RANK SYMMETRIC MATRIX COMPLETION VIA GRADIENT DESCENT

Trung Vu and Raviv Raich

School of EECS, Oregon State University, Corvallis, OR 97331-5501, USA
{vutru, raich}@oregonstate.edu

ABSTRACT

Factorization-based gradient descent is a scalable and efficient algorithm for solving low-rank matrix completion. Recent progress in structured non-convex optimization has offered global convergence guarantees for gradient descent under certain statistical assumptions on the low-rank matrix and the sampling set. However, while the theory suggests gradient descent enjoys fast linear convergence to a global solution of the problem, the universal nature of the bounding technique prevents it from obtaining an accurate estimate of the rate of convergence. In this paper, we perform a local analysis of the exact linear convergence rate of gradient descent for factorization-based matrix completion for symmetric matrices. Without any additional assumptions on the underlying model, we identify the deterministic condition for local convergence of gradient descent, which only depends on the solution matrix and the sampling set. More crucially, our analysis provides a closed-form expression of the asymptotic rate of convergence that matches exactly with the linear convergence observed in practice. To the best of our knowledge, our result is the first one that offers the exact rate of convergence of gradient descent for matrix factorization in Euclidean space for matrix completion.

Index Terms— Low-rank matrix completion, matrix factorization, local convergence analysis, gradient descent.

1. INTRODUCTION

Matrix completion is the problem of recovering a low-rank matrix from a sampling of its entries. In machine learning and signal processing, it has a wide range of applications including collaborative filtering [1], system identification [2] and dimension reduction [3]. In the era of big data, matrix completion has been proven to be an efficient and powerful framework to handle the enormous amount of information by exploiting low-rank structure of the data matrix.

Let $\mathbf{M} \in \mathbb{R}^{n \times m}$ be a rank r matrix with $1 \leq r \leq \min(n, m)$, and $\Omega = \{(i, j) \mid M_{ij} \text{ is observed}\}$ be an index subset of cardinality s such that $s \leq nm$. The goal is to recover the unknown entries of \mathbf{M} . Matrix completion can be formulated as a linearly constrained rank minimization or a rank-constrained least squares problem [4]. Two popular approaches for solving the aforementioned matrix completion problem formulations are convex relaxation via nuclear norm and non-convex factorization. The former approach, motivated by the success of compressed sensing, replaces the matrix rank by its convex surrogate (the nuclear norm). Extensive work on designing convex optimization algorithms with guarantees can be found in [4–8]. While on the theoretical side, the solutions of the relaxed problem and the original problem can be shown to coincide with high probability, on the practical side, computational complexity concerns diminish the applicability of these algorithms. When

the size of the matrix grows rapidly, storing and optimizing over a matrix variable become computationally expensive and even infeasible. In addition, it is evident this approach suffers from slow convergence [9, 10]. In the second approach, the original rank-constrained optimization is studied. Interestingly, by reparametrizing the $n \times m$ matrix as the product of two smaller matrices $\mathbf{M} = \mathbf{X}\mathbf{Y}^\top$, for $\mathbf{X} \in \mathbb{R}^{n \times r}$ and $\mathbf{Y} \in \mathbb{R}^{m \times r}$, the resulting equivalent problem is unconstrained and more computationally efficient to solve [11]. While this problem is non-convex, recent progress shows that for such problem any local minimum is also a global minimum [12, 13]. Thus, basic optimization algorithms such as gradient descent [12, 14, 15] and alternating minimization [16–19] can provably solve matrix completion under a specific sampling regime. Alternatively, the original rank-constrained optimization problem can be solved without the aforementioned reparameterization via the truncated singular value decomposition. [10, 20–25].

Among the aforementioned algorithms, let us draw our attention to the gradient descent method due to its outstanding simplicity and scalability. The first global convergence guarantee is attributed to Sun and Luo [12]. The authors proved that gradient descent with appropriate regularization can converge to the global optima of a factorization-based formulation at a linear rate. Later on, Ma *et al.* [15] proposed that even in the absence of explicit regularization, gradient descent recovers the underlying low-rank matrix by implicitly regularizing its iterates. The aforementioned results, while establishing powerful guarantees on the convergence behavior of gradient descent, impose several limitations. For some methods, the linear convergence rate depends on constants that are not in closed-form and are hard to verify in numerical experiments even when the underlying matrix is known. Second, a solution-independent analysis of the convergence rate typically offers a loose bound when considered for a specific solution. Third, the global convergence analysis requires certain assumptions on the underlying model which largely restrict the setting of the matrix completion problem in practice. Among such assumptions, one would consider the incoherence of the target matrix, the randomness of the sampling set, and the fact that the rank r and the condition number of \mathbf{M} are small constants as $n, m \rightarrow \infty$.

To address these issues, we consider the local convergence analysis of gradient descent for factorization-based matrix completion. In the scope of this paper, we restrict our attention to the symmetric case. We identify the condition for linear convergence of gradient descent that depends only on the solution \mathbf{M} and the sampling set Ω . In addition, we provide a closed-form expression for the asymptotic convergence rate that matches well with the convergence of the algorithm in practice. The proposed analysis does not require an asymptotic setting for matrix completion, e.g., large matrices of small rank. We believe that our analysis can be useful in both theoretical and practical aspects of the matrix completion problem.

Algorithm 1 (Non-convex) Gradient Descent

Input: $\mathbf{X}^0, \mathcal{P}_\Omega(\mathbf{M}), \eta$ **Output:** $\{\mathbf{X}^k\}$ 1: **for** $k = 0, 1, 2, \dots$ **do**2: $\mathbf{X}^{k+1} = \mathbf{X}^k - \eta \mathcal{P}_\Omega(\mathbf{X}^k \mathbf{X}^{k\top} - \mathbf{M}) \mathbf{X}^k$

2. GRADIENT DESCENT FOR MATRIX COMPLETION

Notations. Throughout the paper, we use the notations $\|\cdot\|_F$ and $\|\cdot\|_2$ to denote the Frobenius norm and the spectral norm of a matrix, respectively. On the other hand, $\|\cdot\|_2$ is used on a vector to denote the Euclidean norm. Boldfaced symbols are reserved for vectors and matrices. In addition, the $t \times t$ identity matrix is denoted by \mathbf{I}_t . \otimes denotes the Kronecker product between two matrices, and $\text{vec}(\cdot)$ denotes the vectorization of a matrix by stacking its columns on top of one another. Let \mathbf{X} be some matrix and $\mathbf{F}(\mathbf{X})$ be a matrix-valued function of \mathbf{X} . Then, for some positive number k , we use $\mathbf{F}(\mathbf{X}) = \mathcal{O}(\|\mathbf{X}\|_F^k)$ to imply $\lim_{\delta \rightarrow 0} \sup_{\|\mathbf{X}\|_F = \delta} \|\mathbf{F}(\mathbf{X})\|_F / \|\mathbf{X}\|_F^k < \infty$.

We begin by introducing the low-rank matrix completion problem. For simplicity, we focus on the *symmetric* case where \mathbf{M} is an $n \times n$ positive semi-definite (PSD) rank- r matrix and the sampling set Ω is symmetric.¹ Let the rank- r economy version of the eigendecomposition of \mathbf{M} be given by

$$\mathbf{M} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top,$$

where $\mathbf{U} \in \mathbb{R}^{n \times r}$ is a semi-orthogonal matrix and $\mathbf{\Lambda} \in \mathbb{R}^{r \times r}$ is a diagonal matrix containing r non-zero eigenvalues of \mathbf{M} , i.e., $\lambda_1 \geq \dots \geq \lambda_r > 0$. Since \mathbf{M} can be represented as

$$\mathbf{M} = (\mathbf{U} \mathbf{\Lambda}^{1/2})(\mathbf{U} \mathbf{\Lambda}^{1/2})^\top,$$

we can write $\mathbf{M} = \mathbf{X}^* \mathbf{X}^{*\top}$, such that $\mathbf{X}^* = \mathbf{U} \mathbf{\Lambda}^{1/2} \in \mathbb{R}^{n \times r}$. Therefore, the factorization-based formulation for matrix completion can be described using the following non-convex optimization:

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times r}} \frac{1}{4} \sum_{(i,j) \in \Omega} ([\mathbf{X} \mathbf{X}^\top]_{ij} - M_{ij})^2. \quad (1)$$

Denote $\mathcal{P}_\Omega : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ the projection onto the set of matrices supported in Ω , i.e.,

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

We can rewrite the objective function as $f(\mathbf{X}) = \frac{1}{4} \|\mathcal{P}_\Omega(\mathbf{X} \mathbf{X}^\top - \mathbf{M})\|_F^2$. The gradient of $f(\mathbf{X})$ is given by

$$\nabla f(\mathbf{X}) = \mathcal{P}_\Omega(\mathbf{X} \mathbf{X}^\top - \mathbf{M}) \mathbf{X}. \quad (2)$$

Starting from an initial \mathbf{X}^0 (usually through spectral initialization [15]), the gradient descent algorithm (see Algorithm 1) simply updates the value of \mathbf{X} by taking steps proportional to the negative of the gradient $\nabla f(\mathbf{X})$.

¹If the sampling set is not symmetric, one can symmetrize it by adding (j, i) , for any $(i, j) \in \Omega$, to Ω since $M_{ji} = M_{ij}$.

3. LOCAL CONVERGENCE ANALYSIS

This section presents the local convergence result of Algorithm 1. While recent work on the global guarantees of the algorithm has shown the linear behavior under certain statistical models, we emphasize that no closed-form expression of the convergence rate was provided. Our result in this paper, on the other hand, does not make any assumption about the underlying model for \mathbf{M} and Ω , and provides an exact expression of the asymptotic rate of convergence. Let us first introduce some critical concepts used in our derivation.

Definition 1. Denote $\bar{\Omega} = \{(i-1)n + j \mid (i, j) \in \Omega\}$. The row selection matrix \mathbf{S} is an $s \times n^2$ matrix obtained from a subset of rows corresponding to the elements in $\bar{\Omega}$ from the $n^2 \times n^2$ identity matrix \mathbf{I}_{n^2} .

Definition 2. The orthogonal projection onto the null space of \mathbf{M} is defined by $\mathbf{P}_{\mathbf{U}^\perp} = \mathbf{I}_n - \mathbf{U} \mathbf{U}^\top$.

Definition 3. Let \mathbf{T}_{n^2} be an $n^2 \times n^2$ matrix where the (i, j) th block of \mathbf{T}_{n^2} is the $n \times n$ matrix $\mathbf{e}_j \mathbf{e}_i^\top$ for $1 \leq i, j \leq n$. Then \mathbf{T}_{n^2} can be used to represent the transpose operator as follows:

$$\text{vec}(\mathbf{E}^\top) = \mathbf{T}_{n^2} \text{vec}(\mathbf{E}) \quad \text{for any } \mathbf{E} \in \mathbb{R}^{n \times n}.$$

We are now in position to state our main result on the asymptotic linear convergence rate of Algorithm 1.

Theorem 1. Denote $\mathbf{P}_1 = \mathbf{I}_{n^2} - \mathbf{P}_{\mathbf{U}^\perp} \otimes \mathbf{P}_{\mathbf{U}^\perp}$, $\mathbf{P}_2 = \frac{1}{2}(\mathbf{I}_{n^2} + \mathbf{T}_{n^2})$, and $\mathbf{P} = \mathbf{P}_1 \mathbf{P}_2$. In addition, let

$$\mathbf{H} = \mathbf{P} \left(\mathbf{I}_{n^2} - \eta(\mathbf{M} \oplus \mathbf{M})(\mathbf{S}^\top \mathbf{S}) \right) \mathbf{P}, \quad (3)$$

where $\mathbf{M} \oplus \mathbf{M} = \mathbf{M} \otimes \mathbf{I}_n + \mathbf{I}_n \otimes \mathbf{M}$ is the Kronecker sum. Define the spectral radius of \mathbf{H} , $\rho(\mathbf{H})$, as the largest absolute value of the eigenvalues of \mathbf{H} . If $\rho(\mathbf{H}) < 1$, then Algorithm 1 produces a sequence of matrices $\mathbf{X}^k \mathbf{X}^{k\top}$ converging to \mathbf{M} at an asymptotic linear rate $\rho(\mathbf{H})$. Formally, there exists a neighborhood $\mathcal{N}(\mathbf{M})$ of \mathbf{M} such that for any $\mathbf{X}^0 \mathbf{X}^{0\top} \in \mathcal{N}(\mathbf{M})$,

$$\|\mathbf{X}^k \mathbf{X}^{k\top} - \mathbf{M}\|_F \leq C \|\mathbf{X}^0 \mathbf{X}^{0\top} - \mathbf{M}\|_F \rho(\mathbf{H})^k, \quad (4)$$

for some numerical constant $C > 0$.

Remark 1. Theorem 1 provides a closed-form expression of the asymptotic linear convergence rate of Algorithm 1, which only depends on \mathbf{M} , Ω and the choice of step-size η . We note that the condition for linear convergence, $\rho(\mathbf{H}) < 1$, is fully determined given \mathbf{M} , Ω , and η . It would be interesting to establish a connection between this condition and the standard statistical model for matrix completion. For instance, how the incoherence of \mathbf{M} and the randomness of Ω would affect the spectral radius of \mathbf{H} ? This exploration is left as future work.

In our approach, the following lemma plays a pivotal role in the derivation of Theorem 1, establishing the recursion on the error matrix $\mathbf{X}^k \mathbf{X}^{k\top} - \mathbf{M}$:²

Lemma 1. Let $\mathbf{E}^k = \mathbf{X}^k \mathbf{X}^{k\top} - \mathbf{M}$. Then

$$\mathbf{E}^{k+1} = \mathbf{E}^k - \eta(\mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{M} + \mathbf{M} \mathcal{P}_\Omega(\mathbf{E}^k)) + \mathcal{O}(\|\mathbf{E}^k\|_F^2).$$

Furthermore, denote $\mathbf{A} = \mathbf{I}_{n^2} - \eta(\mathbf{M} \oplus \mathbf{M})(\mathbf{S}^\top \mathbf{S})$ and $\mathbf{e}^k = \text{vec}(\mathbf{E}^k)$, the matrix recursion can be rewritten compactly as

$$\mathbf{e}^{k+1} = \mathbf{A} \mathbf{e}^k + \mathcal{O}(\|\mathbf{e}^k\|_2^2). \quad (5)$$

²We provide proofs of all the lemmas in the Appendix.




Fig. 1. Linear convergence of gradient descent for matrix completion. The decrease in the norm of error matrix \mathbf{E}^k through iterations is shown in the blue dashed line with triangle markers. The black solid line with square markers and the red dotted line with circle markers represent first-order approximations of the error using \mathbf{A} and \mathbf{H} , respectively. Finally, the green dash-dot line is the theoretical bound (up to a constant) given by $\|e^0\|_2 \rho(\mathbf{H})^k$. We use different markers, i.e., triangle versus circle, to better distinguish the blue line from the red line, respectively.

Remark 2. Figure 1 illustrates the effectiveness of the proposed bound on the asymptotic rate of convergence given by Theorem 1. In Fig. 1, the low-rank solution matrix \mathbf{M} is generated by taking the product of a 20×3 matrix \mathbf{X} and its transpose, where \mathbf{X} has i.i.d. normally distributed entries. The sampling set Ω is obtained by randomly selecting the entries of \mathbf{M} based on a Bernoulli model with probability 0.3. Next, we run the economy-SVD on \mathbf{M} to compute $\mathbf{X}^* = \mathbf{U}\mathbf{\Lambda}^{1/2}$. The initialization \mathbf{X}^0 is obtained by adding i.i.d. normally distributed noise with standard deviation $\sigma = 10^{-2}$ to the entries of \mathbf{X}^* . Then we run Algorithm 1 with \mathbf{X}^0 , $\mathcal{P}_\Omega(\mathbf{M})$, and $\eta = 0.5/\|\mathbf{M}\|_2$. It is noticeable from Fig. 1 that our theoretical bound $\|e^0\|_2 \rho(\mathbf{H})^k$ given by the green line predicts successfully the rate of decrease in $\|\mathbf{E}^k\|_F$, running parallel to the blue line as soon as $\|\mathbf{E}^k\|_F < 10^{-2}$. As far as the approximations are concerned, we compare the changes in the error modeled by $e^{k+1} = \mathbf{A}e^k$ and the error modeled by $e^{k+1} = \mathbf{H}e^k$. While the former (represented by $\|\mathbf{A}^k e^0\|_2$ in black) fails to approximate $\|\mathbf{E}^k\|_F$ for $\|\mathbf{E}^k\|_F < 10^{-2}$, the later (represented by $\|\mathbf{H}^k e^0\|_2$ in red) matches $\|\mathbf{E}^k\|_F$ surprisingly well.

In the rest of this section, we shall derive the proof of Theorem 1. First, we present a major challenge met by the traditional approach that uses (5) to characterize the convergence of the error towards zero. Next, we describe our proposed technique to overcome this difficulty. Finally, we show that our bounding technique recovers the exact rate of local convergence of Algorithm 1.

3.1. A challenge of establishing the error contraction

The stability of the nonlinear difference equation (5) is the key to analyze the convergence of Algorithm 1. In essence, linear convergence rate is obtained by the following lemma:

Lemma 2. (Rephrased from the supplemental material of [26]) Let $(a_n)_{n \in \mathbb{N}} \subset \mathbb{R}_+$ be the sequence defined by

$$a_{n+1} \leq \rho a_n + q a_n^2 \quad \text{for } n = 0, 1, 2, \dots,$$

where $0 \leq \rho < 1$ and $q \geq 0$. Then (a_n) converges to 0 if and only if $a_0 < \frac{1-\rho}{q}$. A simple linear convergence bound can be derived for $a_0 < \rho \frac{1-\rho}{q}$ in the form of

$$a_n \leq a_0 K \rho^n, \quad \text{for } K = \left(1 - \frac{a_0 q}{\rho(1-\rho)}\right)^{-1}.$$

In order to apply Lemma 2 to (5), one natural way is to perform the eigendecomposition $\mathbf{A} = \mathbf{Q}_A \mathbf{\Lambda}_A \mathbf{Q}_A^{-1}$, where \mathbf{Q}_A is the square matrix whose columns are n^2 eigenvectors of \mathbf{A} , and $\mathbf{\Lambda}_A$ is the diagonal matrix whose diagonal elements are the corresponding eigenvalues of \mathbf{A} . Then, left-multiplying both sides of (5) by \mathbf{Q}_A^{-1} yields

$$\mathbf{Q}_A^{-1} e^{k+1} = \mathbf{\Lambda}_A \mathbf{Q}_A^{-1} e^k + \mathcal{O}(\|e^k\|_2^2),$$

where \mathbf{Q}_A^{-1} does not affect the \mathcal{O} term since its norm is constant. Applying the triangle inequality³ to the last equation leads to

$$\|\mathbf{Q}_A^{-1} e^{k+1}\|_2 = \|\mathbf{\Lambda}_A \mathbf{Q}_A^{-1} e^k\|_2 + \mathcal{O}(\|e^k\|_2^2). \quad (6)$$

With the definition of the spectral radius of \mathbf{A} using the spectral norm of $\mathbf{\Lambda}_A$, we have

$$\rho(\mathbf{A}) = \|\mathbf{\Lambda}_A\|_2 = \sup \left\{ \frac{\|\mathbf{\Lambda}_A \tilde{e}\|_2}{\|\tilde{e}\|_2} : \tilde{e} \in \mathbb{R}^{n^2}, \tilde{e} \neq \mathbf{0} \right\}. \quad (7)$$

Now, using (7) and the fact that $\mathcal{O}(\|e^k\|_2^2) = \mathcal{O}(\|\mathbf{Q}_A^{-1} e^k\|_2^2)$, (6) can be upper-bounded by

$$\|\mathbf{Q}_A^{-1} e^{k+1}\|_2 \leq \rho(\mathbf{A}) \|\mathbf{Q}_A^{-1} e^k\|_2 + \mathcal{O}(\|\mathbf{Q}_A^{-1} e^k\|_2^2). \quad (8)$$

If $\rho(\mathbf{A}) < 1$, then by Lemma 2, the sequence $\|\mathbf{Q}_A^{-1} e^k\|_2$ converges to 0 linearly at rate $\rho(\mathbf{A})$. Unfortunately, one can verify that $\rho(\mathbf{A}) \geq 1$ by taking any vector $\mathbf{v} \in \mathbb{R}^{n^2}$ such that $v_i = 0$ for all $i \in \bar{\Omega}$. Since $\mathbf{A}\mathbf{v} = \mathbf{v}$, 1 must be an eigenvalue of \mathbf{A} .

The failure of the aforementioned bounding technique is it overlooks the fact that $\mathbf{E}^k = \mathbf{X}^k \mathbf{X}^{k\top} - \mathbf{M}$. By defining $\mathcal{E} = \{\mathbf{X} \mathbf{X}^\top - \mathbf{M} \mid \mathbf{X} \in \mathbb{R}^{n \times r}\}$ and $\tilde{\mathcal{E}}_A = \{\mathbf{Q}_A^{-1} \text{vec}(\mathbf{E}) \mid \mathbf{E} \in \mathcal{E}\}$, a tighter bound on $\|\mathbf{\Lambda}_A \mathbf{Q}_A^{-1} e^k\|_2 / \|\mathbf{Q}_A^{-1} e^k\|_2$ can be obtained by

$$\rho^\mathcal{E}(\mathbf{A}, \delta) = \sup \left\{ \frac{\|\mathbf{\Lambda}_A \tilde{e}\|_2}{\|\tilde{e}\|_2} : \tilde{e} \in \tilde{\mathcal{E}}_A, \tilde{e} \neq \mathbf{0}, \|\tilde{e}\|_2 \leq \delta \right\}, \quad (9)$$

for some constant $\delta > 0$. Taking into account the structure of \mathbf{E}^k , one would expect $\rho^\mathcal{E}(\mathbf{A}) = \lim_{\delta \rightarrow 0} \rho^\mathcal{E}(\mathbf{A}, \delta)$ is a more reliable estimate of the asymptotic rate of convergence for (5). Nonetheless, (9) is a non-trivial optimization problem that has no closed-form solution to the best of our knowledge.

³Given $a = b + c$, by triangle inequality, we have $\|a\| \leq \|b\| + \|c\|$ and $\|a\| \geq \|b\| - \|c\|$ (since $b = a + (-c)$ and hence $\|b\| \leq \|a\| + \|-c\| = \|a\| + \|c\|$ or $\|a\| \geq \|b\| - \|c\|$). Consequently, we can write $\|a\| - \|b\| \leq \|c\|$ and hence $\|a\| = \|b\| + \mathcal{O}(\|c\|)$.

3.2. Integrating structural constraints

To address the aforementioned issue, we propose to integrate the structural constraint on \mathbf{E}^k into the recursion (5). As we shall show in the next subsection, this integration enables the application of Lemma 2 to the new recursion in order to obtain a tight bound on the convergence rate. First, let us characterize the feasible set of error matrices \mathcal{E} as follows:

Lemma 3. $\mathbf{E} \in \mathcal{E}$ if and only if the following conditions hold simultaneously:

- (C1) $\mathcal{P}_r(\mathbf{M} + \mathbf{E}) = \mathbf{M} + \mathbf{E}$, where \mathcal{P}_r is the truncated singular value decomposition of order r [27].
- (C2) $\mathbf{E}^\top = \mathbf{E}$.
- (C3) $\mathbf{v}^\top(\mathbf{M} + \mathbf{E})\mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{R}^n$.

Our strategy is to integrate three conditions in Lemma 3 into the linear operator \mathbf{A} so that the resulting recursion will implicitly enforce \mathbf{E}^k to remain in \mathcal{E} . Specifically, for condition (C1), we linearize \mathcal{P}_r using the first-order perturbation analysis of the truncated singular value decomposition [28]. For condition (C2), we leverage the linearity of the transpose operator. Finally, while handling condition (C3) is non-trivial, it turns out that this condition can be ignored. In the following lemma, we introduce the linear projection that ensures the updated error \mathbf{E}^k remains near \mathcal{E} .

Lemma 4. Recall that $\mathbf{P}_1 = \mathbf{I}_{n^2} - \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp}$, $\mathbf{P}_2 = \frac{1}{2}(\mathbf{I}_{n^2} + \mathbf{T}_{n^2})$. Then, the following statements hold:

1. \mathbf{P}_1 corresponds to the orthogonal projection onto the tangent plane of the set of rank- r matrices at \mathbf{M} .
2. \mathbf{P}_2 corresponds to the orthogonal projection onto the space of symmetric matrices.
3. \mathbf{P}_1 and \mathbf{P}_2 commute, and $\mathbf{P} = \mathbf{P}_1\mathbf{P}_2 = \mathbf{P}_2\mathbf{P}_1$ is also an orthogonal projection.
4. For any $\mathbf{E} \in \mathcal{E}$, $\text{vec}(\mathbf{E}) = \mathbf{P} \text{vec}(\mathbf{E}) + \mathcal{O}(\|\mathbf{E}\|_F^2)$.

By Lemma 4-4, we have $\mathbf{e}^k = \mathbf{P}\mathbf{e}^k + \mathcal{O}(\|\mathbf{e}^k\|_2^2)$ for all k . Using this result with $k+1$ instead of k and replacing \mathbf{e}^{k+1} from (5) into the first term on the RHS, we have

$$\mathbf{e}^{k+1} = \mathbf{P}(\mathbf{A}\mathbf{e}^k + \mathcal{O}(\|\mathbf{e}^k\|_2^2)) + \mathcal{O}(\|\mathbf{e}^{k+1}\|_2^2).$$

Substituting $\mathbf{e}^k = \mathbf{P}\mathbf{e}^k + \mathcal{O}(\|\mathbf{e}^k\|_2^2)$ and using $\mathbf{e}^{k+1} = \mathcal{O}(\|\mathbf{e}^k\|_2)$, we obtain

$$\mathbf{e}^{k+1} = \mathbf{P}\mathbf{A}\mathbf{P}\mathbf{e}^k + \mathcal{O}(\|\mathbf{e}^k\|_2^2). \quad (10)$$

It can be seen from Lemma 4-1 and Lemma 4-2 that the projection \mathbf{P} enforces the error vector \mathbf{e}^k to lie in the space under conditions (C1) and (C2) in Lemma 3. Now replacing the definition $\mathbf{H} = \mathbf{P}\mathbf{A}\mathbf{P}$, (10) can be rewritten as

$$\mathbf{e}^{k+1} = \mathbf{H}\mathbf{e}^k + \mathcal{O}(\|\mathbf{e}^k\|_2^2). \quad (11)$$

Similar to the derivation with \mathbf{A} , let $\mathbf{H} = \mathbf{Q}_H\mathbf{\Lambda}_H\mathbf{Q}_H^{-1}$ be the eigendecomposition of \mathbf{H} and define $\tilde{\mathbf{e}}^k = \mathbf{Q}_H^{-1}\mathbf{e}^k$. Then, we have

$$\|\tilde{\mathbf{e}}^{k+1}\|_2 = \|\mathbf{\Lambda}_H\tilde{\mathbf{e}}^k\|_2 + \mathcal{O}(\|\tilde{\mathbf{e}}^k\|_2^2). \quad (12)$$

In addition, denote $\tilde{\mathcal{E}}_H = \{\mathbf{Q}_H^{-1}\text{vec}(\mathbf{E}) \mid \mathbf{E} \in \mathcal{E}\}$, we can define

$$\rho(\mathbf{H}) = \sup \left\{ \frac{\|\mathbf{\Lambda}_H\tilde{\mathbf{e}}\|_2}{\|\tilde{\mathbf{e}}\|_2} : \tilde{\mathbf{e}} \in \mathbb{R}^{n^2}, \tilde{\mathbf{e}} \neq \mathbf{0} \right\} \text{ and} \quad (13)$$

$$\rho^\mathcal{E}(\mathbf{H}, \delta) = \sup \left\{ \frac{\|\mathbf{\Lambda}_H\tilde{\mathbf{e}}\|_2}{\|\tilde{\mathbf{e}}\|_2} : \tilde{\mathbf{e}} \in \tilde{\mathcal{E}}_H, \tilde{\mathbf{e}} \neq \mathbf{0}, \|\tilde{\mathbf{e}}\|_2 \leq \delta \right\}. \quad (14)$$

Since (5) and (11) are two different systems that describes the same dynamic for $\mathbf{E}^k \in \mathcal{E}$, one would expect they share the same asymptotic behavior. In particular, their linear rates of convergence should agree when the constraint $\mathbf{E}^k \in \mathcal{E}$ is considered.

Lemma 5. Let $\rho^\mathcal{E}(\mathbf{H}) = \lim_{\delta \rightarrow 0} \rho^\mathcal{E}(\mathbf{H}, \delta)$. Then,

$$\rho^\mathcal{E}(\mathbf{H}) = \rho^\mathcal{E}(\mathbf{A}).$$

While using \mathbf{H} instead of \mathbf{A} preserves the system dynamic over \mathcal{E} , it provides updates of the error that ensure that it remains in \mathcal{E} . Consequently, we can ignore the constraints that are implicitly satisfied in our analysis when using \mathbf{H} .

3.3. Asymptotic bound on the linear convergence rate

We have seen in Subsection 3.1 that applying Lemma 2 to (5) fails to estimate the convergence rate due to the gap between $\rho^\mathcal{E}(\mathbf{A})$ and $\rho(\mathbf{A})$. In this subsection, we show that integrating the structural constraint helps eliminating the gap between $\rho^\mathcal{E}(\mathbf{H})$ and $\rho(\mathbf{H})$ (even when condition (C3) is omitted). Therefore, applying Lemma 2 to (12) yields $\rho(\mathbf{H})$ as a tight bound on the convergence rate. To that end, our goal is to prove the following lemma:

Lemma 6. As δ approaches 0, we have $\rho(\mathbf{H}) - \rho^\mathcal{E}(\mathbf{H}, \delta) = \mathcal{O}(\delta)$. Consequently, it holds that $\rho(\mathbf{H}) = \rho^\mathcal{E}(\mathbf{H})$.

Let us briefly present the key ideas and lemmas we use to prove Lemma 6. Our proof relies on two critical considerations: (i) $\rho^\mathcal{E}(\mathbf{H}, \delta) \leq \rho(\mathbf{H})$, (ii) there exists a maximizer $\tilde{\mathbf{e}}^*$ of the supremum in (13) such that the distance from $\tilde{\mathbf{e}}^*$ to $\tilde{\mathcal{E}}_H$ is $\mathcal{O}(\delta^2)$. While (i) is trivial from (13) and (14), (ii) is proven by introducing \mathcal{F}_δ as a surrogate for the set \mathcal{E} as follows:

Lemma 7. Denote the eigenvector of \mathbf{H} corresponding to the largest (in magnitude) eigenvalue by \mathbf{q}_1 . Define \mathbf{G} as the $n \times n$ matrix satisfying $\text{vec}(\mathbf{G}) = \delta\mathbf{q}_1$. Let \mathcal{F}_δ be the set of $n \times n$ matrices satisfying the following conditions: (i) $\|\mathbf{F}\|_F \leq 2\delta$; (ii) $\mathbf{F}^\top = \mathbf{F}$; (iii) $\|\mathbf{P}_{U_\perp}\mathbf{F}\mathbf{P}_{U_\perp}\|_F \leq \frac{2}{\lambda_r}\delta^2$; and (iv) $\mathbf{v}^\top(\mathbf{M} + \mathbf{F})\mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{R}^n$. Then, there exists $\mathbf{F} \in \mathcal{F}_\delta$ satisfying

$$\|\mathbf{F} - \mathbf{G}\|_F = \mathcal{O}(\delta^2).$$

Lemma 8. For any $\mathbf{F} \in \mathcal{F}_\delta$, there exists $\mathbf{E} \in \mathcal{E}$ satisfying

$$\|\mathbf{E} - \mathbf{F}\|_F = \mathcal{O}(\delta^2).$$

From (i) and (ii), it follows that the difference between $\rho^\mathcal{E}(\mathbf{H}, \delta)$ and $\rho(\mathbf{H})$ is $\mathcal{O}(\delta)$. Thus, $\rho(\mathbf{H}) = \rho^\mathcal{E}(\mathbf{H})$ when taking the limit of $\rho^\mathcal{E}(\mathbf{H}, \delta)$ as $\delta \rightarrow 0$. Our derivation of Theorem 1 is completed by directly applying Lemma 2 to (12).

4. CONCLUSION AND FUTURE WORK

We presented a framework for analyzing the convergence of the existing gradient descent approach for low-rank matrix completion. In our analysis, we restricted our focus to the symmetric matrix completion case. We proved that the algorithm converges linearly. Different to other approaches, we made no assumption on the rank of the matrix or fraction of available entries. Instead, we derived an expression for the linear convergence rate via the spectral norm of a closed-form matrix. As future work, using random matrix theory, the closed-form expression for the convergence rate can be further related to the rank, the number of available entries, and the matrix dimensions. Additionally, this work can be extended to the non-symmetric case.

5. APPENDIX

5.1. Proof of Lemma 1

Recall the gradient descent update in Algorithm 1:

$$\begin{aligned} \mathbf{X}^{k+1} &= \mathbf{X}^k - \eta \mathcal{P}_\Omega(\mathbf{X}^k \mathbf{X}^{k\top} - \mathbf{M}) \mathbf{X}^k \\ &= (\mathbf{I}_n - \eta \mathcal{P}_\Omega(\mathbf{E}^k)) \mathbf{X}^k. \end{aligned} \quad (15)$$

Substituting (15) into the definition of \mathbf{E}^{k+1} , we have

$$\begin{aligned} \mathbf{E}^{k+1} &= \mathbf{X}^{k+1} \mathbf{X}^{k+1\top} - \mathbf{M} \\ &= (\mathbf{I}_n - \eta \mathcal{P}_\Omega(\mathbf{E}^k)) \mathbf{X}^k \mathbf{X}^{k\top} (\mathbf{I}_n - \eta \mathcal{P}_\Omega(\mathbf{E}^k))^\top - \mathbf{M}. \end{aligned}$$

From the fact that \mathbf{E}^k is symmetric and Ω is a symmetric sampling, the last equation can be further expanded as

$$\begin{aligned} \mathbf{E}^{k+1} &= \mathbf{X}^k \mathbf{X}^{k\top} - \eta \mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{X}^k \mathbf{X}^{k\top} \\ &\quad - \eta \mathbf{X}^k \mathbf{X}^{k\top} \mathcal{P}_\Omega(\mathbf{E}^k) + \eta^2 \mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{X}^k \mathbf{X}^{k\top} \mathcal{P}_\Omega(\mathbf{E}^k) - \mathbf{M}. \end{aligned} \quad (16)$$

Since $\mathbf{X}^k \mathbf{X}^{k\top} = \mathbf{M} + \mathbf{E}^k$, (16) is equivalent to

$$\begin{aligned} \mathbf{E}^{k+1} &= \mathbf{E}^k - \eta (\mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{M} + \mathbf{M} \mathcal{P}_\Omega(\mathbf{E}^k)) \\ &\quad - \eta (\mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{E}^k + \mathbf{E}^k \mathcal{P}_\Omega(\mathbf{E}^k)) \\ &\quad + \eta^2 \mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{M} \mathcal{P}_\Omega(\mathbf{E}^k) + \eta^2 \mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{E}^k \mathcal{P}_\Omega(\mathbf{E}^k). \end{aligned} \quad (17)$$

Note that $\|\mathcal{P}_\Omega(\mathbf{E}^k)\|_F \leq \|\mathbf{E}^k\|_F$. Hence, collecting terms that are of second order and higher, with respect to $\|\mathbf{E}^k\|_F$, on the RHS of (17) yields

$$\mathbf{E}^{k+1} = \mathbf{E}^k - \eta (\mathcal{P}_\Omega(\mathbf{E}^k) \mathbf{M} + \mathbf{M} \mathcal{P}_\Omega(\mathbf{E}^k)) + \mathcal{O}(\|\mathbf{E}^k\|_F^2).$$

Now by Definition 1, it is easy to verify that

$$\mathbf{S} \mathbf{S}^\top = \mathbf{I}_{n^2} \quad \text{and} \quad \text{vec}(\mathcal{P}_\Omega(\mathbf{E}^k)) = \mathbf{S}^\top \mathbf{S} \mathbf{e}^k.$$

Using the property $\text{vec}(\mathbf{A} \mathbf{B} \mathbf{C}) = (\mathbf{C}^\top \otimes \mathbf{A}) \text{vec}(\mathbf{B})$, (5) can be vectorized as follows:

$$\begin{aligned} \mathbf{e}^{k+1} &= \mathbf{e}^k - \eta (\mathbf{M} \otimes \mathbf{I}_n) \text{vec}(\mathcal{P}_\Omega(\mathbf{E}^k)) \\ &\quad - \eta (\mathbf{I}_n \otimes \mathbf{M}) \text{vec}(\mathcal{P}_\Omega(\mathbf{E}^k)) + \mathcal{O}(\|\mathbf{e}^k\|_2^2). \end{aligned}$$

The last equation can be reorganized as

$$\mathbf{e}^{k+1} = (\mathbf{I}_{n^2} - \eta (\mathbf{M} \oplus \mathbf{M}) (\mathbf{S}^\top \mathbf{S})) \mathbf{e}^k + \mathcal{O}(\|\mathbf{e}^k\|_2^2).$$

5.2. Proof of Lemma 3

(\Rightarrow) Suppose $\mathbf{E} \in \mathcal{E}$. Then for (C1), i.e., $\mathbf{E}^\top = \mathbf{E}$, $\mathbf{E} = \mathbf{X} \mathbf{X}^\top - \mathbf{M}$ is symmetric since both $\mathbf{X} \mathbf{X}^\top$ and \mathbf{M} are symmetric. For (C2), i.e., $\mathcal{P}_r(\mathbf{M} + \mathbf{E}) = \mathbf{M} + \mathbf{E}$, stems from the fact $\mathbf{M} + \mathbf{E} = \mathbf{X} \mathbf{X}^\top$ has rank no greater than r for $\mathbf{X} \in \mathbb{R}^{n \times r}$. Finally, for any $\mathbf{v} \in \mathbb{R}^n$, we have

$$\mathbf{v}^\top (\mathbf{M} + \mathbf{E}) \mathbf{v} = \mathbf{v}^\top (\mathbf{X} \mathbf{X}^\top) \mathbf{v} = \|\mathbf{X}^\top \mathbf{v}\|_2^2 \geq 0.$$

(\Leftarrow) From conditions (C1) and (C3), $\mathbf{M} + \mathbf{E}$ is a PSD matrix. In addition, $\mathcal{P}_r(\mathbf{M} + \mathbf{E}) = \mathbf{M} + \mathbf{E}$ implies $\mathbf{M} + \mathbf{E}$ must have rank no greater r . Since any PSD matrix \mathbf{A} with rank less than or equal to r can be factorized as $\mathbf{A} = \mathbf{Y} \mathbf{Y}^\top$ for some $\mathbf{Y} \in \mathbb{R}^{n \times r}$, we conclude that $\mathbf{E} \in \mathcal{E}$.

5.3. Proof of Lemma 4

First, recall that any matrix $\mathbf{\Pi} \in \mathbb{R}^{n^2 \times n^2}$ is an orthogonal projection if and only if $\mathbf{\Pi}^2 = \mathbf{\Pi}$ and $\mathbf{\Pi} = \mathbf{\Pi}^\top$. Since $\mathbf{P}_{U_\perp}^\top = \mathbf{P}_{U_\perp}$, we have

$$\begin{aligned} \mathbf{P}_1^\top &= (\mathbf{I}_{n^2} - \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp})^\top \\ &= \mathbf{I}_{n^2}^\top - \mathbf{P}_{U_\perp}^\top \otimes \mathbf{P}_{U_\perp}^\top \\ &= \mathbf{I}_{n^2} - \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp} = \mathbf{P}_1. \end{aligned}$$

In addition, since $\mathbf{P}_{U_\perp}^2 = \mathbf{P}_{U_\perp}$, we have

$$\begin{aligned} \mathbf{P}_1^2 &= (\mathbf{I}_{n^2} - \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp}) (\mathbf{I}_{n^2} - \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp})^\top \\ &= \mathbf{I}_{n^2}^2 - 2 \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp} + (\mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp})^2 \\ &= \mathbf{I}_{n^2} - 2 \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp} + (\mathbf{P}_{U_\perp}^2 \otimes \mathbf{P}_{U_\perp}^2) \\ &= \mathbf{I}_{n^2} - 2 \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp} + \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp} \\ &= \mathbf{I}_{n^2} - \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp} = \mathbf{P}_1. \end{aligned}$$

Second, using the fact that $\mathbf{T}_{n^2}^2 = \mathbf{I}_{n^2}$ and \mathbf{T}_{n^2} is symmetric, we can derive similar result:

$$\mathbf{P}_2^\top = \left(\frac{\mathbf{I}_{n^2} + \mathbf{T}_{n^2}}{2} \right)^\top = \frac{\mathbf{I}_{n^2} + \mathbf{T}_{n^2}}{2} = \mathbf{P}_2,$$

and

$$\begin{aligned} \mathbf{P}_2^2 &= \frac{(\mathbf{I}_{n^2} + \mathbf{T}_{n^2})^2}{4} \\ &= \frac{\mathbf{I}_{n^2}^2 + 2 \mathbf{T}_{n^2} + \mathbf{T}_{n^2}^2}{4} \\ &= \frac{2 \mathbf{I}_{n^2} + 2 \mathbf{T}_{n^2}}{4} \\ &= \frac{\mathbf{I}_{n^2} + \mathbf{T}_{n^2}}{2} = \mathbf{P}_2. \end{aligned}$$

Third, we observe that \mathbf{P}_1 and \mathbf{P}_2 are the vectorized version of the linear operators

$$\mathbf{\Pi}_1(\mathbf{E}) = \mathbf{E} - \mathbf{P}_{U_\perp} \mathbf{E} \mathbf{P}_{U_\perp}$$

and

$$\mathbf{\Pi}_2(\mathbf{E}) = \frac{1}{2}(\mathbf{E} + \mathbf{E}^\top),$$

respectively, for any $\mathbf{E} \in \mathbb{R}^{n \times n}$. Hence, in order to prove that \mathbf{P}_1 and \mathbf{P}_2 commute, it is sufficient to show that operators $\mathbf{\Pi}_1$ and $\mathbf{\Pi}_2$ commute. Indeed, we have

$$\begin{aligned} \mathbf{\Pi}_2 \mathbf{\Pi}_1(\mathbf{E}) &= \frac{1}{2} ((\mathbf{E} - \mathbf{P}_{U_\perp} \mathbf{E} \mathbf{P}_{U_\perp}) + (\mathbf{E} - \mathbf{P}_{U_\perp} \mathbf{E} \mathbf{P}_{U_\perp})^\top) \\ &= \frac{1}{2} (\mathbf{E} + \mathbf{E}^\top) - \mathbf{P}_{U_\perp} \frac{1}{2} (\mathbf{E} + \mathbf{E}^\top) \mathbf{P}_{U_\perp} \\ &= \mathbf{\Pi}_1 \mathbf{\Pi}_2(\mathbf{E}). \end{aligned}$$

This implies $\mathbf{\Pi}_1$ and $\mathbf{\Pi}_2$ commute. Since \mathbf{P} is the product of two commuting orthogonal projections, it is also an orthogonal projection.

Finally, let us restrict \mathbf{E} to belong to \mathcal{E} and denote $\mathbf{e} = \text{vec}(\mathbf{E})$. Using Theorem 3 in [28], we have

$$\mathcal{P}_r(\mathbf{M} + \mathbf{E}) = \mathbf{M} + \mathbf{E} - \mathbf{P}_{U_\perp} \mathbf{E} \mathbf{P}_{U_\perp} + \mathcal{O}(\|\mathbf{E}\|_F^2). \quad (18)$$

Since $\mathcal{P}_r(M + E) = M + E$, it follows from (18) that

$$P_{U_\perp} E P_{U_\perp} = \mathcal{O}(\|E\|_F^2).$$

Vectorizing the last equation, we obtain

$$(P_{U_\perp} \otimes P_{U_\perp})e = \mathcal{O}(\|E\|_F^2). \quad (19)$$

On the other hand, since E is symmetric,

$$e = T_{n^2} e = \left(\frac{I_{n^2} + T_{n^2}}{2} \right) e. \quad (20)$$

From (19) and (20), we have

$$\begin{aligned} e &= (I_{n^2} - P_{U_\perp} \otimes P_{U_\perp})e + \mathcal{O}(\|E\|_F^2) \\ &= (I_{n^2} - P_{U_\perp} \otimes P_{U_\perp}) \left(\frac{I_{n^2} + T_{n^2}}{2} \right) e + \mathcal{O}(\|E\|_F^2). \end{aligned} \quad (21)$$

Substituting

$$P = P_1 P_2 = (I_{n^2} - P_{U_\perp} \otimes P_{U_\perp}) \left(\frac{I_{n^2} + T_{n^2}}{2} \right)$$

into (21) completes our proof of the lemma.

5.4. Proof of Lemma 5

Let $\tilde{\mathcal{E}} = \{\text{vec}(E) \mid E \in \mathcal{E}\}$. Recall that for any $e \in \tilde{\mathcal{E}}$,

$$e = P e + \mathcal{O}(\|e\|_2^2).$$

Therefore, by the triangle inequality, we obtain

$$\begin{aligned} \|Ae\|_2 &= \|A(Pe + \mathcal{O}(\|e\|_2^2))\| \\ &\leq \|APe\|_2 + \|A\mathcal{O}(\|e\|_2^2)\|_2. \end{aligned}$$

Since the second term on the RHS of the last inequality is $\mathcal{O}(\|e\|_2^2)$, it is also $\mathcal{O}(\delta^2)$ for any $e \in \tilde{\mathcal{E}}$ such that $\|e\|_2 \leq \delta$. In other words,

$$\|Ae\|_2 \leq \|APe\|_2 + \mathcal{O}(\delta^2). \quad (22)$$

Similarly, we also have,

$$\begin{aligned} \|Ae\|_2 &\geq \|APe\|_2 - \|A\mathcal{O}(\|e\|_2^2)\|_2 \\ &= \|APe\|_2 - \mathcal{O}(\delta^2). \end{aligned} \quad (23)$$

From (22) and (23), it follows that

$$\frac{\|Ae\|_2}{\|e\|_2} = \frac{\|APe\|_2}{\|e\|_2} + \mathcal{O}(\delta). \quad (24)$$

Taking the limit of the supremum of (24) as $\delta \rightarrow 0$ yields

$$\begin{aligned} \rho^\mathcal{E}(A) &= \lim_{\delta \rightarrow 0} \sup_{\substack{e \in \tilde{\mathcal{E}} \\ e \neq 0 \\ \|e\|_2 \leq \delta}} \frac{\|Ae\|_2}{\|e\|_2} \\ &= \lim_{\delta \rightarrow 0} \sup_{\substack{e \in \tilde{\mathcal{E}} \\ e \neq 0 \\ \|e\|_2 \leq \delta}} \frac{\|APe\|_2}{\|e\|_2} = \rho^\mathcal{E}(AP). \end{aligned} \quad (25)$$

Now following similar argument in Lemma 6, we have

$$\begin{cases} \rho^\mathcal{E}(AP) = \rho(AP), \\ \rho^\mathcal{E}(PAP) = \rho(PAP). \end{cases} \quad (26)$$

Given (25) and (26), it remains to show that $\rho(AP) = \rho(PAP)$. Indeed, using Gelfand's formula [29], we have

$$\begin{aligned} \rho(AP) &= \lim_{k \rightarrow \infty} \|(AP)^k\|_2^{1/k} \\ \text{and } \rho(PAP) &= \lim_{k \rightarrow \infty} \|(PAP)^k\|_2^{1/k}. \end{aligned}$$

By the property of operator norms,

$$\|(AP)^k\|_2 = \|A(PAP)^{k-1}\|_2 \leq \|A\|_2 \|(PAP)^{k-1}\|_2.$$

Thus,

$$\|(AP)^k\|_2^{1/k} \leq \|A\|_2^{1/k} \left(\|(PAP)^{k-1}\|_2^{1/(k-1)} \right)^{(k-1)/k}.$$

Taking the limit of both sides of the last inequality as $k \rightarrow \infty$ yields $\rho(AP) \leq \rho(PAP)$. Similarly, since

$$\|(PAP)^k\|_2 = \|P(AP)^k\|_2 \leq \|(AP)^k\|_2,$$

we also obtain $\rho(PAP) \leq \rho(AP)$. This concludes our proof of the lemma.

5.5. Proof of Lemma 6

Without loss of generality, assume λ_1 is the eigenvalue with largest magnitude, i.e., $|\lambda_1| = \rho(H)$. By the definition of G , we have $\|G\|_F = \delta$. Since $H \text{vec}(G) = \lambda_1 \text{vec}(G)$ and $H = Q_H \Lambda_H Q_H^{-1}$, it follows that

$$Q_H \Lambda_H Q_H^{-1} \text{vec}(G) = \lambda_1 \text{vec}(G). \quad (27)$$

Multiplying both sides of (27) by Q_H^{-1} , we obtain

$$\Lambda_H Q_H^{-1} \text{vec}(G) = \lambda_1 Q_H^{-1} \text{vec}(G).$$

Taking the L_2 -norm and reorganizing the equation yields

$$\frac{\|\Lambda_H Q_H^{-1} \text{vec}(G)\|_2}{\|Q_H^{-1} \text{vec}(G)\|_2} = |\lambda_1| = \rho(H). \quad (28)$$

Therefore, G leads to a solution of the supremum in (13). We now prove that G is symmetric and $(P_{U_\perp} \otimes P_{U_\perp}) \text{vec}(G) = \mathbf{0}$. First, since P_1, P_2 and $P = P_1 P_2$ are orthogonal projections, we have

$$\begin{aligned} P_2 H &= P_2 P A P \\ &= P_2 P_2 P_1 A P \\ &= P_2 P_1 A P \\ &= P_1 P_2 A P \\ &= P A P = H. \end{aligned}$$

Thus,

$$\begin{aligned} \lambda_1 \text{vec}(G) &= H \text{vec}(G) \\ &= P_2 H \text{vec}(G) \\ &= \lambda_1 P_2 \text{vec}(G). \end{aligned} \quad (29)$$

Substituting $P_2 = \frac{1}{2}(I_{n^2} + T_{n^2})$ into (29) yields

$$\text{vec}(G^\top) = T_{n^2} \text{vec}(G) \quad \text{or} \quad G = G^\top.$$

Second, since $\mathbf{P}_1 \mathbf{H} = \mathbf{H}$, we obtain

$$\begin{aligned}\lambda_1 \text{vec}(\mathbf{G}) &= \mathbf{H} \text{vec}(\mathbf{G}) \\ &= \mathbf{P}_1 \mathbf{H} \text{vec}(\mathbf{G}) \\ &= \lambda_1 \mathbf{P}_1 \text{vec}(\mathbf{G}).\end{aligned}\quad (30)$$

Substituting $\mathbf{P}_1 = \mathbf{I}_{n^2} - \mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp}$ into (30) yields

$$(\mathbf{P}_{U_\perp} \otimes \mathbf{P}_{U_\perp}) \text{vec}(\mathbf{G}) = \mathbf{0} \quad \text{or} \quad \mathbf{P}_{U_\perp} \mathbf{G} \mathbf{P}_{U_\perp} = \mathbf{0}.$$

Since $\|\mathbf{E} - \mathbf{G}\|_F \leq \|\mathbf{E} - \mathbf{F}\|_F + \|\mathbf{F} - \mathbf{G}\|_F$ (by the triangle inequality), Lemmas 7 and 8 imply the existence of $\tilde{\mathbf{E}} \in \mathcal{E}$ such that $\|\mathbf{E} - \mathbf{G}\|_F = \mathcal{O}(\delta^2)$. Denote $\tilde{\mathbf{e}} = \mathbf{Q}_H^{-1} \text{vec}(\tilde{\mathbf{E}}) \in \tilde{\mathcal{E}}_H$, we have

$$\begin{aligned}\mathbf{\Lambda}_H \tilde{\mathbf{e}} &= \lambda_1 \tilde{\mathbf{e}} - (\lambda_1 \mathbf{I}_{n^2} - \mathbf{\Lambda}_H) \tilde{\mathbf{e}} \\ &= \lambda_1 \tilde{\mathbf{e}} - (\lambda_1 \mathbf{I}_{n^2} - \mathbf{\Lambda}_H) \mathbf{Q}_H^{-1} \text{vec}(\tilde{\mathbf{E}}) \\ &= \lambda_1 \tilde{\mathbf{e}} - (\lambda_1 \mathbf{I}_{n^2} - \mathbf{\Lambda}_H) \mathbf{Q}_H^{-1} \text{vec}(\mathbf{E} - \mathbf{G}).\end{aligned}$$

where the last equality stems from the fact that $\lambda_1 \mathbf{Q}_H^{-1} \text{vec}(\mathbf{G}) = \mathbf{\Lambda}_H \mathbf{Q}_H^{-1} \text{vec}(\mathbf{G})$. Next, using the triangle inequality, we obtain

$$\begin{aligned}\|\mathbf{\Lambda}_H \tilde{\mathbf{e}}\|_2 &\geq \|\lambda_1 \tilde{\mathbf{e}}\|_2 - \|(\lambda_1 \mathbf{I}_{n^2} - \mathbf{\Lambda}_H) \mathbf{Q}_H^{-1} \text{vec}(\mathbf{E} - \mathbf{G})\|_2 \\ &\geq \rho(\mathbf{H}) \|\tilde{\mathbf{e}}\|_2 - \|\lambda_1 \mathbf{I}_{n^2} - \mathbf{\Lambda}_H\|_2 \|\mathbf{Q}_H^{-1}\|_2 \|\text{vec}(\mathbf{E} - \mathbf{G})\|_2.\end{aligned}$$

Dividing both sides by $\|\tilde{\mathbf{e}}\|_2$ yields

$$\frac{\|\mathbf{\Lambda}_H \tilde{\mathbf{e}}\|_2}{\|\tilde{\mathbf{e}}\|_2} \geq \rho(\mathbf{H}) - \frac{\|\lambda_1 \mathbf{I}_{n^2} - \mathbf{\Lambda}_H\|_2 \|\mathbf{Q}_H^{-1}\|_2 \|\text{vec}(\mathbf{E} - \mathbf{G})\|_2}{\|\tilde{\mathbf{e}}\|_2}.\quad (31)$$

Since $\|\mathbf{E} - \mathbf{G}\|_F = \mathcal{O}(\delta^2)$, (31) can be rewritten as

$$\frac{\|\mathbf{\Lambda}_H \tilde{\mathbf{e}}\|_2}{\|\tilde{\mathbf{e}}\|_2} \geq \rho(\mathbf{H}) - \mathcal{O}(\delta^2).\quad (32)$$

On the other hand, for any $\tilde{\mathbf{e}} \in \tilde{\mathcal{E}}_H$, we also have

$$\frac{\|\mathbf{\Lambda}_H \tilde{\mathbf{e}}\|_2}{\|\tilde{\mathbf{e}}\|_2} \leq \rho^\mathcal{E}(\mathbf{H}, \delta) \leq \rho(\mathbf{H}).\quad (33)$$

Combining (32) and (33) yields $\rho(\mathbf{H}) - \rho^\mathcal{E}(\mathbf{H}, \delta) = \mathcal{O}(\delta)$.

5.6. Proof of Lemma 7

Denote $\mathbf{P}_U = \mathbf{U}\mathbf{U}^\top$, for any $\mathbf{v} \in \mathbb{R}^n$, we can decompose \mathbf{v} into two orthogonal component:

$$\mathbf{v} = \mathbf{v}_U + \mathbf{v}_\perp,$$

where $\mathbf{v}_U = \mathbf{P}_U \mathbf{v}$ and $\mathbf{v}_\perp = \mathbf{P}_{U_\perp} \mathbf{v}$. Without loss of generality, assume that $\|\mathbf{v}\|_2 = \|\mathbf{v}_U\|_2^2 + \|\mathbf{v}_\perp\|_2^2 = 1$. Thus, we have

$$\begin{aligned}\mathbf{v}^\top (\mathbf{M} + \mathbf{G}) \mathbf{v} &= (\mathbf{v}_U + \mathbf{v}_\perp)^\top (\mathbf{M} + \mathbf{G}) (\mathbf{v}_U + \mathbf{v}_\perp) \\ &= \mathbf{v}_U^\top \mathbf{M} \mathbf{v}_U + \mathbf{v}_U^\top \mathbf{G} \mathbf{v}_U + \mathbf{v}_U^\top \mathbf{G} \mathbf{v}_\perp \\ &\quad + \mathbf{v}_\perp^\top \mathbf{G} \mathbf{v}_U + \mathbf{v}_\perp^\top \mathbf{G} \mathbf{v}_\perp,\end{aligned}\quad (34)$$

where the last equation stems from the fact that $\mathbf{M} = \mathcal{P}_U \mathbf{M} \mathcal{P}_U$ and $\mathcal{P}_U \mathcal{P}_{U_\perp} = \mathbf{0}$. Since $\mathcal{P}_{U_\perp} \mathbf{G} \mathcal{P}_{U_\perp} = \mathbf{0}$, we have

$$\mathbf{v}_\perp^\top \mathbf{G} \mathbf{v}_\perp = \mathbf{v}^\top \mathcal{P}_{U_\perp} \mathbf{G} \mathcal{P}_{U_\perp} \mathbf{v} = 0.$$

Thus, (34) is equivalent to

$$\mathbf{v}^\top (\mathbf{M} + \mathbf{G}) \mathbf{v} = \mathbf{v}_U^\top \mathbf{M} \mathbf{v}_U + \mathbf{v}_U^\top \mathbf{G} \mathbf{v}_U + 2\mathbf{v}_U^\top \mathbf{G} \mathbf{v}_\perp.\quad (35)$$

Now let us lower-bound each term on the RHS of (35) as follows. First, by the Rayleigh quotient, we have

$$\mathbf{v}_U^\top \mathbf{M} \mathbf{v}_U \geq \lambda_r \|\mathbf{v}_U\|_2^2,\quad (36)$$

and

$$\mathbf{v}_U^\top \mathbf{G} \mathbf{v}_U \geq \lambda_{\min}(\mathbf{G}) \|\mathbf{v}_U\|_2^2 \geq -\|\mathbf{G}\|_F \|\mathbf{v}_U\|_2^2.\quad (37)$$

Next, by Cauchy-Schwarz inequality,

$$\mathbf{v}_U^\top \mathbf{G} \mathbf{v}_\perp \geq -\|\mathbf{G}\|_2 \|\mathbf{v}_U\|_2 \|\mathbf{v}_\perp\|_2 \geq -\|\mathbf{G}\|_F \|\mathbf{v}_U\|_2.\quad (38)$$

From (36), (37), and (38), we obtain

$$\mathbf{v}^\top (\mathbf{M} + \mathbf{G}) \mathbf{v} \geq (\lambda_r - \|\mathbf{G}\|_F) \|\mathbf{v}_U\|_2^2 - 2\|\mathbf{G}\|_F \|\mathbf{v}_U\|_2.\quad (39)$$

Note that $\|\mathbf{G}\|_F = \delta$ and the quadratic $g(t) = (\lambda_r - \delta)t^2 - 2\delta t$ is minimized at

$$t_* = \frac{\delta}{\lambda_r - \delta}, \quad g(t_*) = -\frac{\delta^2}{\lambda_r - \delta}.$$

Combining this with (39) yields

$$\mathbf{v}^\top (\mathbf{M} + \mathbf{G}) \mathbf{v} \geq -\frac{2}{\lambda_r} \delta^2,$$

for sufficiently small δ . Let $\mathbf{F} = \mathbf{G} + \frac{2}{\lambda_r} \delta^2 \mathbf{I}_n$. Now we can easily verify that $\|\mathbf{F} - \mathbf{G}\|_F = \mathcal{O}(\delta^2)$ and $\mathbf{F} \in \mathcal{F}$.

5.7. Proof of Lemma 8

We shall show that the matrix $\mathbf{E} = \mathcal{P}_r(\mathbf{M} + \mathbf{F}) - \mathbf{M}$ belongs to \mathcal{E} and satisfies

$$\|\mathbf{E} - \mathbf{F}\|_F = \mathcal{O}(\delta^2).\quad (40)$$

First, since $\mathbf{F} \in \mathcal{F}_\delta$, $\mathbf{M} + \mathbf{F}$ must be PSD. Thus, $\mathcal{P}_r(\mathbf{M} + \mathbf{F})$ is a PSD matrix of rank no greater than r and it admits a rank- r factorization $\mathcal{P}_r(\mathbf{M} + \mathbf{F}) = \mathbf{Z}\mathbf{Z}^\top$, for some $\mathbf{Z} \in \mathbb{R}^{n \times r}$. Therefore, by the definition of \mathcal{E} ,

$$\mathbf{E} = \mathcal{P}_r(\mathbf{M} + \mathbf{F}) - \mathbf{M} = \mathbf{Z}\mathbf{Z}^\top - \mathbf{M} \in \mathcal{E}.$$

Next, using (18), we have

$$\begin{aligned}\mathbf{E} - \mathbf{F} &= \mathcal{P}_r(\mathbf{M} + \mathbf{F}) - \mathbf{M} - \mathbf{F} \\ &= \mathcal{P}_{U_\perp} \mathbf{F} \mathcal{P}_{U_\perp} + \mathcal{O}(\|\mathbf{F}\|_F^2).\end{aligned}$$

Since $\mathbf{F} \in \mathcal{F}_\delta$ implies $\mathcal{P}_{U_\perp} \mathbf{F} \mathcal{P}_{U_\perp} = \mathcal{O}(\|\mathbf{F}\|_F^2)$, we conclude that $\mathbf{E} - \mathbf{F} = \mathcal{O}(\|\mathbf{F}\|_F^2)$.

6. REFERENCES

- [1] Jasson DM Rennie and Nathan Srebro, "Fast maximum margin matrix factorization for collaborative prediction," in *International Conference on Machine Learning*, 2005, pp. 713–719.
- [2] Zhang Liu and Lieven Vandenberghe, "Interior-point method for nuclear norm approximation with application to system identification," *SIAM Journal on Matrix Analysis and Applications*, vol. 31, no. 3, pp. 1235–1256, 2009.
- [3] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright, "Robust principal component analysis?," *Journal of the ACM*, vol. 58, no. 3, pp. 11, 2011.

- [4] Emmanuel J Candès and Benjamin Recht, “Exact matrix completion via convex optimization,” *Foundations of Computational mathematics*, vol. 9, no. 6, pp. 717, 2009.
- [5] Shuiwang Ji and Jieping Ye, “An accelerated gradient method for trace norm minimization,” in *International Conference on Machine Learning*, 2009, pp. 457–464.
- [6] Kim-Chuan Toh and Sangwoon Yun, “An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems,” *Pacific Journal of Optimization*, vol. 6, pp. 615–640, 2010.
- [7] Jian-Feng Cai, Emmanuel J Candès, and Zuowei Shen, “A singular value thresholding algorithm for matrix completion,” *SIAM Journal on Optimization*, vol. 20, no. 4, pp. 1956–1982, 2010.
- [8] Shiqian Ma, Donald Goldfarb, and Lifeng Chen, “Fixed point and Bregman iterative methods for matrix rank minimization,” *Mathematical Programming*, vol. 128, no. 1-2, pp. 321–353, 2011.
- [9] Yehuda Koren, “The BellKor solution to the Netflix grand prize,” *Netflix prize documentation*, vol. 81, no. 2009, pp. 1–10, 2009.
- [10] Trung Vu and Raviv Raich, “Accelerating iterative hard thresholding for low-rank matrix completion via adaptive restart,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 2917–2921.
- [11] Samuel Burer and Renato DC Monteiro, “Local minima and convergence in low-rank semidefinite programming,” *Mathematical Programming*, vol. 103, no. 3, pp. 427–444, 2005.
- [12] Ruoyu Sun and Zhi-Quan Luo, “Guaranteed matrix completion via non-convex factorization,” *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 6535–6579, 2016.
- [13] Rong Ge, Jason D Lee, and Tengyu Ma, “Matrix completion has no spurious local minimum,” in *Advances in Neural Information Processing Systems*, 2016, pp. 2973–2981.
- [14] Yudong Chen and Martin J Wainwright, “Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees,” *arXiv preprint arXiv:1509.03025*, 2015.
- [15] Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen, “Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval and matrix completion,” in *International Conference on Machine Learning*, 2018, pp. 3345–3354.
- [16] Caihua Chen, Bingsheng He, and Xiaoming Yuan, “Matrix completion via an alternating direction method,” *IMA Journal of Numerical Analysis*, vol. 32, no. 1, pp. 227–245, 2012.
- [17] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi, “Low-rank matrix completion using alternating minimization,” in *Annual ACM Symposium on Theory of Computing*, 2013, pp. 665–674.
- [18] Moritz Hardt, “Understanding alternating minimization for matrix completion,” in *IEEE Annual Symposium on Foundations of Computer Science*, 2014, pp. 651–660.
- [19] Moritz Hardt and Mary Wootters, “Fast matrix completion without the condition number,” in *Conference on Learning Theory*, 2014, pp. 638–678.
- [20] Prateek Jain, Raghu Meka, and Inderjit S Dhillon, “Guaranteed rank minimization via singular value projection,” in *Advances in Neural Information Processing Systems*, 2010, pp. 937–945.
- [21] Donald Goldfarb and Shiqian Ma, “Convergence of fixed-point continuation algorithms for matrix rank minimization,” *Foundations of Computational Mathematics*, vol. 11, no. 2, pp. 183–210, 2011.
- [22] Jared Tanner and Ke Wei, “Normalized iterative hard thresholding for matrix completion,” *SIAM Journal on Scientific Computing*, vol. 35, no. 5, pp. S104–S125, 2013.
- [23] Prateek Jain and Praneeth Netrapalli, “Fast exact matrix completion with finite samples,” in *Conference on Learning Theory*, 2015, pp. 1007–1034.
- [24] Evgenia Chunikhina, Raviv Raich, and Thanh Nguyen, “Performance analysis for matrix completion via iterative hard-thresholded SVD,” in *IEEE Workshop on Statistical Signal Processing*, 2014, pp. 392–395.
- [25] Trung Vu and Raviv Raich, “Local convergence of the Heavy Ball method in iterative hard thresholding for low-rank matrix completion,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 3417–3421.
- [26] Trung Vu and Raviv Raich, “Adaptive step size momentum method for deconvolution,” in *IEEE Statistical Signal Processing Workshop*, 2018, pp. 438–442.
- [27] Carl Eckart and Gale Young, “The approximation of one matrix by another of lower rank,” *Psychometrika*, vol. 1, no. 3, pp. 211–218, 1936.
- [28] Trung Vu, Evgenia Chunikhina, and Raviv Raich, “Perturbation expansions and error bounds for the truncated singular value decomposition,” *arXiv preprint arXiv:2009.07542*, 2020.
- [29] Izrail Gelfand, “Normierte ringe,” *Mathematics Sbornik*, vol. 9, no. 1, pp. 3–24, 1941.