

Provable Efficient Online Matrix Completion via Non-convex Stochastic Gradient Descent

Chi Jin* Sham M. Kakade† Praneeth Netrapalli‡

May 27, 2016

Abstract

Matrix completion, where we wish to recover a low rank matrix by observing a few entries from it, is a widely studied problem in both theory and practice with wide applications. Most of the provable algorithms so far on this problem have been restricted to the offline setting where they provide an estimate of the unknown matrix using all observations simultaneously. However, in many applications, the online version, where we observe one entry at a time and dynamically update our estimate, is more appealing. While existing algorithms are efficient for the offline setting, they could be highly inefficient for the online setting.

In this paper, we propose the first provable, efficient online algorithm for matrix completion. Our algorithm starts from an initial estimate of the matrix and then performs non-convex stochastic gradient descent (SGD). After every observation, it performs a fast update involving only one row of two tall matrices, giving near linear total runtime. Our algorithm can be naturally used in the offline setting as well, where it gives competitive sample complexity and runtime to state of the art algorithms. Our proofs introduce a general framework to show that SGD updates tend to stay away from saddle surfaces and could be of broader interests for other non-convex problems to prove tight rates.

1 Introduction

Low rank matrix completion refers to the problem of recovering a low rank matrix by observing the values of only a tiny fraction of its entries. This problem arises in several applications such as video denoising [14], phase retrieval [3] and most famously in movie recommendation engines [16]. In the context of recommendation engines for instance, the matrix we wish to recover would be user-item rating matrix where each row corresponds to a user and each column corresponds to an item. Each entry of the matrix is the rating given by a user to an item. Low rank assumption on the matrix is inspired by the intuition that rating of an item by a user depends on only a few hidden factors, which are much fewer than the number of users or items. The goal is to estimate the ratings of all items by users given only partial ratings of items by users, which would then be helpful in recommending new items to users.

The seminal works of Candès and Recht [4] first identified regularity conditions under which low rank matrix completion can be solved in polynomial time using convex relaxation – low rank matrix

*UC Berkeley. Email: chijin@cs.berkeley.edu

†University of Washington. Email: sham@cs.washington.edu

‡Microsoft Research New England. Email: praneeth@microsoft.com

completion could be ill-posed and NP-hard in general without such regularity assumptions [10]. Since then, a number of works have studied various algorithms under different settings for matrix completion: weighted and noisy matrix completion, fast convex solvers, fast iterative non-convex solvers, parallel and distributed algorithms and so on.

Most of this work however deals only with the offline setting where all the observed entries are revealed at once and the recovery procedure does computation using all these observations simultaneously. However in several applications [5, 19], we encounter the online setting where observations are only revealed sequentially and at each step the recovery algorithm is required to maintain an estimate of the low rank matrix based on the observations so far. Consider for instance recommendation engines, where the low rank matrix we are interested in is the user-item rating matrix. While we make an observation only when a user rates an item, at any point of time, we should have an estimate of the user-item rating matrix based on all prior observations so as to be able to continuously recommend items to users. Moreover, this estimate should get better as we observe more ratings.

Algorithms for offline matrix completion can be used to solve the online version by rerunning the algorithm after every additional observation. However, performing so much computation for every observation seems wasteful and is also impractical. For instance, using alternating minimization, which is among the fastest known algorithms for the offline problem, would mean that we take several passes of the entire data for every additional observation. This is simply not feasible in most settings. Another natural approach is to group observations into batches and do an update only once for each batch. This however induces a lag between observations and estimates which is undesirable. To the best of our knowledge, there is no known *provable, efficient, online algorithm* for matrix completion.

On the other hand, in order to deal with the online matrix completion scenario in practical applications, several heuristics (with no convergence guarantees) have been proposed in literature [2, 20]. Most of these approaches are based on starting with an estimate of the matrix and doing fast updates of this estimate whenever a new observation is presented. One of the update procedures used in this context is that of stochastic gradient descent (SGD) applied to the following non-convex optimization problem

$$\min_{\mathbf{U}, \mathbf{V}} \|\mathbf{M} - \mathbf{UV}^\top\|_{\text{F}}^2 \quad \text{s.t.} \quad \mathbf{U} \in \mathbb{R}^{d_1 \times k}, \mathbf{V} \in \mathbb{R}^{d_2 \times k}, \quad (1)$$

where \mathbf{M} is the unknown matrix of size $d_1 \times d_2$, k is the rank of \mathbf{M} and \mathbf{UV}^\top is a low rank factorization of \mathbf{M} we wish to obtain. The algorithm starts with some \mathbf{U}_0 and \mathbf{V}_0 , and given a new observation $(\mathbf{M})_{ij}$, SGD updates the i^{th} -row and the j^{th} -row of the current iterates \mathbf{U}_t and \mathbf{V}_t respectively by

$$\begin{aligned} \mathbf{U}_{t+1}^{(i)} &= \mathbf{U}_t^{(i)} - 2\eta d_1 d_2 \left(\mathbf{U}_t \mathbf{V}_t^\top - \mathbf{M} \right)_{ij} \mathbf{V}_t^{(j)}, \text{ and,} \\ \mathbf{V}_{t+1}^{(j)} &= \mathbf{V}_t^{(j)} - 2\eta d_1 d_2 \left(\mathbf{U}_t \mathbf{V}_t^\top - \mathbf{M} \right)_{ij} \mathbf{U}_t^{(i)}, \end{aligned} \quad (2)$$

where η is an appropriately chosen stepsize, and $\mathbf{U}^{(i)}$ denote the i^{th} row of matrix \mathbf{U} . Note that each update modifies only one row of the factor matrices \mathbf{U} and \mathbf{V} , and the computation only involves one row of \mathbf{U}, \mathbf{V} and the new observed entry $(\mathbf{M})_{ij}$ and hence are extremely fast. These fast updates make SGD extremely appealing in practice. Moreover, SGD, in the context of matrix completion, is also useful for parallelization and distributed implementation [24].

1.1 Our Contributions

In this work we present the first provable efficient algorithm for online matrix completion by showing that SGD (2) with a good initialization converges to a true factorization of \mathbf{M} at a geometric rate. Our main contributions are as follows.

- We provide the *first provable, efficient, online* algorithm for matrix completion. Starting with a good initialization, after each observation, the algorithm makes quick updates each taking time $O(k^3)$ and requires $O(\mu d k \kappa^4 (k + \log \frac{\|\mathbf{M}\|_{\text{F}}}{\epsilon}) \log d)$ observations to reach ϵ accuracy, where μ is the incoherence parameter, $d = \max(d_1, d_2)$, k is the rank and κ is the condition number of \mathbf{M} .
- Moreover, our result features both sample complexity and total runtime *linear in d* , and is *competitive* to even the best existing offline results for matrix completion. (either improve over or is incomparable, i.e., better in some parameters and worse in others, to these results). See Table 1 for the comparison.
- To obtain our results, we introduce *a general framework to show SGD updates tend to stay away from saddle surfaces*. In order to do so, we consider distances from saddle surfaces, show that they behave like sub-martingales under SGD updates and use martingale convergence techniques to conclude that the iterates stay away from saddle surfaces. While [25] shows that SGD updates stay away from saddle surfaces, the stepsizes they can handle are quite small (scaling as $1/\text{poly}(d_1, d_2)$), leading to suboptimal computational complexity. Our framework makes it possible to establish the same statement for much larger step sizes, giving us near-optimal runtime. We believe these techniques may be applicable in other non-convex settings as well.

1.2 Related Work

In this section we will mention some more related work.

Offline matrix completion: There has been a lot of work on designing offline algorithms for matrix completion, we provide the detailed comparison with our algorithm in Table 1. The nuclear norm relaxation algorithm [23] has near-optimal sample complexity for this problem but is computationally expensive. Motivated by the empirical success of non-convex heuristics, a long line of works, [15, 9, 13, 25] and so on, has obtained convergence guarantees for alternating minimization, gradient descent, projected gradient descent etc. Even the best of these are suboptimal in sample complexity by $\text{poly}(k, \kappa)$ factors. Our sample complexity is better than that of [15] and is incomparable to those of [9, 13]. To the best of our knowledge, the only provable online algorithm for this problem is that of Sun and Luo [25]. However the stepsizes they suggest are quite small, leading to suboptimal computational complexity by factors of $\text{poly}(d_1, d_2)$. The runtime of our algorithm is linear in d , which makes $\text{poly}(d)$ improvements over it.

Other models for online matrix completion: Another variant of online matrix completion studied in the literature is where observations are made on a column by column basis e.g., [17, 27]. These models can give improved offline performance in terms of space and could potentially work under relaxed regularity conditions. However, they do not tackle the version where only entries (as opposed to columns) are observed.

Non-convex optimization: Over the last few years, there has also been a significant amount of work in designing other efficient algorithms for solving non-convex problems. Examples include

Algorithm	Sample complexity	Total runtime	Online?
Nuclear Norm [23]	$\tilde{O}(\mu dk)$	$\tilde{O}(d^3/\sqrt{\epsilon})$	No
Alternating minimization [15]	$\tilde{O}(\mu dk \kappa^8 \log \frac{1}{\epsilon})$	$\tilde{O}(\mu dk^2 \kappa^8 \log \frac{1}{\epsilon})$	No
Alternating minimization [9]	$\tilde{O}(\mu dk^2 \kappa^2 (k + \log \frac{1}{\epsilon}))$	$\tilde{O}(\mu dk^3 \kappa^2 (k + \log \frac{1}{\epsilon}))$	No
Projected gradient descent [13]	$\tilde{O}(\mu dk^5)$	$\tilde{O}(\mu dk^7 \log \frac{1}{\epsilon})$	No
SGD [25]	$\tilde{O}(\mu^2 dk^7 \kappa^6)$	$\text{poly}(\mu, d, k, \kappa) \log \frac{1}{\epsilon}$	Yes
SGD [8] ¹	$d \cdot \text{poly}(\mu, k, \kappa)$	$\text{poly}(\mu, d, k, \kappa, \frac{1}{\epsilon})$	Yes
Our result	$\tilde{O}(\mu dk \kappa^4 (k + \log \frac{1}{\epsilon}))$	$\tilde{O}(\mu dk^4 \kappa^4 \log \frac{1}{\epsilon})$	Yes

Table 1: Comparison of sample complexity and runtime of our algorithm with existing algorithms in order to obtain Frobenius norm error ϵ . $\tilde{O}(\cdot)$ hides $\log d$ factors. See Section 1.2 for more discussion.

eigenvector computation [6, 12], sparse coding [21, 1] etc. For general non-convex optimization, an interesting line of recent work is that of [7], which proves gradient descent with noise can also escape saddle point, but they only provide polynomial rate without explicit dependence. Later [18, 22] show that without noise, the space of points from where gradient descent converges to a saddle point is a measure zero set. However, they do not provide a rate of convergence. Another related piece of work to ours is [11], proves global convergence along with rates of convergence, for the special case of computing matrix squareroot. During the preparation of this draft, the recent work [8] was announced which proves the global convergence of SGD for matrix completion and can also be applied to the online setting. However, their result only deals with the case where \mathbf{M} is positive semidefinite (PSD) and their rate is still suboptimal by factors of $\text{poly}(d_1, d_2)$.

1.3 Outline

The rest of the paper is organized as follows. In Section 2 we formally describe the problem and all relevant parameters. In Section 3, we present our algorithms, results and some of the key intuition behind our results. In Section 4 we give proof outline for our main results. We conclude in Section 5. All formal proofs are deferred to the Appendix.

2 Preliminaries

In this section, we introduce our notation, formally define the matrix completion problem and regularity assumptions that make the problem tractable.

¹This result only applies to the case where \mathbf{M} is symmetric PSD

2.1 Notation

We use $[d]$ to denote $\{1, 2, \dots, d\}$. We use bold capital letters \mathbf{A}, \mathbf{B} to denote matrices and bold lowercase letters \mathbf{u}, \mathbf{v} to denote vectors. \mathbf{A}_{ij} means the $(i, j)^{\text{th}}$ entry of matrix \mathbf{A} . $\|\mathbf{w}\|$ denotes the ℓ_2 -norm of vector \mathbf{w} and $\|\mathbf{A}\|/\|\mathbf{A}\|_{\text{F}}/\|\mathbf{A}\|_{\infty}$ denotes the spectral/Frobenius/infinity norm of matrix \mathbf{A} . $\sigma_i(\mathbf{A})$ denotes the i^{th} largest singular value of \mathbf{A} and $\sigma_{\min}(\mathbf{A})$ denotes the smallest singular value of \mathbf{A} . We also let $\kappa(\mathbf{A}) = \|\mathbf{A}\|/\sigma_{\min}(\mathbf{A})$ denote the condition number of \mathbf{A} (i.e., the ratio of largest to smallest singular value). Finally, for orthonormal bases of a subspace \mathbf{W} , we also use $\mathcal{P}_{\mathbf{W}} = \mathbf{W}\mathbf{W}^{\top}$ to denote the projection to the subspace spanned by \mathbf{W} .

2.2 Problem statement and assumptions

Consider a general rank k matrix $\mathbf{M} \in \mathbb{R}^{d_1 \times d_2}$. Let $\Omega \subset [d_1] \times [d_2]$ be a subset of coordinates, which are sampled uniformly and independently from $[d_1] \times [d_2]$. We denote $\mathcal{P}_{\Omega}(\mathbf{M})$ to be the projection of \mathbf{M} on set Ω so that:

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

Low rank matrix completion is the task of recovering \mathbf{M} by only observing $\mathcal{P}_{\Omega}(\mathbf{M})$. This task is ill-posed and NP-hard in general [10]. In order to make this tractable, we make by now standard assumptions about the structure of \mathbf{M} .

Definition 2.1. Let $\mathbf{W} \in \mathbb{R}^{d \times k}$ be an orthonormal basis of a subspace of \mathbb{R}^d of dimension k . The coherence of \mathbf{W} is defined to be

$$\mu(\mathbf{W}) \stackrel{\text{def}}{=} \frac{d}{k} \max_{1 \leq i \leq d} \|\mathcal{P}_{\mathbf{W}} \mathbf{e}_i\|^2 = \frac{d}{k} \max_{1 \leq i \leq d} \|\mathbf{e}_i^{\top} \mathbf{W}\|^2$$

Assumption 2.2 (μ -incoherence[4, 23]). *We assume \mathbf{M} is μ -incoherent, i.e., $\max\{\mu(\mathbf{X}), \mu(\mathbf{Y})\} \leq \mu$, where $\mathbf{X} \in \mathbb{R}^{d_1 \times k}$, $\mathbf{Y} \in \mathbb{R}^{d_2 \times k}$ are the left and right singular vectors of \mathbf{M} .*

3 Main Results

In this section, we present our main result. We will first state result for a special case where \mathbf{M} is a symmetric positive semi-definite (PSD) matrix, where the algorithm and analysis are much simpler. We will then discuss the general case.

3.1 Symmetric PSD Case

Consider the special case where \mathbf{M} is symmetric PSD. We let $d \stackrel{\text{def}}{=} d_1 = d_2$, and we can parametrize a rank k symmetric PSD matrix by $\mathbf{U}\mathbf{U}^{\top}$ where $\mathbf{U} \in \mathbb{R}^{d \times k}$. Our algorithm for this case is given in Algorithm 1. The following theorem provides guarantees on the performance of Algorithm 1. The algorithm starts by using an initial set of samples Ω_{init} to construct a crude approximation to the low rank factorization of \mathbf{M} . It then observes samples from \mathbf{M} one at a time and updates its factorization after every observation. Note that each update step modifies two rows of \mathbf{U}_t and hence takes time $O(k)$.

Algorithm 1 Online Algorithm for PSD Matrix Completion.

Input: Initial set of uniformly random samples Ω_{init} of a symmetric PSD matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$, learning rate η , iterations T
Output: \mathbf{U} such that $\mathbf{U}\mathbf{U}^\top \approx \mathbf{M}$
 $\mathbf{U}_0\mathbf{U}_0^\top \leftarrow$ top k SVD of $\frac{d^2}{|\Omega_{\text{init}}|} \mathcal{P}_{\Omega_{\text{init}}}(\mathbf{M})$
for $t = 0, \dots, T - 1$ **do**
 Observe \mathbf{M}_{ij} where $(i, j) \sim \text{Unif}([d] \times [d])$
 $\mathbf{U}_{t+1} \leftarrow \mathbf{U}_t - 2\eta d^2 (\mathbf{U}_t \mathbf{U}_t^\top - \mathbf{M})_{ij} (\mathbf{e}_i \mathbf{e}_j^\top + \mathbf{e}_j \mathbf{e}_i^\top) \mathbf{U}_t$
end for
Return \mathbf{U}_T

Theorem 3.1. *Let $\mathbf{M} \in \mathbb{R}^{d \times d}$ be a rank k , symmetric PSD matrix with μ -incoherence. There exist some absolute constants c_0 and c such that if $|\Omega_{\text{init}}| \geq c_0 \mu d k^2 \kappa^2(\mathbf{M}) \log d$, learning rate $\eta \leq \frac{c}{\mu d k \kappa^3(\mathbf{M}) \|\mathbf{M}\| \log d}$, then for any fixed $T \geq 1$, with probability at least $1 - \frac{T}{d^{10}}$, we will have for all $t \leq T$ that:*

$$\|\mathbf{U}_t \mathbf{U}_t^\top - \mathbf{M}\|_F^2 \leq \left(1 - \frac{1}{2} \eta \cdot \sigma_{\min}(\mathbf{M})\right)^t \left(\frac{1}{10} \sigma_{\min}(\mathbf{M})\right)^2.$$

Remarks:

- The algorithm uses an initial set of observations Ω_{init} to produce a warm start iterate \mathbf{U}_0 , then enters the online stage, where it performs SGD.
- The sample complexity of the warm start phase is $O(\mu d k^2 \kappa^2(\mathbf{M}) \log d)$. The initialization consists of a top- k SVD on a sparse matrix, whose runtime is $O(\mu d k^3 \kappa^2(\mathbf{M}) \log d)$.
- For the online phase (SGD), if we choose $\eta = \frac{c}{\mu d k \kappa^3(\mathbf{M}) \|\mathbf{M}\| \log d}$, the number of observations T required for the error $\|\mathbf{U}_T \mathbf{U}_T^\top - \mathbf{M}\|_F$ to be smaller than ϵ is $O(\mu d k \kappa(\mathbf{M})^4 \log d \log \frac{\sigma_{\min}(\mathbf{M})}{\epsilon})$.
- Since each SGD step modifies two rows of \mathbf{U}_t , its runtime is $O(k)$ with a total runtime for online phase of $O(kT)$.

Our proof approach is to essentially show that the objective function is well-behaved (i.e., is smooth and strongly convex) in a local neighborhood of the warm start region, and then use standard techniques to show that SGD obtains geometric convergence in this setting. The most challenging and novel part of our analysis comprises of showing that the iterate does not leave this local neighborhood while performing SGD updates. Refer Section 4 for more details on the proof outline.

3.2 General Case

Let us now consider the general case where $\mathbf{M} \in \mathbb{R}^{d_1 \times d_2}$ can be factorized as $\mathbf{U}\mathbf{V}^\top$ with $\mathbf{U} \in \mathbb{R}^{d_1 \times k}$ and $\mathbf{V} \in \mathbb{R}^{d_2 \times k}$. In this scenario, we denote $d = \max\{d_1, d_2\}$. We recall our remarks from the previous section that our analysis of the performance of SGD depends on the smoothness and strong convexity properties of the objective function in a local neighborhood of the iterates. Having $\mathbf{U} \neq \mathbf{V}$ introduces additional challenges in this approach since for any nonsingular k -by- k matrix \mathbf{C} , and $\mathbf{U}' \stackrel{\text{def}}{=} \mathbf{U}\mathbf{C}^\top$, $\mathbf{V}' \stackrel{\text{def}}{=} \mathbf{V}\mathbf{C}^{-1}$, we have $\mathbf{U}'\mathbf{V}'^\top = \mathbf{U}\mathbf{V}^\top$. Suppose for instance \mathbf{C} is a very small scalar times the identity i.e., $\mathbf{C} = \delta \mathbf{I}$ for some small $\delta > 0$. In this case, \mathbf{U}' will be large while \mathbf{V}'

will be small. This drastically deteriorates the smoothness and strong convexity properties of the objective function in a neighborhood of $(\mathbf{U}', \mathbf{V}')$.

Algorithm 2 Online Algorithm for Matrix Completion (Theoretical)

Input: Initial set of uniformly random samples Ω_{init} of $\mathbf{M} \in \mathbb{R}^{d_1 \times d_2}$, learning rate η , iterations T

Output: \mathbf{U}, \mathbf{V} such that $\mathbf{UV}^\top \approx \mathbf{M}$

$\mathbf{U}_0 \mathbf{V}_0^\top \leftarrow$ top k SVD of $\frac{d_1 d_2}{|\Omega_{\text{init}}|} \mathcal{P}_{\Omega_{\text{init}}}(\mathbf{M})$

for $t = 0, \dots, T - 1$ **do**

$\mathbf{W}_U \mathbf{D} \mathbf{W}_V^\top \leftarrow$ SVD($\mathbf{U}_t \mathbf{V}_t^\top$)

$\tilde{\mathbf{U}}_t \leftarrow \mathbf{W}_U \mathbf{D}^{\frac{1}{2}}, \quad \tilde{\mathbf{V}}_t \leftarrow \mathbf{W}_V \mathbf{D}^{\frac{1}{2}}$

Observe \mathbf{M}_{ij} where $(i, j) \sim \text{Unif}([d] \times [d])$

$\mathbf{U}_{t+1} \leftarrow \tilde{\mathbf{U}}_t - 2\eta d_1 d_2 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij} \mathbf{e}_i \mathbf{e}_j^\top \tilde{\mathbf{V}}_t$

$\mathbf{V}_{t+1} \leftarrow \tilde{\mathbf{V}}_t - 2\eta d_1 d_2 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij} \mathbf{e}_j \mathbf{e}_i^\top \tilde{\mathbf{U}}_t$

end for

Return $\mathbf{U}_T, \mathbf{V}_T$.

To preclude such a scenario, we would ideally like to renormalize after each step by doing $\tilde{\mathbf{U}}_t \leftarrow \mathbf{W}_U \mathbf{D}^{\frac{1}{2}}, \tilde{\mathbf{V}}_t \leftarrow \mathbf{W}_V \mathbf{D}^{\frac{1}{2}}$, where $\mathbf{W}_U \mathbf{D} \mathbf{W}_V^\top$ is the SVD of matrix $\mathbf{U}_t \mathbf{V}_t^\top$. This algorithm is described in Algorithm 2. However, a naive implementation of Algorithm 2, especially the SVD step, would incur $O(\min\{d_1, d_2\})$ computation per iteration, resulting in a runtime overhead of $O(d)$ over both the online PSD case (i.e., Algorithm 1) as well as the near linear time offline algorithms (see Table 1). It turns out that we can take advantage of the fact that in each iteration we only update a single row of \mathbf{U}_t and a single row of \mathbf{V}_t , and do efficient (but more complicated) update steps instead of doing an SVD on $d_1 \times d_2$ matrix. The resulting algorithm is given in Algorithm 3. The key idea is that in order to implement the updates, it suffices to do an SVD of $\mathbf{U}_t^\top \mathbf{U}_t$ and $\mathbf{V}_t^\top \mathbf{V}_t$ which are $k \times k$ matrices. So the runtime of each iteration is at most $O(k^3)$. The following lemma shows the equivalence between Algorithms 2 and 3.

Algorithm 3 Online Algorithm for Matrix Completion (Practical)

Input: Initial set of uniformly random samples Ω_{init} of $\mathbf{M} \in \mathbb{R}^{d_1 \times d_2}$, learning rate η , iterations T

Output: \mathbf{U}, \mathbf{V} such that $\mathbf{UV}^\top \approx \mathbf{M}$

$\mathbf{U}_0 \mathbf{V}_0^\top \leftarrow$ top k SVD of $\frac{d_1 d_2}{|\Omega_{\text{init}}|} \mathcal{P}_{\Omega_{\text{init}}}(\mathbf{M})$

for $t = 0, \dots, T - 1$ **do**

$\mathbf{R}_U \mathbf{D}_U \mathbf{R}_U^\top \leftarrow$ SVD($\mathbf{U}_t^\top \mathbf{U}_t$)

$\mathbf{R}_V \mathbf{D}_V \mathbf{R}_V^\top \leftarrow$ SVD($\mathbf{V}_t^\top \mathbf{V}_t$)

$\mathbf{Q}_U \mathbf{D} \mathbf{Q}_V^\top \leftarrow$ SVD($\mathbf{D}_U^{\frac{1}{2}} \mathbf{R}_U^\top \mathbf{R}_V (\mathbf{D}_V^{\frac{1}{2}})^\top$)

Observe \mathbf{M}_{ij} where $(i, j) \sim \text{Unif}([d] \times [d])$

$\mathbf{U}_{t+1} \leftarrow \mathbf{U}_t - 2\eta d_1 d_2 (\mathbf{U}_t \mathbf{V}_t^\top - \mathbf{M})_{ij} \mathbf{e}_i \mathbf{e}_j^\top \mathbf{V}_t \mathbf{R}_V \mathbf{D}_V^{-\frac{1}{2}} \mathbf{Q}_V \mathbf{Q}_U^\top \mathbf{D}_U^{\frac{1}{2}} \mathbf{R}_U^\top$

$\mathbf{V}_{t+1} \leftarrow \mathbf{V}_t - 2\eta d_1 d_2 (\mathbf{U}_t \mathbf{V}_t^\top - \mathbf{M})_{ij} \mathbf{e}_j \mathbf{e}_i^\top \mathbf{U}_t \mathbf{R}_U \mathbf{D}_U^{-\frac{1}{2}} \mathbf{Q}_U \mathbf{Q}_V^\top \mathbf{D}_V^{\frac{1}{2}} \mathbf{R}_V^\top$

end for

Return $\mathbf{U}_T, \mathbf{V}_T$.

Lemma 3.2. *Algorithm 2 and Algorithm 3 are equivalent in the sense that: given same observations*

from \mathbf{M} and other inputs, the outputs of Algorithm 2, \mathbf{U}, \mathbf{V} and those of Algorithm 3, \mathbf{U}', \mathbf{V}' satisfy $\mathbf{U}\mathbf{V}^\top = \mathbf{U}'\mathbf{V}'^\top$.

Since the output of both algorithms is the same, we can analyze Algorithm 2 (which is easier than that of Algorithm 3), while implementing Algorithm 3 in practice. The following theorem is the main result of our paper which presents guarantees on the performance of Algorithm 2.

Theorem 3.3. *Let $\mathbf{M} \in \mathbb{R}^{d_1 \times d_2}$ be a rank k matrix with μ -incoherence and let $d \stackrel{\text{def}}{=} \max(d_1, d_2)$. There exist some absolute constants c_0 and c such that if $|\Omega_{\text{init}}| \geq c_0 \mu d k^2 \kappa^2(\mathbf{M}) \log d$, learning rate $\eta \leq \frac{c}{\mu d k \kappa^3(\mathbf{M}) \|\mathbf{M}\| \log d}$, then for any fixed $T \geq 1$, with probability at least $1 - \frac{T}{d^{10}}$, we will have for all $t \leq T$ that:*

$$\|\mathbf{U}_t \mathbf{V}_t^\top - \mathbf{M}\|_F^2 \leq \left(1 - \frac{1}{2} \eta \cdot \sigma_{\min}(\mathbf{M})\right)^t \left(\frac{1}{10} \sigma_{\min}(\mathbf{M})\right)^2.$$

Remarks:

- Just as in the case of PSD matrix completion (Theorem 3.1), Algorithm 2 needs a an initial set of observations Ω_{init} to provide a warm start \mathbf{U}_0 and \mathbf{V}_0 after which it performs SGD.
- The sample complexity and runtime of the warm start phase are the same as in symmetric PSD case. The stepsize η and the number of observations T to achieve ϵ error in online phase (SGD) are also the same as in symmetric PSD case.
- However, runtime of each update step in online phase is $O(k^3)$ with total runtime for online phase $O(k^3 T)$.

The proof of this theorem again follows a similar line of reasoning as that of Theorem 3.1 by first showing that the local neighborhood of warm start iterate has good smoothness and strong convexity properties and then use them to show geometric convergence of SGD. Proof of the fact that iterates do not move away from this local neighborhood however is significantly more challenging due to renormalization steps in the algorithm. Please see Appendix C for the full proof.

4 Proof Sketch

In this section we will provide the intuition and proof sketch for our main results. For simplicity and highlighting the most essential ideas, we will mostly focus on the symmetric PSD case (Theorem 3.1). For the asymmetric case, though the high-level ideas are still valid, a lot of additional effort is required to address the renormalization step in Algorithm 2. This makes the proof more involved.

First, note that our algorithm for the PSD case consists of an initialization and then stochastic descent steps. The following lemma provides guarantees on the error achieved by the initial iterate \mathbf{U}_0 .

Lemma 4.1. *Let $\mathbf{M} \in \mathbb{R}^{d \times d}$ be a rank- k PSD matrix with μ -incoherence. There exists a constant c_0 such that if $|\Omega_{\text{init}}| \geq c_0 \mu d k^2 \kappa^2(\mathbf{M}) \log d$, then with probability at least $1 - \frac{1}{d^{10}}$, the top- k SVD of $\frac{d^2}{|\Omega_{\text{init}}|} \mathcal{P}_{\Omega_{\text{init}}}(\mathbf{M})$ satisfies Then there exists universal constant c_0 , for any $m \geq$, we have:*

$$\|\mathbf{M} - \mathbf{U}_0 \mathbf{U}_0^\top\|_F \leq \frac{1}{20} \sigma_{\min}(\mathbf{M}) \quad \text{and} \quad \max_j \left\| \mathbf{e}_j^\top \mathbf{U}_0 \right\|^2 \leq \frac{10 \mu k \kappa(\mathbf{M})}{d} \|\mathbf{M}\| \quad (3)$$

By Lemma 4.1, we know the initialization algorithm already gives \mathbf{U}_0 in the local region given by Eq.(3). Intuitively, stochastic descent steps should keep doing local search within this local region.

To establish linear convergence on $\|\mathbf{U}_t \mathbf{U}_t^\top - \mathbf{M}\|_F^2$ and obtain final result, we first establish several important lemmas describing the properties of this local regions. Throughout this section, we always denote $\text{SVD}(\mathbf{M}) = \mathbf{X}\mathbf{S}\mathbf{X}^\top$, where $\mathbf{X} \in \mathbb{R}^{d \times k}$, and diagonal matrix $\mathbf{S} \in \mathbb{R}^{k \times k}$. We postpone all the formal proofs in Appendix.

Lemma 4.2. *For function $f(\mathbf{U}) = \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F^2$ and any $\mathbf{U}_1, \mathbf{U}_2 \in \{\mathbf{U} \mid \|\mathbf{U}\| \leq \Gamma\}$, we have:*

$$\|\nabla f(\mathbf{U}_1) - \nabla f(\mathbf{U}_2)\|_F \leq 16 \max\{\Gamma^2, \|\mathbf{M}\|\} \cdot \|\mathbf{U}_1 - \mathbf{U}_2\|_F$$

Lemma 4.3. *For function $f(\mathbf{U}) = \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F^2$ and any $\mathbf{U} \in \{\mathbf{U} \mid \sigma_{\min}(\mathbf{X}^\top \mathbf{U}) \geq \gamma\}$, we have:*

$$\|\nabla f(\mathbf{U})\|_F^2 \geq 4\gamma^2 f(\mathbf{U})$$

Lemma 4.2 tells function f is smooth if spectral norm of \mathbf{U} is not very large. On the other hand, $\sigma_{\min}(\mathbf{X}^\top \mathbf{U})$ not too small requires both $\sigma_{\min}(\mathbf{U}^\top \mathbf{U})$ and $\sigma_{\min}(\mathbf{X}^\top \mathbf{W})$ are not too small, where \mathbf{W} is top-k eigenspace of $\mathbf{U}\mathbf{U}^\top$. That is, Lemma 4.3 tells function f has a property similar to strongly convex in standard optimization literature, if \mathbf{U} is rank k in a robust sense ($\sigma_k(\mathbf{U})$ is not too small), and the angle between the top k eigenspace of $\mathbf{U}\mathbf{U}^\top$ and the top k eigenspace \mathbf{M} is not large.

Lemma 4.4. *Within the region $\mathcal{D} = \{\mathbf{U} \mid \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F \leq \frac{1}{10}\sigma_k(\mathbf{M})\}$, we have:*

$$\|\mathbf{U}\| \leq \sqrt{2\|\mathbf{M}\|}, \quad \sigma_{\min}(\mathbf{X}^\top \mathbf{U}) \geq \sqrt{\sigma_k(\mathbf{M})/2}$$

Lemma 4.4 tells inside region $\{\mathbf{U} \mid \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F \leq \frac{1}{10}\sigma_k(\mathbf{M})\}$, matrix \mathbf{U} always has a good spectral property which gives preconditions for both Lemma 4.2 and 4.3, where $f(\mathbf{U})$ is both smooth and has a property very similar to strongly convex.

With above three lemmas, we already been able to see the intuition behind linear convergence in Theorem 3.1. Denote stochastic gradient

$$SG(\mathbf{U}) = 2d^2(\mathbf{U}\mathbf{U}^\top - \mathbf{M})_{ij}(\mathbf{e}_i \mathbf{e}_j^\top + \mathbf{e}_j \mathbf{e}_i^\top) \mathbf{U} \quad (4)$$

where $SG(\mathbf{U})$ is a random matrix depends on the randomness of sample (i, j) of matrix \mathbf{M} . Then, the stochastic update step in Algorithm 1 can be rewritten as:

$$\mathbf{U}_{t+1} \leftarrow \mathbf{U}_t - \eta SG(\mathbf{U}_t)$$

Let $f(\mathbf{U}) = \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F^2$, By easy calculation, we know $\mathbb{E}SG(\mathbf{U}) = \nabla f(\mathbf{U})$, that is $SG(\mathbf{U})$ is unbiased. Combine Lemma 4.4 with Lemma 4.2 and Lemma 4.3, we know within region \mathcal{D} specified by Lemma 4.4, we have function $f(\mathbf{U})$ is $32\|\mathbf{M}\|$ -smooth, and $\|\nabla f(\mathbf{U})\|_F^2 \geq 2\sigma_{\min}(\mathbf{M})f(\mathbf{U})$.

Let's suppose ideally, we always have $\mathbf{U}_0, \dots, \mathbf{U}_t$ inside region \mathcal{D} , this directly gives:

$$\begin{aligned} \mathbb{E}f(\mathbf{U}_{t+1}) &\leq \mathbb{E}f(\mathbf{U}_t) - \eta \mathbb{E}\langle \nabla f(\mathbf{U}_t), SG(\mathbf{U}_t) \rangle + 16\eta^2 \|\mathbf{M}\| \cdot \mathbb{E}\|SG(\mathbf{U}_t)\|_F^2 \\ &= \mathbb{E}f(\mathbf{U}_t) - \eta \mathbb{E}\|\nabla f(\mathbf{U}_t)\|_F^2 + 16\eta^2 \|\mathbf{M}\| \cdot \mathbb{E}\|SG(\mathbf{U}_t)\|_F^2 \\ &\leq (1 - 2\eta\sigma_{\min}(\mathbf{M}))\mathbb{E}f(\mathbf{U}_t) + 16\eta^2 \|\mathbf{M}\| \cdot \mathbb{E}\|SG(\mathbf{U}_t)\|_F^2 \end{aligned}$$

One interesting aspect of our main result is that we actually show linear convergence under the presence of noise in gradient. This is true because for the second-order (η^2) term above, we can roughly see from Eq.(4) that $\|SG(\mathbf{U})\|_F^2 \leq h(\mathbf{U}) \cdot f(\mathbf{U})$, where $h(\mathbf{U})$ is a factor depends on \mathbf{U} and always bounded. That is, $SG(\mathbf{U})$ enjoys self-bounded property — $\|SG(\mathbf{U})\|_F^2$ will goes to zero, as objective function $f(\mathbf{U})$ goes to zero. Therefore, by choosing learning rate η appropriately small, we can have the first-order term always dominate the second-order term, which establish the linear convergence.

Now, the only remaining issue is to prove that “ $\mathbf{U}_0, \dots, \mathbf{U}_t$ always stay inside local region \mathcal{D} ”. In reality, we can only prove this statement with high probability due to the stochastic nature of the update. This is also the most challenging part in our proof, which makes our analysis different from standard convex analysis, and uniquely required due to non-convex setting.

Our key theorem is presented as follows:

Theorem 4.5. *Let $f(\mathbf{U}) = \|\mathbf{U}\mathbf{U}^\top - \mathbf{M}\|_F^2$ and $g_i(\mathbf{U}) = \|\mathbf{e}_i^\top \mathbf{U}\|^2$. Suppose initial \mathbf{U}_0 satisfying:*

$$f(\mathbf{U}_0) \leq \left(\frac{\sigma_{\min}(\mathbf{M})}{20} \right)^2, \quad \max_i g_i(\mathbf{U}_0) \leq \frac{10\mu k \kappa(\mathbf{M})^2}{d} \|\mathbf{M}\|$$

Then, there exist some absolute constant c such that for any learning rate $\eta < \frac{c}{\mu d k \kappa^3(\mathbf{M}) \|\mathbf{M}\| \log d}$, with at least $1 - \frac{T}{d^{10}}$ probability, we will have for all $t \leq T$ that:

$$f(\mathbf{U}_t) \leq \left(1 - \frac{1}{2}\eta\sigma_{\min}(\mathbf{M})\right)^t \left(\frac{\sigma_{\min}(\mathbf{M})}{10} \right)^2, \quad \max_i g_i(\mathbf{U}_t) \leq \frac{20\mu k \kappa(\mathbf{M})^2}{d} \|\mathbf{M}\| \quad (5)$$

Note function $\max_i g_i(\mathbf{U})$ indicates the incoherence of matrix \mathbf{U} . Theorem 4.5 guarantees if initial \mathbf{U}_0 is in the local region which is incoherent and $\mathbf{U}_0\mathbf{U}_0^\top$ is close to \mathbf{M} , then with high probability for all steps $t \leq T$, $\mathbf{U}_t, \mathbf{U}_t$ will always stay in a slightly relaxed local region, and $f(\mathbf{U}_t)$ has linear convergence.

It is not hard to show that all saddle point of $f(\mathbf{U})$ satisfies $\sigma_k(\mathbf{U}) = 0$, and all local minima are global minima. Since $\mathbf{U}_0, \dots, \mathbf{U}_t$ automatically stay in region $f(\mathbf{U}) \leq \left(\frac{\sigma_{\min}(\mathbf{M})}{10}\right)^2$ with high probability, we know \mathbf{U}_t also stay away from all saddle points. The claim that $\mathbf{U}_0, \dots, \mathbf{U}_t$ stays incoherent is essential to better control the variance and probability 1 bound of $SG(\mathbf{U}_t)$, so that we can have large step size and tight convergence rate.

The major challenging in proving Theorem 4.5 is to both prove \mathbf{U}_t stays in the local region, and achieve good sample complexity and running time (linear in d) in the same time. This also requires the learning rate η in Algorithm 1 to be relatively large. Let the event \mathfrak{E}_t denote the good event where $\mathbf{U}_0, \dots, \mathbf{U}_t$ satisfies Eq.(5). Theorem 4.5 is claiming that $P(\mathfrak{E}_T)$ is large. The essential steps in the proof is constructing two supermartingales related to $f(\mathbf{U}_t)1_{\mathfrak{E}_t}$ and $g_i(\mathbf{U}_t)1_{\mathfrak{E}_t}$ (where $1_{(\cdot)}$ denote indicator function), and use Bernstein inequality to show the concentration of supermartingales. The $1_{\mathfrak{E}_t}$ term allow us the claim all previous $\mathbf{U}_0, \dots, \mathbf{U}_t$ have all desired properties inside local region.

Finally, we see Theorem 3.1 as a immediate corollary of Theorem 4.5.

5 Conclusion

In this paper, we presented the first provable, efficient online algorithm for matrix completion, based on nonconvex SGD. In addition to the online setting, our results are also competitive with state of

the art results in the offline setting. We obtain our results by introducing a general framework that helps us show how SGD updates self-regulate to stay away from saddle points. We hope our paper and results help generate interest in online matrix completion, and our techniques and framework prompt tighter analysis for other nonconvex problems.

References

- [1] Sanjeev Arora, Rong Ge, Tengyu Ma, and Ankur Moitra. Simple, efficient, and neural algorithms for sparse coding. *arXiv preprint arXiv:1503.00778*, 2015.
- [2] Matthew Brand. Fast online svd revisions for lightweight recommender systems. In *SDM*, pages 37–46. SIAM, 2003.
- [3] Emmanuel J Candes, Yonina C Eldar, Thomas Strohmer, and Vladislav Voroninski. Phase retrieval via matrix completion. *SIAM Review*, 57(2):225–251, 2015.
- [4] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, December 2009.
- [5] James Davidson, Benjamin Liebald, Junning Liu, Palash Nandy, Taylor Van Vleet, Ullas Gargi, Sujoy Gupta, Yu He, Mike Lambert, Blake Livingston, et al. The youtube video recommendation system. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 293–296. ACM, 2010.
- [6] Christopher De Sa, Kunle Olukotun, and Christopher Ré. Global convergence of stochastic gradient descent for some non-convex matrix problems. *arXiv preprint arXiv:1411.1134*, 2014.
- [7] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. *arXiv preprint arXiv:1503.02101*, 2015.
- [8] Rong Ge, Jason D. Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. *arXiv preprint arXiv:1605.07272*, 2016.
- [9] Marcus Hardt. Understanding alternating minimization for matrix completion. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 651–660. IEEE, 2014.
- [10] Moritz Hardt, Raghu Meka, Prasad Raghavendra, and Benjamin Weitz. Computational limits for matrix completion. In *COLT*, pages 703–725, 2014.
- [11] Prateek Jain, Chi Jin, Sham M Kakade, and Praneeth Netrapalli. Computing matrix square-root via non convex local search. *arXiv preprint arXiv:1507.05854*, 2015.
- [12] Prateek Jain, Chi Jin, Sham M Kakade, Praneeth Netrapalli, and Aaron Sidford. Matching matrix bernstein with little memory: Near-optimal finite sample guarantees for oja’s algorithm. *arXiv preprint arXiv:1602.06929*, 2016.
- [13] Prateek Jain and Praneeth Netrapalli. Fast exact matrix completion with finite samples. *arXiv preprint arXiv:1411.1087*, 2014.

- [14] Hui Ji, Chaoqiang Liu, Zuowei Shen, and Yuhong Xu. Robust video denoising using low rank matrix completion. 2010.
- [15] Raghunandan Hulikal Keshavan. *Efficient algorithms for collaborative filtering*. PhD thesis, STANFORD UNIVERSITY, 2012.
- [16] Yehuda Koren. The BellKor solution to the Netflix grand prize, 2009.
- [17] Akshay Krishnamurthy and Aarti Singh. Low-rank matrix and tensor completion via adaptive sampling. In *Advances in Neural Information Processing Systems*, pages 836–844, 2013.
- [18] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent converges to minimizers. *University of California, Berkeley*, 1050:16, 2016.
- [19] G. Linden, B. Smith, and J. York. Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, Jan 2003.
- [20] Xin Luo, Yunni Xia, and Qingsheng Zhu. Incremental collaborative filtering recommender based on regularized matrix factorization. *Knowledge-Based Systems*, 27:271–280, 2012.
- [21] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online learning for matrix factorization and sparse coding. *The Journal of Machine Learning Research*, 11:19–60, 2010.
- [22] Ioannis Panageas and Georgios Piliouras. Gradient descent converges to minimizers: The case of non-isolated critical points. *arXiv preprint arXiv:1605.00405*, 2016.
- [23] Benjamin Recht. A simple approach to matrix completion, 2009.
- [24] Benjamin Recht and Christopher Ré. Parallel stochastic gradient algorithms for large-scale matrix completion. *Mathematical Programming Computation*, 5(2):201–226, 2013.
- [25] Ruoyu Sun and Zhi-Quan Luo. Guaranteed matrix completion via nonconvex factorization. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 270–289. IEEE, 2015.
- [26] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- [27] Se-Young Yun, Marc Lelarge, and Alexandre Proutiere. Streaming, memory limited matrix completion with noise. *arXiv preprint arXiv:1504.03156*, 2015.

A Proof of Initialization

In this section, we will prove Lemma 4.1 and a corresponding lemma for asymmetric case as follows (which will be used to prove Theorem 3.3):

Lemma A.1. *Assume $\mathbf{M} \in \mathbb{R}^{d_1 \times d_2}$ is a rank k matrix with μ -incoherence, and Ω is a subset uniformly i.i.d sampled from all coordinate. Let $\mathbf{U}_0 \mathbf{V}_0^\top$ be the top- k SVD of $\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})$, where $|\Omega| = m$. Let $d = \max\{d_1, d_2\}$. Then there exists universal constant c_0 , for any $m \geq c_0 \mu d k^2 \kappa^2(\mathbf{M}) \log d$, with probability at least $1 - \frac{1}{d^{10}}$, we have:*

$$\begin{aligned} \|\mathbf{M} - \mathbf{U}_0 \mathbf{V}_0^\top\|_F &\leq \frac{1}{20} \sigma_{\min}(\mathbf{M}), \\ \max_i \left\| \mathbf{e}_i^\top \mathbf{U}_0 \mathbf{V}_0^\top \right\|^2 &\leq \frac{10\mu k}{d_1} \|\mathbf{M}\|, \quad \max_j \left\| \mathbf{e}_j^\top \mathbf{V}_0 \mathbf{U}_0^\top \right\|^2 \leq \frac{10\mu k}{d_2} \|\mathbf{M}\| \end{aligned} \quad (6)$$

We will focus mostly on Lemma A.1, and prove Lemma 4.1 as a special case. Most of the argument of this section follows from [15]. We include here for completeness. The remaining of this section can be viewed as proving both the Frobenius norm claim and incoherence claim of Lemma A.1 separately.

In this section, We always denote $d = \max\{d_1, d_2\}$. For simplicity, WLOG, we also assume $\|\mathbf{M}\| = 1$ in all proof. Also, when it's clear from the context, we use κ to specifically to represent $\kappa(\mathbf{M})$. Then $\sigma_{\min}(\mathbf{M}) = \frac{1}{\kappa}$. Also in the proof, we always denote $\text{SVD}(\mathbf{M}) = \mathbf{X} \mathbf{S} \mathbf{Y}^\top$, and $\text{SVD}(\mathbf{U} \mathbf{V}^\top) = \mathbf{W}_U \mathbf{D} \mathbf{W}_V^\top$, where \mathbf{S} and \mathbf{D} are $k \times k$ diagonal matrix.

A.1 Frobenius Norm of Initialization

Theorem A.2 (Matrix Bernstein [26]). *A finite sequence $\{\mathbf{X}_t\}$ of independent, random matrices with dimension $d_1 \times d_2$. Assume that each matrix satisfies:*

$$\mathbb{E} \mathbf{X}_t = 0, \quad \text{and} \quad \|\mathbf{X}_t\| \leq R \quad \text{almost surely}$$

Define

$$\sigma^2 = \max \left\{ \left\| \sum_t \mathbb{E}(\mathbf{X}_t \mathbf{X}_t^\top) \right\|, \left\| \sum_t \mathbb{E}(\mathbf{X}_t^\top \mathbf{X}_t) \right\| \right\}$$

Then, for all $s \geq 0$,

$$\Pr \left(\left\| \sum_t \mathbf{X}_t \right\| \geq s \right) \leq (d_1 + d_2) \cdot \exp \left(\frac{-s^2/2}{\sigma^2 + Rs/3} \right)$$

Lemma A.3. *Let $|\Omega| = m$, then there exists universal constant C, c_0 , for any $m \geq c_0 \mu d k \log d$, with probability at least $1 - \frac{1}{d^{10}}$, we have:*

$$\left\| \mathbf{M} - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) \right\| \leq C \sqrt{\frac{\mu d k \log d}{m}}$$

Proof. We know

$$\left\| \mathbf{M} - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) \right\| = \frac{d_1 d_2}{m} \left\| \mathcal{P}_\Omega(\mathbf{M}) - \frac{m}{d_1 d_2} \mathbf{M} \right\|$$

and note:

$$\mathcal{P}_\Omega(\mathbf{M}) - \frac{m}{d_1 d_2} \mathbf{M} = \sum_{ij} \mathbf{M}_{ij} (Z_{ij} - \frac{m}{d_1 d_2}) \mathbf{e}_i \mathbf{e}_j^\top$$

where Z_{ij} are independence Bernoulli($m/d_1 d_2$) random variables. Let matrix

$$\psi_{ij} = \mathbf{M}_{ij} (Z_{ij} - \frac{m}{d_1 d_2}) \mathbf{e}_i \mathbf{e}_j^\top$$

By construction, we have:

$$\left\| \sum_{ij} \psi_{ij} \right\| = \left\| \mathcal{P}_\Omega(\mathbf{M}) - \frac{m}{d_1 d_2} \mathbf{M} \right\|$$

Clearly $\mathbb{E} \psi_{ij} = 0$. Let $\mathbf{XSY}^\top = \text{SVD}(\mathbf{M})$, then by μ -incoherence of \mathbf{M} , with probability 1:

$$\|\mathbf{M}\|_\infty \leq \max_{ij} |\mathbf{e}_i^\top \mathbf{XSY}^\top \mathbf{e}_j| \leq \|\mathbf{M}\| \frac{\mu k}{\sqrt{d_1 d_2}}$$

Also:

$$\begin{aligned} \left\| \sum_{ij} \mathbb{E}(\psi_{ij} \psi_{ij}^\top) \right\| &= \left\| \sum_{ij} \mathbb{E} \mathbf{M}_{ij}^2 (Z_{ij} - \frac{m}{d_1 d_2})^2 \mathbf{e}_i \mathbf{e}_i^\top \right\| \leq \frac{m}{d_1 d_2} (1 - \frac{m}{d_1 d_2}) \left\| \sum_{ij} \mathbf{M}_{ij}^2 \mathbf{e}_i \mathbf{e}_i^\top \right\| \\ &= \frac{m}{d_1 d_2} (1 - \frac{m}{d_1 d_2}) \max_i \sum_j \mathbf{M}_{ij}^2 \leq \frac{2m}{d_1 d_2} \frac{\mu k}{d_1} \|\mathbf{M}\|^2 = \frac{2m\mu k}{d_1^2 d_2} \|\mathbf{M}\|^2 \\ \left\| \sum_{ij} \mathbb{E}(\psi_{ij}^\top \psi_{ij}) \right\| &= \left\| \sum_{ij} \mathbb{E} \mathbf{M}_{ij}^2 (Z_{ij} - \frac{m}{d_1 d_2})^2 \mathbf{e}_j \mathbf{e}_j^\top \right\| \leq \frac{m}{d_1 d_2} (1 - \frac{m}{d_1 d_2}) \left\| \sum_{ij} \mathbf{M}_{ij}^2 \mathbf{e}_j \mathbf{e}_j^\top \right\| \\ &= \frac{m}{d_1 d_2} (1 - \frac{m}{d_1 d_2}) \max_j \sum_i \mathbf{M}_{ij}^2 \leq \frac{2m}{d_1 d_2} \frac{\mu k}{d_2} \|\mathbf{M}\|^2 = \frac{2m\mu k}{d_1 d_2^2} \|\mathbf{M}\|^2 \end{aligned}$$

Then, by matrix Bernstein (Theorem A.2), we have:

$$\Pr\left(\left\| \sum_{ij} \psi_{ij} \right\| \geq s \right) \leq 2(d_1 + d_2) \cdot \exp\left(\frac{-s^2/2}{\frac{2m\mu dk}{d_1^2 d_2^2} \|\mathbf{M}\|^2 + \|\mathbf{M}\| \frac{\mu k}{3\sqrt{d_1 d_2}} s} \right)$$

That is, with probability at least $1 - \frac{1}{d^{10}}$, for some universal constant C , we have:

$$\left\| \mathcal{P}_\Omega(\mathbf{M}) - \frac{m}{d_1 d_2} \mathbf{M} \right\| \leq C \|\mathbf{M}\| \cdot \max\left\{ \sqrt{\frac{m\mu dk \log d}{d_1^2 d_2^2}}, \frac{\mu k \log d}{\sqrt{d_1 d_2}} \right\}$$

For $m \geq \mu dk \log d$, we finishes the proof. \square

Theorem A.4. Let $\mathbf{U}_0 \mathbf{V}_0^\top$ be the top- k SVD of $\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})$, where $|\Omega| = m$ then there exists universal constant c_0 , for any $m \geq c_0 \mu dk^2 \kappa^2 \log d$, with probability at least $1 - \frac{1}{d^{10}}$, we have:

$$\left\| \mathbf{M} - \mathbf{U}_0 \mathbf{V}_0^\top \right\|_F \leq \frac{1}{20\kappa}$$

Proof. Since \mathbf{M} is a rank k matrix, we know $\sigma_{k+1}(\mathbf{M}) = 0$, thus

$$\sigma_{k+1}\left(\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})\right) \leq \sigma_{k+1}(\mathbf{M}) + \left\| \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) - \mathbf{M} \right\| = \left\| \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) - \mathbf{M} \right\|$$

Therefore:

$$\begin{aligned} \left\| \mathbf{M} - \mathbf{U}_0 \mathbf{V}_0^\top \right\| &\leq \left\| \mathbf{M} - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) \right\| + \left\| \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) - \mathbf{U}_0 \mathbf{V}_0^\top \right\| \\ &\leq \left\| \mathbf{M} - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) \right\| + \sigma_{k+1}\left(\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})\right) \leq 2 \left\| \mathbf{M} - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) \right\| \end{aligned}$$

Meanwhile, since $\text{rank}(\mathbf{M}) = k$, $\text{rank}(\mathbf{U}_0 \mathbf{V}_0^\top) = k$, we know: $\text{rank}(\mathbf{M} - \mathbf{U}_0 \mathbf{V}_0^\top) \leq 2k$, and therefore:

$$\left\| \mathbf{M} - \mathbf{U}_0 \mathbf{V}_0^\top \right\|_F \leq \sqrt{2k} \left\| \mathbf{M} - \mathbf{U}_0 \mathbf{V}_0^\top \right\| \leq 2\sqrt{2k} \left\| \mathbf{M} - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M}) \right\|$$

by choosing $m \geq c_0 \mu d \kappa^2 \log d \cdot \kappa^2$ for large enough constant c_0 and apply Lemma A.3, we finishes the proof. \square

A.2 Incoherence of Initialization

Lemma A.5. *Let $\mathbf{U}\mathbf{V}^\top$ be the top- k SVD of $\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})$, where $|\Omega| = m$. then there exists universal constant c_0 , for any $m \geq c_0 \mu d \kappa^2 \log d$, with probability at least $1 - \frac{1}{d^{10}}$, we have:*

$$\max_j \left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \mathbf{V}\mathbf{U}^\top) \right\| \leq 2\sqrt{\frac{\mu k}{d_2}}$$

Proof. Suppose $\text{SVD}(\mathbf{M}) = \mathbf{X}\mathbf{S}\mathbf{Y}^\top$. Denote $\tilde{\mathbf{X}} = \mathbf{X}\mathbf{S}^{\frac{1}{2}}$ and $\tilde{\mathbf{Y}} = \mathbf{Y}\mathbf{S}^{\frac{1}{2}}$. Also let $\text{SVD}(\mathbf{U}\mathbf{V}^\top) = \mathbf{W}_\mathbf{U}\mathbf{D}\mathbf{W}_\mathbf{V}^\top$.

Then, we have:

$$\begin{aligned} \left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \mathbf{V}\mathbf{U}^\top) \right\| &= \left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})^\top \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top) \right\| \\ &= \left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \mathbf{M}^\top \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top + \mathbf{M}^\top \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})^\top \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top) \right\| \\ &\leq \left\| \mathbf{e}_j^\top \mathbf{M}^\top (\mathbf{I} - \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top) \right\| + \left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})^\top) \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top \right\| \end{aligned}$$

For the first term, since $\mathbf{W}_\mathbf{U}^\top \mathbf{W}_{\mathbf{U},\perp} = 0$, we have:

$$\begin{aligned} \left\| \mathbf{e}_j^\top \mathbf{M}^\top (\mathbf{I} - \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top) \right\| &\leq \left\| \mathbf{e}_j^\top \mathbf{Y} \right\| \left\| \mathbf{S}\mathbf{X}^\top \mathbf{W}_{\mathbf{U},\perp} \mathbf{W}_{\mathbf{U},\perp}^\top \right\| \\ &= \sqrt{\frac{\mu k}{d_2}} \left\| \mathbf{Y}^\top (\mathbf{M}^\top - \mathbf{W}_\mathbf{V} \mathbf{D} \mathbf{W}_\mathbf{U}^\top) \mathbf{W}_{\mathbf{U},\perp} \mathbf{W}_{\mathbf{U},\perp}^\top \right\| \\ &\leq \sqrt{\frac{\mu k}{d_2}} \left\| \mathbf{M}^\top - \mathbf{W}_\mathbf{V} \mathbf{D} \mathbf{W}_\mathbf{U}^\top \right\| \leq \sqrt{\frac{\mu k}{d_2}} \cdot \frac{1}{\kappa} \end{aligned}$$

The last step is due to sample $m \geq \mu d k \kappa^2 \log d$, and theorem A.4.

For the second term, we have:

$$\begin{aligned} & \left\| \mathbf{e}_j^\top \left(\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})^\top - \mathbf{M}^\top \right) \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top \right\| = \left\| \tilde{\mathbf{Y}}_j^\top \left(\frac{d_1 d_2}{m} \sum_{i:(i,j) \in \Omega} \tilde{\mathbf{x}}_i \mathbf{w}_{\mathbf{U},i}^\top - \sum_i \tilde{\mathbf{x}}_i \mathbf{w}_{\mathbf{U},i}^\top \right) \mathbf{W}_\mathbf{U}^\top \right\| \\ & \leq \sqrt{\frac{\mu k}{d_2}} \cdot \frac{d_1 d_2}{m} \cdot \left\| \sum_{i:(i,j) \in \Omega} \tilde{\mathbf{x}}_i \mathbf{w}_{\mathbf{U},i}^\top - \frac{m}{d_1 d_2} \sum_i \tilde{\mathbf{x}}_i \mathbf{w}_{\mathbf{U},i}^\top \right\| \end{aligned} \quad (7)$$

Where $\tilde{\mathbf{x}}_i$ and $\mathbf{w}_{\mathbf{U},i}$ are the i -th row of $\tilde{\mathbf{X}}$ and $\mathbf{W}_\mathbf{U}$ respectively.

Let $\phi_{ij} = \tilde{\mathbf{x}}_i \mathbf{w}_{\mathbf{U},i}^\top (Z_{ij} - \frac{m}{d_1 d_2})$, where Z_{ij} is Bernoulli($\frac{m}{d_1 d_2}$) random variable, $Z_{ij} = 1$ iff $(i, j) \in \Omega$. Clearly, we have $\mathbb{E}\phi = 0$, and with probability 1:

$$\|\phi_{ij}\| \leq 2 \|\tilde{\mathbf{x}}_i\| \|\mathbf{w}_{\mathbf{U},i}\| \leq 2 \sqrt{\frac{\mu k}{d_1}} \max_i \|\mathbf{e}_i^\top \mathbf{W}_\mathbf{U}\|$$

Also, we have variance term:

$$\begin{aligned} \left\| \sum_i \mathbb{E} \phi_{ij}^\top \phi_{ij} \right\| &= \left\| \sum_i \mathbb{E} (Z_{ij} - \frac{m}{d_1 d_2})^2 \|\tilde{\mathbf{x}}_i\|^2 \mathbf{w}_{\mathbf{U},i} \mathbf{w}_{\mathbf{U},i}^\top \right\| \\ &\leq \frac{m}{d_1 d_2} \left(1 - \frac{m}{d_1 d_2}\right) \max_i \|\tilde{\mathbf{x}}_i\|^2 \left\| \sum_i \mathbf{w}_{\mathbf{U},i} \mathbf{w}_{\mathbf{U},i}^\top \right\| \\ &\leq \frac{m}{d_1 d_2} \frac{\mu k}{d_1} \|\mathbf{W}_\mathbf{U}^\top \mathbf{W}_\mathbf{U}\| \leq \frac{\mu k m}{d_1^2 d_2} \\ \left\| \sum_i \mathbb{E} \phi_{ij} \phi_{ij}^\top \right\| &= \left\| \sum_i \mathbb{E} (Z_{ij} - \frac{m}{d_1 d_2})^2 \|\mathbf{w}_{\mathbf{U},i}\|^2 \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \right\| \\ &\leq \frac{m}{d_1 d_2} \max_i \|\mathbf{e}_i^\top \mathbf{W}_\mathbf{U}\|^2 \end{aligned}$$

Therefore, with $m \geq \mu d k \kappa^2 \log d$, by matrix Bernstein, we have with probability at least $1 - \frac{1}{d^{10}}$, we know that for all $j \in [d_2]$, there exists some absolute constant C' so that:

$$\left\| \sum_{i:(i,j) \in \Omega} \tilde{\mathbf{x}}_i \mathbf{w}_{\mathbf{U},i}^\top - \frac{m}{d_1 d_2} \sum_i \tilde{\mathbf{x}}_i \mathbf{w}_{\mathbf{U},i}^\top \right\| \leq C' \sqrt{\frac{m \log d}{d_1 d_2}} \left(\sqrt{\frac{\mu k}{d_1}} + \max_i \|\mathbf{e}_i^\top \mathbf{W}_\mathbf{U}\| \right)$$

Substitute into Eq.(7), this gives:

$$\left\| \mathbf{e}_j^\top \left(\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})^\top - \mathbf{M}^\top \right) \mathbf{W}_\mathbf{U} \mathbf{W}_\mathbf{U}^\top \right\| \leq C' \sqrt{\frac{\mu k d_1 \log d}{m}} \left(\sqrt{\frac{\mu k}{d_1}} + \max_i \|\mathbf{e}_i^\top \mathbf{W}_\mathbf{U}\| \right)$$

On the other hand, we also have:

$$\begin{aligned} \|\mathbf{e}_i^\top \mathbf{W}_\mathbf{U}\| &\leq \|\mathbf{e}_i^\top \mathbf{W}_\mathbf{U} \mathbf{S}\| \|\mathbf{S}^{-1}\| = 2\kappa \|\mathbf{e}_i^\top \mathbf{U} \mathbf{V}^\top\| \leq 2\kappa \left(\|\mathbf{e}_i^\top (\mathbf{U} \mathbf{V}^\top - \mathbf{M})\| + \|\mathbf{e}_i^\top \mathbf{M}\| \right) \\ &\leq 2\kappa \left(\sqrt{\frac{\mu k}{d_1}} + \|\mathbf{e}_i^\top (\mathbf{U} \mathbf{V}^\top - \mathbf{M})\| \right) \end{aligned}$$

This gives overall inequality:

$$\max_j \left\| \mathbf{e}_j^\top (\mathbf{V}\mathbf{U}^\top - \mathbf{M}^\top) \right\| \leq \sqrt{\frac{\mu k}{d_2}} \cdot \frac{1}{\kappa} + C'' \sqrt{\frac{\mu k d_1 \log d}{m}} \kappa \left(\sqrt{\frac{\mu k}{d_1}} + \max_i \left\| \mathbf{e}_i^\top (\mathbf{U}\mathbf{V}^\top - \mathbf{M}) \right\| \right)$$

By symmetry, we will also have:

$$\max_i \left\| \mathbf{e}_i^\top (\mathbf{U}\mathbf{V}^\top - \mathbf{M}) \right\| \leq \sqrt{\frac{\mu k}{d_1}} \cdot \frac{1}{\kappa} + C'' \sqrt{\frac{\mu k d_2 \log d}{m}} \kappa \left(\sqrt{\frac{\mu k}{d_2}} + \max_j \left\| \mathbf{e}_j^\top (\mathbf{V}\mathbf{U}^\top - \mathbf{M}^\top) \right\| \right)$$

Combine above two equations and choose $m \geq c_0 \mu d k \kappa^2 \log d$ for some large enough c_0 . We have:

$$\max_j \left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \mathbf{V}\mathbf{U}^\top) \right\| \leq 2 \sqrt{\frac{\mu k}{d_2}}$$

This finishes the proof. □

Theorem A.6. Let $\mathbf{U}_0 \mathbf{V}_0^\top$ be the top- k SVD of $\frac{d_1 d_2}{m} \mathcal{P}_\Omega(\mathbf{M})$, where $|\Omega| = m$. then there exists universal constant c_0 , for any $m \geq c_0 \mu d k \kappa^2 \log d$, with probability at least $1 - \frac{1}{d^{10}}$, we have:

$$\max_i \left\| \mathbf{e}_i^\top \mathbf{U}_0 \mathbf{V}_0^\top \right\|^2 \leq \frac{9\mu k}{d_1} \quad \text{and} \quad \max_j \left\| \mathbf{e}_j^\top \mathbf{V}_0 \mathbf{U}_0^\top \right\|^2 \leq \frac{9\mu k}{d_2}$$

Proof. By Theorem A.5, we know for any $j \in [d_2]$:

$$\left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \mathbf{V}_0 \mathbf{U}_0^\top) \right\| \leq 2 \sqrt{\frac{\mu k}{d_2}}$$

Therefore, we have:

$$\left\| \mathbf{e}_j^\top \mathbf{V}_0 \mathbf{U}_0^\top \right\| \leq \left[\left\| \mathbf{e}_j^\top \mathbf{M}^\top \right\| + \left\| \mathbf{e}_j^\top (\mathbf{M}^\top - \mathbf{V}_0 \mathbf{U}_0^\top) \right\| \right] \leq 3 \sqrt{\frac{\mu k}{d_2}}$$

By symmetry, we also know for any $i \in [d_1]$

$$\left\| \mathbf{e}_i^\top \mathbf{U}_0 \mathbf{V}_0^\top \right\| \leq 3 \sqrt{\frac{\mu k}{d_1}}$$

Which finishes the proof. □

For the special case where $\mathbf{M} \in \mathbb{R}^{d \times d}$ is symmetric and PSD, we can easily extends to have following:

Corollary A.7. Let $\mathbf{U}_0 \mathbf{U}_0^\top$ be the top- k SVD of $\frac{d^2}{m} \mathcal{P}_\Omega(\mathbf{M})$, where $|\Omega| = m$. then there exists universal constant c_0 , for any $m \geq c_0 \mu d k \kappa^2 \log d$, with probability at least $1 - \frac{1}{d^{10}}$, we have:

$$\max_i \left\| \mathbf{e}_i^\top \mathbf{U}_0 \right\|^2 \leq \frac{10\mu k \kappa}{d}$$

Proof. By Corollary A.6, we have:

$$\max_i \left\| \mathbf{e}_i^\top \mathbf{U}_0 \mathbf{U}_0^\top \right\|^2 \leq \frac{9\mu k}{d}$$

On the other hand, by Theorem A.4, we have:

$$\sigma_{\min}(\mathbf{U}_0^\top \mathbf{U}_0) = \sigma_k(\mathbf{U}_0 \mathbf{U}_0^\top) \geq \sigma_k(\mathbf{M}) - \left\| \mathbf{M} - \mathbf{U}_0 \mathbf{U}_0^\top \right\| \geq \frac{9}{10\kappa}$$

Therefore, for any $i \in [d]$ we have:

$$\left\| \mathbf{e}_i^\top \mathbf{U}_0 \right\|^2 \leq \frac{\left\| \mathbf{e}_i^\top \mathbf{U}_0 \mathbf{U}_0^\top \right\|^2}{\sigma_{\min}(\mathbf{U}_0^\top \mathbf{U}_0)} \leq \frac{10\mu k \kappa}{d}$$

Which finishes the proof. \square

Finally, Lemma A.1 can be easily concluded from Theorem A.4 and Theorem A.6, while Lemma 4.1 is also directly proved by Theorem A.4 and Corollary A.7.

B Proof of Symmetric PSD Case

In this section, we prove Theorem 3.1. WLOG, we continue to assume $\|\mathbf{M}\| = 1$ in all proof. Also, when it's clear from the context, we use κ to specifically to represent $\kappa(\mathbf{M})$. Then $\sigma_{\min}(\mathbf{M}) = \frac{1}{\kappa}$. Also in this section, we always denote $\text{SVD}(\mathbf{M}) = \mathbf{X}\mathbf{S}\mathbf{X}^\top$, and $\text{SVD}(\mathbf{U}\mathbf{U}^\top) = \mathbf{W}\mathbf{D}\mathbf{W}^\top$.

The most essential part to prove Theorem 3.1 is proving following Theorem:

Theorem B.1 (restatement of Theorem 4.5). *Let $f(\mathbf{U}) = \|\mathbf{U}\mathbf{U}^\top - \mathbf{M}\|_F^2$ and $g_i(\mathbf{U}) = \|\mathbf{e}_i^\top \mathbf{U}\|^2$. Suppose after initialization, we have:*

$$f(\mathbf{U}_0) \leq \left(\frac{1}{20\kappa} \right)^2, \quad \max_i g_i(\mathbf{U}_0) \leq \frac{10\mu k \kappa^2}{d}$$

Then, there exist some absolute constant c such that for any learning rate $\eta < \frac{c}{\mu d k \kappa^3 \log d}$, with at least $1 - \frac{\eta}{10}$ probability, we will have for all $t \leq T$ that:

$$f(\mathbf{U}_t) \leq \left(1 - \frac{\eta}{2\kappa}\right)^t \left(\frac{1}{10\kappa} \right)^2, \quad \max_i g_i(\mathbf{U}_t) \leq \frac{20\mu k \kappa^2}{d}$$

Theorem B.1 says once initialization algorithm provides \mathbf{U}_0 in good local region, with high probability \mathbf{U}_t will always stay in this good region and $f(\mathbf{U}_t)$ is linear converging to 0. With this theorem, we can then immediately conclude Theorem 3.1 from Theorem B.1 and Lemma 4.1.

The rest of this section all focus on proving Theorem B.1. First, we prepare with a few lemmas about the property of objective function, and the spectral property of \mathbf{U} in a local Frobenius ball around optimal. Then, we prove Theorem B.1 by constructing two supermartingales related to $f(\mathbf{U}_t), g_i(\mathbf{U}_t)$ each, and applying concentration argument.

For symmetric PSD case, we denote the stochastic gradient as:

$$SG(\mathbf{U}) = 2d^2 (\mathbf{U}\mathbf{U}^\top - \mathbf{M})_{ij} (\mathbf{e}_i \mathbf{e}_j^\top + \mathbf{e}_j \mathbf{e}_i^\top) \mathbf{U}$$

The update in Algorithm 1 can be now written as:

$$\mathbf{U}_{t+1} \leftarrow \mathbf{U}_t - \eta SG(\mathbf{U}_t) \quad (8)$$

We immediately have the property:

$$\mathbb{E}SG(\mathbf{U}) = \nabla f(\mathbf{U}) = 4(\mathbf{U}\mathbf{U}^\top - \mathbf{M})\mathbf{U}$$

B.1 Geometric Properties in Local Region

First, we prove two lemmas w.r.t the smoothness and property similar to strongly convex for objective function:

Lemma B.2. (restatement of Lemma 4.2) *Within the region $\mathcal{D} = \{\mathbf{U} \mid \|\mathbf{U}\| \leq \Gamma\}$, we have function $f(\mathbf{U}) = \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F^2$ satisfying for any $\mathbf{U}_1, \mathbf{U}_2 \in \mathcal{D}$:*

$$\|\nabla f(\mathbf{U}_1) - \nabla f(\mathbf{U}_2)\|_F \leq \beta \|\mathbf{U}_1 - \mathbf{U}_2\|_F$$

where smoothness parameter $\beta = 16 \max\{\Gamma^2, \|\mathbf{M}\|\}$.

Proof. Inside region \mathcal{D} , we have:

$$\begin{aligned} & \|\nabla f(\mathbf{U}_1) - \nabla f(\mathbf{U}_2)\|_F \\ &= \|4(\mathbf{U}_1\mathbf{U}_1^\top - \mathbf{M})\mathbf{U}_1 - 4(\mathbf{U}_2\mathbf{U}_2^\top - \mathbf{M})\mathbf{U}_2\|_F \\ &\leq 4\|\mathbf{U}_1\mathbf{U}_1^\top\mathbf{U}_1 - \mathbf{U}_2\mathbf{U}_2^\top\mathbf{U}_2\|_F + 4\|\mathbf{M}(\mathbf{U}_1 - \mathbf{U}_2)\|_F \\ &= 4\|\mathbf{U}_1\mathbf{U}_1^\top(\mathbf{U}_1 - \mathbf{U}_2) + \mathbf{U}_1(\mathbf{U}_1 - \mathbf{U}_2)^\top\mathbf{U}_2 + (\mathbf{U}_1 - \mathbf{U}_2)\mathbf{U}_2^\top\mathbf{U}_2\|_F + 4\|\mathbf{M}(\mathbf{U}_1 - \mathbf{U}_2)\|_F \\ &\leq 12 \max\{\|\mathbf{U}_1\|^2, \|\mathbf{U}_2\|^2\} \|\mathbf{U}_1 - \mathbf{U}_2\|_F + 4\|\mathbf{M}\| \|\mathbf{U}_1 - \mathbf{U}_2\|_F \\ &\leq 16 \max\{\Gamma^2, \|\mathbf{M}\|\} \|\mathbf{U}_1 - \mathbf{U}_2\|_F \end{aligned}$$

□

Lemma B.3. (restatement of Lemma 4.3) *Within the region $\mathcal{D} = \{\mathbf{U} \mid \sigma_{\min}(\mathbf{X}^\top\mathbf{U}) \geq \gamma\}$, then we have function $f(\mathbf{U}) = \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F^2$ satisfying:*

$$\|\nabla f(\mathbf{U})\|_F^2 \geq \alpha f(\mathbf{U})$$

where constant $\alpha = 4\gamma^2$.

Proof. Inside region \mathcal{D} , recall we denote $\mathbf{W}\mathbf{D}\mathbf{W}^\top = \text{SVD}(\mathbf{U}\mathbf{U}^\top)$, thus we have:

$$\begin{aligned} \|\nabla f(\mathbf{U})\|_F^2 &= 16\|(\mathbf{U}\mathbf{U}^\top - \mathbf{M})\mathbf{U}\|_F^2 \\ &= 16[\|\mathcal{P}_{\mathbf{W}}(\mathbf{U}\mathbf{U}^\top - \mathbf{M})\mathbf{U}\|_F^2 + \|\mathcal{P}_{\mathbf{W}_\perp}(\mathbf{U}\mathbf{U}^\top - \mathbf{M})\mathbf{U}\|_F^2] \\ &\geq 16[\sigma_{\min}(\mathbf{D})\|\mathcal{P}_{\mathbf{W}}(\mathbf{U}\mathbf{U}^\top - \mathbf{M})\mathcal{P}_{\mathbf{W}}\|_F^2 + \|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}\mathbf{U}\|_F^2] \\ &\geq 16[\sigma_{\min}(\mathbf{D})\|\mathbf{U}\mathbf{U}^\top - \mathcal{P}_{\mathbf{W}}\mathbf{M}\mathcal{P}_{\mathbf{W}}\|_F^2 + \|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}\mathbf{U}\|_F^2] \end{aligned}$$

On the other hand, we have:

$$\begin{aligned} \|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}\mathbf{U}\|_F^2 &= \|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{X}\mathbf{X}^\top\mathbf{U}\|_F^2 \geq \sigma_{\min}^2(\mathbf{X}^\top\mathbf{U})\|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{X}\mathbf{X}^\top\|_F^2 \\ &= \sigma_{\min}^2(\mathbf{X}^\top\mathbf{U})\text{tr}(\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}^2\mathcal{P}_{\mathbf{W}_\perp}) = \sigma_{\min}^2(\mathbf{X}^\top\mathbf{U})\|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}\|_F^2 \end{aligned}$$

and

$$\sigma_{\min}(\mathbf{D}) = \lambda_{\min}(\mathbf{U}^\top \mathbf{U}) \geq \lambda_{\min}(\mathbf{U}^\top \mathcal{P}_{\mathbf{X}} \mathbf{U}) = \sigma_{\min}^2(\mathbf{X}^\top \mathbf{U})$$

Therefore, combine all above, we have:

$$\begin{aligned} \|\nabla f(\mathbf{U})\|_{\mathbb{F}}^2 &\geq 16\sigma_{\min}^2(\mathbf{X}^\top \mathbf{U})[\|\mathbf{U}\mathbf{U}^\top - \mathcal{P}_{\mathbf{W}}\mathbf{M}\mathcal{P}_{\mathbf{W}}\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}\|_{\mathbb{F}}^2] \\ &\geq 4\sigma_{\min}^2(\mathbf{X}^\top \mathbf{U})[\|\mathbf{U}\mathbf{U}^\top - \mathcal{P}_{\mathbf{W}}\mathbf{M}\mathcal{P}_{\mathbf{W}}\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}\mathcal{P}_{\mathbf{W}}\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\mathbf{W}}\mathbf{M}\mathcal{P}_{\mathbf{W}_\perp}\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\mathbf{W}_\perp}\mathbf{M}\mathcal{P}_{\mathbf{W}_\perp}\|_{\mathbb{F}}^2] \\ &= 4\sigma_{\min}^2(\mathbf{X}^\top \mathbf{U})\|\mathbf{U}\mathbf{U}^\top - \mathbf{M}\|_{\mathbb{F}}^2 \end{aligned}$$

□

Next, we show as long as we are in some Frobenious ball around optimum, then we have good spectral property over \mathbf{U} which guarantees the preconditions for Lemma B.2 and Lemma B.3.

Lemma B.4. (restatement of Lemma 4.4) *Within the region $\mathcal{D} = \{\mathbf{U} \mid \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_{\mathbb{F}} \leq \frac{1}{10}\sigma_k(\mathbf{M})\}$, we have:*

$$\|\mathbf{U}\| \leq \sqrt{2\|\mathbf{M}\|}, \quad \sigma_{\min}(\mathbf{X}^\top \mathbf{U}) \geq \sqrt{\sigma_k(\mathbf{M})/2}$$

Proof. For spectral norm of \mathbf{U} , we have:

$$\|\mathbf{U}\|^2 \leq \|\mathbf{M}\| + \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\| \leq \|\mathbf{M}\| + \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_{\mathbb{F}} \leq 2\|\mathbf{M}\|$$

For the minimum singular value of $\mathbf{U}^\top \mathbf{U}$, we have:

$$\begin{aligned} \sigma_{\min}(\mathbf{U}^\top \mathbf{U}) = \sigma_k(\mathbf{U}\mathbf{U}^\top) &\geq \sigma_k(\mathbf{M}) - \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\| \\ &\geq \sigma_k(\mathbf{U}\mathbf{U}^\top) \geq \sigma_k(\mathbf{M}) - \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_{\mathbb{F}} \geq \frac{9}{10}\sigma_k(\mathbf{M}) \end{aligned}$$

On the other hand, we have:

$$\begin{aligned} \frac{9}{10}\sigma_k(\mathbf{M})\|\mathbf{X}_\perp \mathbf{W}\|^2 &\leq \sigma_{\min}(\mathbf{D})\|\mathbf{X}_\perp \mathbf{W}\|^2 \leq \|\mathbf{X}_\perp^\top \mathbf{W} \Sigma \mathbf{W}^\top \mathbf{X}_\perp\| \\ &\leq \|\mathbf{X}_\perp^\top \mathbf{U}\mathbf{U}^\top \mathbf{X}_\perp\|_{\mathbb{F}} = \|\mathcal{P}_{\mathbf{X}_\perp}(\mathbf{M} - \mathbf{U}\mathbf{U}^\top)\mathcal{P}_{\mathbf{X}_\perp}\|_{\mathbb{F}} \\ &\leq \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_{\mathbb{F}} \leq \frac{1}{10}\sigma_k(\mathbf{M}) \end{aligned}$$

Let the principal angle between \mathbf{X} and \mathbf{W} to be θ . This gives $\sin^2 \theta = \|\mathbf{X}_\perp^\top \mathbf{W}\|^2 \leq \frac{1}{9}$. Thus $\cos^2 \theta = \sigma_{\min}^2(\mathbf{X}^\top \mathbf{W}) \geq \frac{8}{9}$. Therefore:

$$\sigma_{\min}^2(\mathbf{X}^\top \mathbf{U}) \geq \sigma_{\min}^2(\mathbf{X}^\top \mathbf{W})\sigma_{\min}(\mathbf{U}^\top \mathbf{U}) \geq \sigma_k(\mathbf{M})/2$$

□

B.2 Proof of Theorem B.1

Now, we are ready for our key theorem. By Lemma B.2, Lemma B.3, and Lemma B.4, we already know the function has good property locally in the region $\mathcal{D} = \{\mathbf{U} \mid \|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F \leq \frac{1}{10}\sigma_k(\mathbf{M})\}$ which alludes linear convergence. Then, the work remains and also the most challenging part is to prove that once we initialize inside this region, our algorithm will guarantee \mathbf{U} never leave this region with high probability even with relatively large stepsize. The requirement for tight sample complexity and near optimal runtime makes it more challenging, and require us to further control the incoherence of \mathbf{U}_t over all iterates in addition to the distance $\|\mathbf{M} - \mathbf{U}\mathbf{U}^\top\|_F$.

Following is our formal proof.

Proof of Theorem B.1. Define event $\mathfrak{E}_t = \{\forall \tau \leq t, f(\mathbf{U}_\tau) \leq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2, \max_i g_i(\mathbf{U}_\tau) \leq \frac{20\mu k \kappa^2}{d}\}$. Theorem B.1 is equivalent to prove event \mathfrak{E}_T happens with high probability. The proof achieves this by constructing two supermartingales for $f(\mathbf{U}_t)1_{\mathfrak{E}_t}$ and $g_i(\mathbf{U}_t)1_{\mathfrak{E}_t}$ (where $1_{(\cdot)}$ denote indicator function), applies concentration argument.

The proofs follow the structure of:

1. The constructions of supermartingales
2. Their probability 1 bound and variance bound in order to apply Azuma-Bernstein inequality
3. Final combination of concentration results to conclude the proof

First, let filtration $\mathfrak{F}_t = \sigma\{SG(\mathbf{U}_0), \dots, SG(\mathbf{U}_{t-1})\}$ where $\sigma\{\cdot\}$ denotes the sigma field. Note by definition of \mathfrak{E}_t , we have $\mathfrak{E}_t \subset \mathfrak{F}_t$. Also $\mathfrak{E}_{t+1} \subset \mathfrak{E}_t$, and thus $1_{\mathfrak{E}_{t+1}} \leq 1_{\mathfrak{E}_t}$. Note \mathfrak{E}_t denotes the event which up to time t , \mathbf{U}_τ always stay in a local region which both close to \mathbf{M} and incoherent.

By Lemma B.4, we immediately know that conditioned on \mathfrak{E}_t , we have $\|\mathbf{U}_t\| \leq \sqrt{2}$, $\sigma_{\min}(\mathbf{X}^\top \mathbf{U}_t) \geq 1/\sqrt{2\kappa}$ and $\sigma_{\min}(\mathbf{U}_t^\top \mathbf{U}_t) \geq 1/2\kappa$. We will use this fact throughout the proof.

Construction of supermartingale G : Since $g_i(\mathbf{U}) = \mathbf{e}_i^\top \mathbf{U}\mathbf{U}^\top \mathbf{e}_i$ is a quadratic function, we know for any change $\Delta\mathbf{U}$, we have:

$$g_i(\mathbf{U} + \Delta\mathbf{U}) = g_i(\mathbf{U}) + 2\mathbf{e}_i^\top (\Delta\mathbf{U})\mathbf{U}^\top \mathbf{e}_i + \left\| \mathbf{e}_i^\top \Delta\mathbf{U} \right\|^2$$

We know for any $l \in [d]$:

$$\begin{aligned} \mathbb{E} \left\| \mathbf{e}_i^\top SG(\mathbf{U}) \right\|^2 1_{\mathfrak{E}_t} &\leq \mathbb{E} 16d^4 \delta_{il} (\mathbf{u}_i^\top \mathbf{u}_j - \mathbf{M}_{ij})^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^2 1_{\mathfrak{E}_t} \\ &= 16d^2 \left\| \mathbf{e}_i^\top (\mathbf{U}\mathbf{U}^\top - \mathbf{M}) \right\|^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^2 1_{\mathfrak{E}_t} \leq O(\mu^2 k^2 \kappa^4) 1_{\mathfrak{E}_t} \end{aligned}$$

Therefore, by update Eq.(8), and $\mathbb{E}SG(\mathbf{U}) = \nabla f(\mathbf{U}) = 4(\mathbf{U}\mathbf{U}^\top - \mathbf{M})\mathbf{U}$, we know:

$$\begin{aligned}
& \mathbb{E}[g_i(\mathbf{U}_{t+1})\mathbf{1}_{\mathfrak{E}_t}|\mathfrak{F}_t] \\
&= [g_i(\mathbf{U}_t) - 2\eta\mathbf{e}_i^\top [\mathbb{E}SG(\mathbf{U}_t)]\mathbf{U}_t^\top \mathbf{e}_i + \frac{\eta^2}{2}\mathbb{E}\left\|\mathbf{e}_i^\top SG(\mathbf{U}_t)\right\|^2]\mathbf{1}_{\mathfrak{E}_t} \\
&= [\text{tr}(\mathbf{U}_t^\top \mathbf{e}_i \mathbf{e}_i^\top [\mathbf{I} - 8\eta(\mathbf{U}_t\mathbf{U}_t^\top - \mathbf{M})]\mathbf{U}_t) + \frac{\eta^2}{2}\mathbb{E}\left\|\mathbf{e}_i^\top SG(\mathbf{U}_t)\right\|^2]\mathbf{1}_{\mathfrak{E}_t} \\
&= [\text{tr}(\mathbf{U}_t^\top \mathbf{e}_i \mathbf{e}_i^\top \mathbf{U}_t(\mathbf{I} - 8\eta\mathbf{U}_t^\top \mathbf{U}_t)) + 8\eta\text{tr}(\mathbf{U}_t^\top \mathbf{e}_i \mathbf{e}_i^\top \mathbf{M}\mathbf{U}_t) + \eta^2 O(\mu^2 k^2 \kappa^4)]\mathbf{1}_{\mathfrak{E}_t} \\
&\leq [(1 - 8\eta\sigma_{\min}(\mathbf{U}_t^\top \mathbf{U}_t))g_i(\mathbf{U}_t) + 8\eta\text{tr}(\mathbf{U}_t^\top \mathbf{e}_i \mathbf{e}_i^\top \mathbf{M}\mathbf{U}_t) + \eta^2 O(\mu^2 k^2 \kappa^4)]\mathbf{1}_{\mathfrak{E}_t} \\
&\leq [(1 - \frac{4\eta}{\kappa})g_i(\mathbf{U}_t) + 16\sqrt{10}\frac{\eta\mu k\kappa}{d} + \eta^2 O(\mu^2 k^2 \kappa^4)]\mathbf{1}_{\mathfrak{E}_t} \\
&\leq [(1 - \frac{4\eta}{\kappa})g_i(\mathbf{U}_t) + 60\frac{\eta\mu k\kappa}{d}]\mathbf{1}_{\mathfrak{E}_t}
\end{aligned}$$

The last step is true by choosing constant c in learning rate η to be small enough.

Let $G_{it} = (1 - \frac{4\eta}{\kappa})^{-t}(g_i(\mathbf{U}_t)\mathbf{1}_{\mathfrak{E}_{t-1}} - 15\frac{\mu k\kappa^2}{d})$. This gives:

$$\mathbb{E}G_{i(t+1)} \leq (1 - \frac{4\eta}{\kappa})^{-t}(g_i(\mathbf{U}_t)\mathbf{1}_{\mathfrak{E}_t} - 15\frac{\mu k\kappa^2}{d}) \leq G_{it}$$

That is G_{it} is supermartingale.

Probability 1 bound for G : We also know

$$\begin{aligned}
G_{it} - \mathbb{E}[G_{it}|\mathfrak{F}_{t-1}] &= (1 - \frac{4\eta}{\kappa})^{-t} \left[-\eta\mathbf{e}_i^\top [SG(\mathbf{U}_t) - \mathbb{E}SG(\mathbf{U}_t)]\mathbf{U}_t^\top \mathbf{e}_i \right. \\
&\quad \left. + \frac{\eta^2}{2} [\|\mathbf{e}_i^\top SG(\mathbf{U}_t)\|^2 - \mathbb{E}\|\mathbf{e}_i^\top SG(\mathbf{U}_t)\|^2] \right] \mathbf{1}_{\mathfrak{E}_{t-1}} \tag{9}
\end{aligned}$$

Since when sample (i, j) entry of matrix \mathbf{M} , for any $l \in [d]$, we have:

$$\begin{aligned}
& \mathbf{e}_l^\top [SG(\mathbf{U}_t)]\mathbf{U}_t^\top \mathbf{e}_l \cdot \mathbf{1}_{\mathfrak{E}_{t-1}} = O(1)\text{tr}(\mathbf{U}_t^\top \mathbf{e}_l \mathbf{e}_l^\top SG(\mathbf{U}_t))\mathbf{1}_{\mathfrak{E}_{t-1}} \\
&= O(1)d^2(\mathbf{U}\mathbf{U}^\top - \mathbf{M})_{ij}\text{tr}[\mathbf{U}^\top \mathbf{e}_l \mathbf{e}_l^\top (\mathbf{e}_i \mathbf{u}_j^\top + \mathbf{e}_j \mathbf{u}_i^\top)]\mathbf{1}_{\mathfrak{E}_{t-1}} \\
&\leq O(1)d^2 \left\| \mathbf{U}\mathbf{U}^\top - \mathbf{M} \right\|_\infty \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^2 \mathbf{1}_{\mathfrak{E}_{t-1}} \leq O(\mu^2 k^2 \kappa^4)\mathbf{1}_{\mathfrak{E}_{t-1}}
\end{aligned}$$

and

$$\begin{aligned}
& \|\mathbf{e}_l^\top SG(\mathbf{U}_t)\|^2 \mathbf{1}_{\mathfrak{E}_{t-1}} = O(1) \left\| \mathbf{e}_l^\top SG(\mathbf{U}_t) \right\|^2 \mathbf{1}_{\mathfrak{E}_{t-1}} \\
&= O(1)d^4(\mathbf{U}\mathbf{U}^\top - \mathbf{M})_{ij}^2 \left\| \mathbf{e}_l^\top (\mathbf{e}_i \mathbf{u}_j^\top + \mathbf{e}_j \mathbf{u}_i^\top) \right\|^2 \mathbf{1}_{\mathfrak{E}_{t-1}} \\
&\leq O(1)d^4 \left\| \mathbf{U}\mathbf{U}^\top - \mathbf{M} \right\|_\infty^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^2 \mathbf{1}_{\mathfrak{E}_{t-1}} \leq O(\mu^3 d k^3 \kappa^6)\mathbf{1}_{\mathfrak{E}_{t-1}}
\end{aligned}$$

Therefore, by Eq.(9), we have with probability 1:

$$|G_{it} - \mathbb{E}[G_{it}|\mathfrak{F}_{t-1}]| \leq (1 - \frac{4\eta}{\kappa})^{-t} \eta O(\mu^2 k^2 \kappa^4)\mathbf{1}_{\mathfrak{E}_{t-1}} \tag{10}$$

Variance bound for G : For any $l \in [d]$, we also know

$$\begin{aligned}
& \text{Var}(\mathbf{e}_l^\top [SG(\mathbf{U}_t)] \mathbf{U}_t^\top \mathbf{e}_l \cdot \mathbf{1}_{\mathfrak{E}_{t-1}} | \mathfrak{F}_{t-1}) \leq \mathbb{E}[\langle \nabla g_l(\mathbf{U}_t), SG(\mathbf{U}_t) \rangle^2 \mathbf{1}_{\mathfrak{E}_{t-1}} | \mathfrak{F}_{t-1}] \\
& = O(1) \frac{1}{d^2} \sum_{ij} d^4 (\mathbf{U} \mathbf{U}^\top - \mathbf{M})_{ij}^2 \text{tr}[\mathbf{U}^\top \mathbf{e}_l \mathbf{e}_l^\top (\mathbf{e}_i \mathbf{u}_j^\top + \mathbf{e}_j \mathbf{u}_i^\top)]^2 \mathbf{1}_{\mathfrak{E}_{t-1}} \\
& \leq O(1) d^2 \sum_j (\mathbf{U} \mathbf{U}^\top - \mathbf{M})_{lj}^2 \text{tr}[\mathbf{U}^\top \mathbf{e}_l \mathbf{u}_j^\top]^2 \mathbf{1}_{\mathfrak{E}_{t-1}} \\
& \leq O(1) d^2 \left\| \mathbf{e}_l^\top (\mathbf{U} \mathbf{U}^\top - \mathbf{M}) \right\|^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^4 \mathbf{1}_{\mathfrak{E}_{t-1}} \leq O\left(\frac{\mu^3 k^3 \kappa^6}{d}\right) \mathbf{1}_{\mathfrak{E}_{t-1}}
\end{aligned}$$

and

$$\begin{aligned}
& \text{Var}(\|\mathbf{e}_l^\top SG(\mathbf{U})\|^2 \mathbf{1}_{\mathfrak{E}_{t-1}} | \mathfrak{F}_{t-1}) \leq \mathbb{E}[\nabla^2 g_k(SG(\mathbf{U}_t), SG(\mathbf{U}_t))^2 \mathbf{1}_{\mathfrak{E}_{t-1}} | \mathfrak{F}_{t-1}] \\
& = O(1) \frac{1}{d^2} \sum_{ij} d^8 (\mathbf{U} \mathbf{U}^\top - \mathbf{M})_{ij}^4 \left\| \mathbf{e}_k^\top (\mathbf{e}_i \mathbf{u}_j^\top + \mathbf{e}_j \mathbf{u}_i^\top) \right\|^4 \mathbf{1}_{\mathfrak{E}_{t-1}} \\
& \leq O(1) d^6 \sum_j (\mathbf{U} \mathbf{U}^\top - \mathbf{M})_{kj}^4 \|\mathbf{u}_j\|^4 \mathbf{1}_{\mathfrak{E}_{t-1}} \\
& \leq O(1) d^6 \left\| \mathbf{U} \mathbf{U}^\top - \mathbf{M} \right\|_\infty^2 \left\| \mathbf{e}_k^\top (\mathbf{U} \mathbf{U}^\top - \mathbf{M}) \right\|^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^4 \mathbf{1}_{\mathfrak{E}_{t-1}} \leq O(\mu^5 d k^5 \kappa^{10}) \mathbf{1}_{\mathfrak{E}_{t-1}}
\end{aligned}$$

Therefore, by Eq.(9), we have

$$\text{Var}(G_{it} | \mathfrak{F}_{t-1}) \leq \left(1 - \frac{4\eta}{\kappa}\right)^{-2t} \eta^2 O\left(\frac{\mu^3 k^3 \kappa^6}{d}\right) \mathbf{1}_{\mathfrak{E}_{t-1}} \quad (11)$$

Bernstein's inequality for G : Let $\sigma^2 = \sum_{\tau=1}^t \text{Var}(G_{i\tau} | \mathfrak{F}_{\tau-1})$, and R satisfies, with probability 1 that $|G_{i\tau} - \mathbb{E}[G_{i\tau} | \mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Then By standard Bernstein concentration inequality, we know:

$$P(G_{it} \geq G_{i0} + s) \leq \exp\left(\frac{s^2/2}{\sigma^2 + Rs/3}\right)$$

Since $G_{i0} = g_i(\mathbf{U}_0) - 15 \frac{\mu k \kappa^2}{d}$, let $s' = O(1)(1 - \frac{4\eta}{\kappa})^t [\sqrt{\sigma^2 \log d} + R \log d]$, we know

$$P\left(g_i(\mathbf{U}_t) \mathbf{1}_{\mathfrak{E}_{t-1}} \geq 15 \frac{\mu k \kappa^2}{d} + \left(1 - \frac{4\eta}{\kappa}\right)^t \left(g_i(\mathbf{U}_0) - 15 \frac{\mu k \kappa^2}{d}\right) + s'\right) \leq \frac{1}{2d^{11}}$$

By Eq.(10), we know $R = (1 - \frac{4\eta}{\kappa})^{-t} \eta O(\mu^2 k^2 \kappa^4)$ satisfies that $|G_{i\tau} - \mathbb{E}[G_{i\tau} | \mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Also by Eq. (11), we have:

$$\left(1 - \frac{4\eta}{\kappa}\right)^t \sqrt{\sigma^2 \log d} \leq \eta O\left(\sqrt{\frac{\mu^3 k^3 \kappa^6 \log d}{d}}\right) \sqrt{\sum_{\tau=1}^t \left(1 - \frac{4\eta}{\kappa}\right)^{2t-2\tau}} \leq \sqrt{\eta} O\left(\sqrt{\frac{\mu^3 k^3 \kappa^7 \log d}{d}}\right)$$

by $\eta < \frac{c}{\mu d k \kappa^3 \log d}$ and choosing c to be small enough, we have:

$$s' = \sqrt{\eta} O\left(\sqrt{\frac{\mu^3 k^3 \kappa^7 \log d}{d}}\right) + \eta O(\mu^2 k^2 \kappa^4 \log d) \leq \frac{\mu k \kappa^2}{d}$$

Since initialization gives $\max_i g_i(\mathbf{U}_0) \leq \frac{10\mu k \kappa^2}{d}$, therefore:

$$P(g_i(\mathbf{U}_t)1_{\mathfrak{E}_{t-1}} \geq 20\frac{\mu k \kappa^2}{d}) \leq \frac{1}{2d^{11}}$$

That is equivalent to:

$$P(\mathfrak{E}_{t-1} \cap \{g_i(\mathbf{U}_t) \geq 20\frac{\mu k \kappa^2}{d}\}) \leq \frac{1}{2d^{11}} \quad (12)$$

Construction of supermartingale F : On the other hand, we also have

$$\begin{aligned} \mathbb{E}\|SG(\mathbf{U}_t)\|_F^2 1_{\mathfrak{E}_t} &\leq \mathbb{E}16d^4(\mathbf{u}_i^\top \mathbf{u}_j - \mathbf{M}_{ij})^2 \max_i \|\mathbf{e}_i^\top \mathbf{U}\|^2 1_{\mathfrak{E}_t} \\ &\leq 16d^2 \|\mathbf{U}_t \mathbf{U}_t^\top - \mathbf{M}\|_F^2 \max_i \|\mathbf{e}_i^\top \mathbf{U}_t\|^2 1_{\mathfrak{E}_t} \leq O(\mu d k \kappa^2) f(\mathbf{U}_t) 1_{\mathfrak{E}_t} \end{aligned}$$

Therefore, by update function Eq.(8),

$$\begin{aligned} &\mathbb{E}[f(\mathbf{U}_{t+1})1_{\mathfrak{E}_t} | \mathfrak{F}_t] \\ &\leq [f(\mathbf{U}_t) - \mathbb{E}\langle \nabla f(\mathbf{U}_t), \eta SG(\mathbf{U}_t) \rangle + \eta^2 \mathbb{E}\|SG(\mathbf{U}_t)\|_F^2] 1_{\mathfrak{E}_t} \\ &= [f(\mathbf{U}_t) - \eta \|\nabla f(\mathbf{U}_t)\|_F^2 + \eta^2 \mathbb{E}\|SG(\mathbf{U}_t)\|_F^2] 1_{\mathfrak{E}_t} \\ &\leq [(1 - \frac{2\eta}{\kappa})f(\mathbf{U}_t) + \eta^2 O(\mu d k \kappa^2) f(\mathbf{U}_t)] 1_{\mathfrak{E}_t} \\ &\leq (1 - \frac{\eta}{\kappa})f(\mathbf{U}_t) 1_{\mathfrak{E}_t} \end{aligned}$$

Let $F_t = (1 - \frac{\eta}{\kappa})^{-t} f(\mathbf{U}_t) 1_{\mathfrak{E}_{t-1}}$, we know F_t is also a supermartingale.

Probability 1 bound for F : With probability 1, we also have:

$$\begin{aligned} F_t - \mathbb{E}[F_t | \mathfrak{F}_{t-1}] &= (1 - \frac{\eta}{\kappa})^{-t} [-\eta \langle \nabla f(\mathbf{U}_t), SG(\mathbf{U}_t) - \mathbb{E}SG(\mathbf{U}_t) \rangle \\ &\quad + \frac{\eta^2}{2} (\nabla^2 f(\zeta_t)(SG(\mathbf{U}_t), SG(\mathbf{U}_t)) - \mathbb{E}\nabla^2 f(\zeta_t)(SG(\mathbf{U}_t), SG(\mathbf{U}_t)))] 1_{\mathfrak{E}_{t-1}} \quad (13) \end{aligned}$$

where ζ_t depends on $SG(\mathbf{U}_t)$.

First, recall we denote $\text{SVD}(\mathbf{M}) = \mathbf{X}\mathbf{S}\mathbf{X}^\top$, and $\text{SVD}(\mathbf{U}\mathbf{U}^\top) = \mathbf{W}\mathbf{D}\mathbf{W}^\top$, and observe that:

$$\begin{aligned} &\|\mathbf{U}\mathbf{U}^\top - \mathbf{M}\|_\infty 1_{\mathfrak{E}_{t-1}} = \max_{ij} |\text{tr}(\mathbf{e}_i^\top (\mathbf{U}\mathbf{U}^\top - \mathbf{M}) \mathbf{e}_j)| 1_{\mathfrak{E}_{t-1}} \\ &= \max_{ij} |\text{tr}(\mathbf{e}_i^\top (\mathcal{P}_{\mathbf{X}} + \mathcal{P}_{\mathbf{X}_\perp})(\mathbf{U}\mathbf{U}^\top - \mathbf{M}) \mathbf{e}_j)| 1_{\mathfrak{E}_{t-1}} \\ &\leq \max_{ij} |\text{tr}(\mathbf{e}_i^\top \mathcal{P}_{\mathbf{X}}(\mathbf{U}\mathbf{U}^\top - \mathbf{M}) \mathbf{e}_j)| 1_{\mathfrak{E}_{t-1}} + \max_{ij} |\text{tr}(\mathbf{e}_i^\top \mathcal{P}_{\mathbf{X}_\perp} \mathbf{U}\mathbf{U}^\top \mathbf{e}_j)| 1_{\mathfrak{E}_{t-1}} \\ &\leq \max_i \|\mathbf{e}_i^\top \mathbf{X}\| \|\mathbf{U}\mathbf{U}^\top - \mathbf{M}\|_F 1_{\mathfrak{E}_{t-1}} + \max_i \|\mathbf{e}_i^\top \mathbf{W}\| \|\mathbf{U}\mathbf{U}^\top - \mathbf{M}\|_F 1_{\mathfrak{E}_{t-1}} \\ &\leq O(\sqrt{\frac{\mu k \kappa^3}{d}}) \sqrt{f(\mathbf{U}_t)} \end{aligned}$$

Then, when sample (i, j) entry of matrix \mathbf{M} , we have:

$$\begin{aligned} \langle \nabla f(\mathbf{U}_t), SG(\mathbf{U}_t) \rangle 1_{\mathbf{e}_{t-1}} &\leq O(1) \|\nabla f(\mathbf{U}_t)\|_F \|SG(\mathbf{U}_t)\|_F 1_{\mathbf{e}_{t-1}} \\ &\leq O(1) d^2 \sqrt{f(\mathbf{U}_t)} (\mathbf{U}\mathbf{U}^\top - \mathbf{M})_{ij} \left\| \mathbf{e}_i \mathbf{u}_j^\top + \mathbf{e}_j \mathbf{u}_i^\top \right\|_F^2 1_{\mathbf{e}_{t-1}} \\ &\leq O(1) d^2 \sqrt{f(\mathbf{U}_t)} \left\| \mathbf{U}\mathbf{U}^\top - \mathbf{M} \right\|_\infty \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\| 1_{\mathbf{e}_{t-1}} \leq O(\mu d k \kappa^{2.5}) f(\mathbf{U}_t) 1_{\mathbf{e}_{t-1}} \end{aligned}$$

and

$$\begin{aligned} \nabla^2 f(\zeta_t)(SG(\mathbf{U}_t), SG(\mathbf{U}_t)) 1_{\mathbf{e}_{t-1}} &\leq O(1) \|SG(\mathbf{U}_t)\|_F^2 \\ &\leq O(1) d^4 \left\| \mathbf{U}\mathbf{U}^\top - \mathbf{M} \right\|_\infty^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^2 1_{\mathbf{e}_{t-1}} \leq O(\mu^2 d^2 k^2 \kappa^5) f(\mathbf{U}_t) 1_{\mathbf{e}_{t-1}} \end{aligned}$$

Therefore, by decomposition Eq.(13), we have with probability 1:

$$|F_t - \mathbb{E}[F_t | \mathfrak{F}_{t-1}]| \leq (1 - \frac{\eta}{\kappa})^{-t} \eta O(\mu d k \kappa^{2.5}) f(\mathbf{U}_{t-1}) 1_{\mathbf{e}_{t-1}} \leq (1 - \frac{\eta}{\kappa})^{-t} (1 - \frac{\eta}{2\kappa})^t \eta O(\mu d k \kappa^{0.5}) 1_{\mathbf{e}_{t-1}} \quad (14)$$

Variance bound for F : We also know

$$\begin{aligned} \text{Var}(\langle \nabla f(\mathbf{U}_t), SG(\mathbf{U}_t) \rangle 1_{\mathbf{e}_{t-1}} | \mathfrak{F}_{t-1}) &\leq \mathbb{E}[\langle \nabla f(\mathbf{U}_t), SG(\mathbf{U}_t) \rangle^2 1_{\mathbf{e}_{t-1}} | \mathfrak{F}_{t-1}] \\ &\leq \|\nabla f(\mathbf{U}_t)\|_F^2 \mathbb{E} \|SG(\mathbf{U}_t)\|_F^2 1_{\mathbf{e}_{t-1}} \leq O(1) d^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^2 f^2(\mathbf{U}_{t-1}) 1_{\mathbf{e}_{t-1}} \\ &\leq O(\mu d k \kappa^2) f^2(\mathbf{U}_{t-1}) 1_{\mathbf{e}_{t-1}} \end{aligned}$$

and

$$\begin{aligned} \text{Var}(\nabla^2 f(\zeta_t)(SG(\mathbf{U}_t), SG(\mathbf{U}_t)) 1_{\mathbf{e}_{t-1}} | \mathfrak{F}_{t-1}) &\leq \mathbb{E}[\nabla^2 f(\zeta_t)(SG(\mathbf{U}_t), SG(\mathbf{U}_t))^2 1_{\mathbf{e}_{t-1}} | \mathfrak{F}_{t-1}] \\ &\leq O(1) \mathbb{E} \|SG(\mathbf{U}_t)\|_F^4 = O(1) \mathbb{E} d^8 (\mathbf{U}\mathbf{U}^\top - \mathbf{M})_{ij}^4 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^4 1_{\mathbf{e}_{t-1}} \\ &\leq O(1) d^6 \left\| \mathbf{U}\mathbf{U}^\top - \mathbf{M} \right\|_\infty^2 \left\| \mathbf{U}\mathbf{U}^\top - \mathbf{M} \right\|_F^2 \max_i \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^4 1_{\mathbf{e}_{t-1}} \\ &\leq O(\mu^3 d^3 k^3 \kappa^7) f^2(\mathbf{U}_{t-1}) 1_{\mathbf{e}_{t-1}} \end{aligned}$$

Therefore, by decomposition Eq.(13), we have:

$$\text{Var}(F_t | \mathfrak{F}_{t-1}) \leq (1 - \frac{\eta}{\kappa})^{-2t} \eta^2 O(\mu d k \kappa^2) f^2(\mathbf{U}_{t-1}) 1_{\mathbf{e}_{t-1}} \leq (1 - \frac{\eta}{\kappa})^{-2t} (1 - \frac{\eta}{2\kappa})^{2t} \eta^2 O(\frac{\mu d k}{\kappa^2}) 1_{\mathbf{e}_{t-1}} \quad (15)$$

Bernstein's inequality for F : Let $\sigma^2 = \sum_{\tau=1}^t \text{Var}(F_\tau | \mathfrak{F}_{\tau-1})$, and R satisfies, with probability 1 that $|F_\tau - \mathbb{E}[F_\tau | \mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Then By standard Bernstein concentration inequality, we know:

$$P(F_t \geq F_0 + s) \leq \exp\left(\frac{s^2/2}{\sigma^2 + Rs/3}\right)$$

Let $s' = O(1)(1 - \frac{\eta}{\kappa})^t [\sqrt{\sigma^2 \log d} + R \log d]$, this gives:

$$P(f(\mathbf{U}_t) 1_{\mathbf{e}_{t-1}} \geq (1 - \frac{\eta}{\kappa})^t f(\mathbf{U}_0) + s') \leq \frac{1}{2d^{10}}$$

By Eq.(14), we know $R = (1 - \frac{\eta}{\kappa})^{-t}(1 - \frac{\eta}{2\kappa})^t \eta O(\mu dk \kappa^{0.5})$ satisfies that $|F_\tau - \mathbb{E}[F_\tau | \mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Also by Eq. (15), we have:

$$\begin{aligned} (1 - \frac{\eta}{\kappa})^t \sqrt{\sigma^2 \log d} &\leq \eta O(\sqrt{\frac{\mu dk \log d}{\kappa^2}}) \sqrt{\sum_{\tau=1}^t (1 - \frac{\eta}{\kappa})^{2t-2\tau} (1 - \frac{\eta}{2\kappa})^{2\tau}} \\ &\leq (1 - \frac{\eta}{2\kappa})^t \eta O(\sqrt{\frac{\mu dk \log d}{\kappa^2}}) \sqrt{\sum_{\tau=1}^t (1 - \frac{\eta}{\kappa})^{2t-2\tau} (1 - \frac{\eta}{2\kappa})^{2\tau-2t}} \leq (1 - \frac{\eta}{2\kappa})^t \sqrt{\eta} O(\sqrt{\frac{\mu dk \log d}{\kappa}}) \end{aligned}$$

by $\eta < \frac{c}{\mu dk \kappa^3 \log d}$ and choosing c to be small enough, we have:

$$s' = (1 - \frac{\eta}{2\kappa})^t [\sqrt{\eta} O(\sqrt{\frac{\mu dk \log d}{\kappa}}) + \eta O(\mu dk \kappa^{0.5})] \leq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{20\kappa})^2$$

Since $F_0 = f(\mathbf{U}_0) \leq (\frac{1}{20\kappa})^2$, therefore:

$$P(f(\mathbf{U}_t) \mathbf{1}_{\mathfrak{E}_{t-1}} \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2) \leq \frac{1}{2d^{10}}$$

That is equivalent to:

$$P(\mathfrak{E}_{t-1} \cap \{f(\mathbf{U}_t) \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2\}) \leq \frac{1}{2d^{10}} \quad (16)$$

Probability for event \mathfrak{E}_T : Finally, combining the concentration result for martingale G (Eq.(12)) and martingale F (Eq.(16)), we conclude:

$$\begin{aligned} P(\mathfrak{E}_{t-1} \cap \bar{\mathfrak{E}}_t) &= P\left[\mathfrak{E}_{t-1} \cap \left(\cup_i \{g_i(\mathbf{U}_t) \geq 20 \frac{\mu k \kappa^2}{d}\} \cup \{f(\mathbf{U}_t) \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2\}\right)\right] \\ &\leq \sum_{i=1}^d P(\mathfrak{E}_{t-1} \cap \{g_i(\mathbf{U}_t) \geq 20 \frac{\mu k \kappa^2}{d}\}) + P(\mathfrak{E}_{t-1} \cap \{f(\mathbf{U}_t) \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2\}) \leq \frac{1}{d^{10}} \end{aligned}$$

Since

$$P(\bar{\mathfrak{E}}_T) = \sum_{t=1}^T P(\mathfrak{E}_{t-1} \cap \bar{\mathfrak{E}}_t) \leq \frac{T}{d^{10}}$$

We finishes the proof. \square

C Proof of General Asymmetric Case

In this section, we first prove Lemma 3.2, set up the equivalence between Algorithm 2 and Algorithm 3. Then we prove the main theorem for general asymmetric matrix (Theorem 3.3). WLOG, we continue to assume $\|\mathbf{M}\| = 1$ in all proof. Also, when it's clear from the context, we use κ to specifically to represent $\kappa(\mathbf{M})$. Then $\sigma_{\min}(\mathbf{M}) = \frac{1}{\kappa}$. Also in this section, we always use $d = \max\{d_1, d_2\}$ and denote $\text{SVD}(\mathbf{M}) = \mathbf{XSY}^\top$, and $\text{SVD}(\mathbf{UV}^\top) = \mathbf{W}_U \mathbf{D} \mathbf{W}_V^\top$.

Proof of Lemma 3.2. Let us always denote the iterates in Algorithm 2 by $\mathbf{U}_t, \mathbf{V}_t$, and denote the corresponding iterates in Algorithm 3 by $\mathbf{U}'_t, \mathbf{V}'_t$ using prime version. We use induction to prove the equivalence. Assume at time t we have $\mathbf{U}_t \mathbf{V}_t^\top = \mathbf{U}'_t \mathbf{V}'_t{}^\top$. Recall in Algorithm 2, we renormalize $\mathbf{U}_t, \mathbf{V}_t$ to $\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t$, this set up the correspondence:

$$\begin{aligned}\tilde{\mathbf{U}}_t &= \mathbf{U}'_t \mathbf{R}'_U \mathbf{D}'_U{}^{-\frac{1}{2}} \mathbf{Q}'_U \mathbf{D}'_U{}^{\frac{1}{2}} \\ \tilde{\mathbf{V}}_t &= \mathbf{V}'_t \mathbf{R}'_V \mathbf{D}'_V{}^{-\frac{1}{2}} \mathbf{Q}'_V \mathbf{D}'_V{}^{\frac{1}{2}}\end{aligned}$$

Denote $\mathbf{P}'_U = \mathbf{R}'_U \mathbf{D}'_U{}^{-\frac{1}{2}} \mathbf{Q}'_U \mathbf{D}'_U{}^{\frac{1}{2}}$, and $\mathbf{P}'_V = \mathbf{V}'_t \mathbf{R}'_V \mathbf{D}'_V{}^{-\frac{1}{2}} \mathbf{Q}'_V \mathbf{D}'_V{}^{\frac{1}{2}}$. Clearly $\mathbf{P}'_U \mathbf{P}'_V{}^\top = \mathbf{I}$. Then we have $\tilde{\mathbf{U}}_t = \mathbf{U}'_t \mathbf{P}'_U$, $\tilde{\mathbf{V}}_t = \mathbf{P}'_V$ and thus:

$$\begin{aligned}& \mathbf{U}_{t+1} \mathbf{V}_{t+1}^\top \\ &= (\tilde{\mathbf{U}}_t - 2\eta d_1 d_2 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij} \mathbf{e}_i \mathbf{e}_j^\top \tilde{\mathbf{V}}_t) (\tilde{\mathbf{V}}_t - 2\eta d_1 d_2 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij} \mathbf{e}_j \mathbf{e}_i^\top \tilde{\mathbf{U}}_t)^\top \\ &= (\mathbf{U}'_t \mathbf{P}'_U - 2\eta d_1 d_2 (\mathbf{U}'_t \mathbf{V}'_t{}^\top - \mathbf{M})_{ij} \mathbf{e}_i \mathbf{e}_j^\top \mathbf{V}'_t \mathbf{P}'_V) (\mathbf{V}'_t \mathbf{P}'_V - 2\eta d_1 d_2 (\mathbf{U}'_t \mathbf{V}'_t{}^\top - \mathbf{M})_{ij} \mathbf{e}_j \mathbf{e}_i^\top \mathbf{U}'_t \mathbf{P}'_U)^\top \\ &= (\mathbf{U}'_t - 2\eta d_1 d_2 (\mathbf{U}'_t \mathbf{V}'_t{}^\top - \mathbf{M})_{ij} \mathbf{e}_i \mathbf{e}_j^\top \mathbf{V}'_t \mathbf{P}'_V \mathbf{P}'_U{}^{-1}) \\ &\quad \cdot (\mathbf{V}'_t - 2\eta d_1 d_2 (\mathbf{U}'_t \mathbf{V}'_t{}^\top - \mathbf{M})_{ij} \mathbf{e}_j \mathbf{e}_i^\top \mathbf{U}'_t \mathbf{P}'_U \mathbf{P}'_V{}^{-1})^\top \\ &= \mathbf{U}'_{t+1} \mathbf{V}'_{t+1}{}^\top\end{aligned}$$

Clearly with same initialization algorithm, we have $\mathbf{U}_0 \mathbf{V}_0^\top = \mathbf{U}'_0 \mathbf{V}'_0{}^\top$, by induction, we finish the proof. \square

Now we proceed to prove Theorem 3.3. Since Algorithm 2 and Algorithm 3 are equivalent, we will focus our analysis on Algorithm 2 which is more theoretical appealing. As for the symmetric PSD case, we first present the essential ingredient:

Theorem C.1. *Let $f(\mathbf{U}, \mathbf{V}) = \|\mathbf{U}\mathbf{V}^\top - \mathbf{M}\|_F^2$, $g_i(\mathbf{U}, \mathbf{V}) = \|\mathbf{e}_i^\top \mathbf{U}\mathbf{V}^\top\|^2$, and $h_j(\mathbf{U}, \mathbf{V}) = \|\mathbf{e}_j^\top \mathbf{V}\mathbf{U}^\top\|^2$, for $i \in [d_1]$ and $j \in [d_2]$. Suppose after initialization, we have:*

$$f(\mathbf{U}_0, \mathbf{V}_0) \leq \left(\frac{1}{20\kappa}\right)^2, \quad \max_i g_i(\mathbf{U}_0, \mathbf{V}_0) \leq \frac{10\mu k \kappa^2}{d_1}, \quad \max_j h_j(\mathbf{U}_0, \mathbf{V}_0) \leq \frac{10\mu k \kappa^2}{d_2}$$

Then, there exist some absolute constant c such that for any learning rate $\eta < \frac{c}{\mu d k \kappa^3 \log d}$, with at least $1 - \frac{T}{d^{10}}$ probability, we will have for all $t \leq T$ that:

$$f(\mathbf{U}_t, \mathbf{V}_t) \leq \left(1 - \frac{\eta}{2\kappa}\right)^t \left(\frac{1}{10\kappa}\right)^2, \quad \max_i g_i(\mathbf{U}_t, \mathbf{V}_t) \leq \frac{100\mu k \kappa^2}{d_1}, \quad \max_j h_j(\mathbf{U}_t, \mathbf{V}_t) \leq \frac{100\mu k \kappa^2}{d_2}$$

Theorem 3.3 can easily be concluded from Theorem C.1 and Lemma A.1. Theorem C.1 also provides similar guarantees as Theorem B.1 in symmetric case. However, due to the additional invariance between \mathbf{U} and \mathbf{V} , Theorem C.1 need to keep track of more complicated potential function $g_i(\mathbf{U}, \mathbf{V})$ and $h_j(\mathbf{U}, \mathbf{V})$ to control the incoherence, which makes the proof more involved.

The rest of this section all focus on proving Theorem C.1. Similar to symmetric PSD case, we also first prepare with a few lemmas about the property of objective function, and the spectral property of \mathbf{U}, \mathbf{V} in a local Frobenius ball around optimal. Then, we prove Theorem C.1 by constructing three supermartingales related to $f(\mathbf{U}_t, \mathbf{V}_t), g_i(\mathbf{U}_t, \mathbf{V}_t), h_j(\mathbf{U}_t, \mathbf{V}_t)$ each, and applying concentration argument.

To make the notation clear, denote gradient $\nabla f(\mathbf{U}, \mathbf{V}) \in \mathbb{R}^{(d_1+d_2) \times k}$:

$$\nabla f(\mathbf{U}, \mathbf{V}) = \begin{pmatrix} \frac{\partial}{\partial \mathbf{U}} f(\mathbf{U}, \mathbf{V}) \\ \frac{\partial}{\partial \mathbf{V}} f(\mathbf{U}, \mathbf{V}) \end{pmatrix}$$

Also denote the stochastic gradient $SG(\mathbf{U}, \mathbf{V})$ by (if we sampled entry (i, j) of matrix \mathbf{M})

$$\begin{aligned} SG(\mathbf{U}, \mathbf{V}) &= 2d_1d_2(\mathbf{UV}^\top - \mathbf{M})_{ij} \begin{pmatrix} \mathbf{e}_i \mathbf{e}_j^\top \mathbf{V} \\ \mathbf{e}_j \mathbf{e}_i^\top \mathbf{U} \end{pmatrix} \\ \mathbb{E}SG(\mathbf{U}, \mathbf{V}) &= \nabla f(\mathbf{U}, \mathbf{V}) = 2 \begin{pmatrix} (\mathbf{UV}^\top - \mathbf{M})\mathbf{V} \\ (\mathbf{UV}^\top - \mathbf{M})^\top \mathbf{U} \end{pmatrix} \end{aligned}$$

By update function, we know:

$$\begin{pmatrix} \mathbf{U}_{t+1} \\ \mathbf{V}_{t+1} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{U}}_t \\ \tilde{\mathbf{V}}_t \end{pmatrix} - \eta SG(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t)$$

and $\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top = \mathbf{U}_t \mathbf{V}_t^\top$ is the renormalized version of $\mathbf{U}_t \mathbf{V}_t^\top$.

C.1 Geometric Properties in Local Region

Similar to symmetric PSD case, we first prove two lemmas w.r.t the smoothness and property similar to strongly convex for objective function:

Lemma C.2. *Within the region $\mathcal{D} = \{(\mathbf{U}, \mathbf{V}) \mid \|\mathbf{U}\| \leq \Gamma, \|\mathbf{V}\| \leq \Gamma\}$, we have function $f(\mathbf{U}, \mathbf{V}) = \|\mathbf{M} - \mathbf{UV}^\top\|_F^2$ satisfying:*

$$\|\nabla f(\mathbf{U}_1, \mathbf{V}_1) - \nabla f(\mathbf{U}_2, \mathbf{V}_2)\|_F^2 \leq \beta^2 (\|\mathbf{U}_1 - \mathbf{U}_2\|_F^2 + \|\mathbf{V}_1 - \mathbf{V}_2\|_F^2)$$

where smoothness parameter $\beta = 8 \max\{\Gamma^2, \|\mathbf{M}\|\}$.

Proof. Inside region \mathcal{D} , we have:

$$\begin{aligned} & \|\nabla f(\mathbf{U}_1, \mathbf{V}_1) - \nabla f(\mathbf{U}_2, \mathbf{V}_2)\|_F^2 \\ &= \left\| \frac{\partial}{\partial \mathbf{U}} f(\mathbf{U}_1, \mathbf{V}_1) - \frac{\partial}{\partial \mathbf{U}} f(\mathbf{U}_2, \mathbf{V}_2) \right\|_F^2 + \left\| \frac{\partial}{\partial \mathbf{V}} f(\mathbf{U}_1, \mathbf{V}_1) - \frac{\partial}{\partial \mathbf{V}} f(\mathbf{U}_2, \mathbf{V}_2) \right\|_F^2 \\ &= 4 \left(\|(\mathbf{U}_1 \mathbf{V}_1^\top - \mathbf{M})\mathbf{V}_1 - (\mathbf{U}_2 \mathbf{V}_2^\top - \mathbf{M})\mathbf{V}_2\|_F^2 + \|(\mathbf{U}_1 \mathbf{V}_1^\top - \mathbf{M})^\top \mathbf{U}_1 - (\mathbf{U}_2 \mathbf{V}_2^\top - \mathbf{M})^\top \mathbf{U}_2\|_F^2 \right) \\ &\leq 64 \max\{\Gamma^4, \|\mathbf{M}\|^2\} (\|\mathbf{U}_1 - \mathbf{U}_2\|_F^2 + \|\mathbf{V}_1 - \mathbf{V}_2\|_F^2) \end{aligned}$$

The last step is by similar technics as in the proof of Lemma B.2, by expanding

$$\mathbf{U}_1 \mathbf{V}_1^\top \mathbf{V}_1 - \mathbf{U}_2 \mathbf{V}_2^\top \mathbf{V}_2 = (\mathbf{U}_1 - \mathbf{U}_2) \mathbf{V}_1^\top \mathbf{V}_1 + \mathbf{U}_2 (\mathbf{V}_1 - \mathbf{V}_2)^\top \mathbf{V}_1 + \mathbf{U}_2 \mathbf{V}_2^\top (\mathbf{V}_1 - \mathbf{V}_2)$$

□

Lemma C.3. *Within the region $\mathcal{D} = \{(\mathbf{U}, \mathbf{V}) \mid \sigma_{\min}(\mathbf{X}^\top \mathbf{U}) \geq \gamma, \sigma_{\min}(\mathbf{Y}^\top \mathbf{V}) \geq \gamma\}$, then we have function $f(\mathbf{U}, \mathbf{V}) = \|\mathbf{M} - \mathbf{UV}^\top\|_F^2$ satisfying:*

$$\|\nabla f(\mathbf{U}, \mathbf{V})\|_F^2 \geq \alpha f(\mathbf{U}, \mathbf{V})$$

where constant $\alpha = 4\gamma^2$.

Proof. Let $\hat{\mathbf{U}}, \hat{\mathbf{V}}$ be the left singular vectors of \mathbf{U}, \mathbf{V} . Inside region \mathcal{D} , we have:

$$\begin{aligned}
& \|(\mathbf{UV}^\top - \mathbf{M})\mathbf{V}\|_F^2 \\
&= \|\mathcal{P}_{\hat{\mathbf{U}}}(\mathbf{UV}^\top - \mathbf{M})\mathbf{V}\|_F^2 + \|\mathcal{P}_{\hat{\mathbf{U}}^\perp}(\mathbf{UV}^\top - \mathbf{M})\mathbf{V}\|_F^2 \\
&\geq \sigma_k(\mathbf{V})^2 \|\mathcal{P}_{\hat{\mathbf{U}}}(\mathbf{UV}^\top - \mathbf{M})\mathcal{P}_{\hat{\mathbf{V}}}\|_F^2 + \|\mathcal{P}_{\hat{\mathbf{U}}^\perp} \mathbf{M}\mathbf{V}\|_F^2 \\
&\geq \sigma_k(\mathbf{V})^2 \|\mathcal{P}_{\hat{\mathbf{U}}}(\mathbf{UV}^\top - \mathbf{M})\mathcal{P}_{\hat{\mathbf{V}}}\|_F^2 + \sigma_{\min}(\mathbf{Y}^\top \mathbf{V})^2 \|\mathcal{P}_{\hat{\mathbf{U}}^\perp} \mathbf{X}\mathbf{\Sigma}\|_F^2 \\
&= \sigma_k(\mathbf{V})^2 \|\mathcal{P}_{\hat{\mathbf{U}}}(\mathbf{UV}^\top - \mathbf{M})\mathcal{P}_{\hat{\mathbf{V}}}\|_F^2 + \sigma_{\min}(\mathbf{Y}^\top \mathbf{V})^2 \|\mathcal{P}_{\hat{\mathbf{U}}^\perp} \mathbf{M}\|_F^2
\end{aligned}$$

Therefore, by symmetry, we have:

$$\begin{aligned}
\|\nabla f(\mathbf{U}, \mathbf{V})\|_F^2 &= 4(\|(\mathbf{UV}^\top - \mathbf{M})\mathbf{V}\|_F^2 + \|(\mathbf{UV}^\top - \mathbf{M})^\top \mathbf{U}\|_F^2) \\
&\geq 4\gamma^2(2\|\mathcal{P}_{\hat{\mathbf{U}}}(\mathbf{UV}^\top - \mathbf{M})\mathcal{P}_{\hat{\mathbf{V}}}\|_F^2 + \|\mathcal{P}_{\hat{\mathbf{U}}^\perp} \mathbf{M}\|_F^2 + \|\mathbf{M}\mathcal{P}_{\hat{\mathbf{V}}^\perp}\|_F^2) \\
&\geq 4\gamma^2\|\mathbf{UV}^\top - \mathbf{M}\|_F^2
\end{aligned}$$

□

Next, we show as long as we are in some Frobenious ball around optimum, then we have good spectral property over \mathbf{U}, \mathbf{V} which guarantees the preconditions for Lemma C.2 and Lemma C.3.

Lemma C.4. *Within the region $\mathcal{D} = \{(\mathbf{U}, \mathbf{V}) \mid \|\mathbf{M} - \mathbf{UV}^\top\|_F \leq \frac{1}{10}\sigma_k(\mathbf{M})\}$, and for $\mathbf{U} = \mathbf{W}_U \mathbf{D}^{\frac{1}{2}}, \mathbf{V} = \mathbf{W}_V \mathbf{D}^{\frac{1}{2}}$ where $\mathbf{W}_U \mathbf{D} \mathbf{W}_V = \text{SVD}(\mathbf{UV}^\top)$, we have:*

$$\begin{aligned}
\|\mathbf{U}\| &\leq \sqrt{2\|\mathbf{M}\|}, & \sigma_{\min}(\mathbf{X}^\top \mathbf{U}) &\geq \sqrt{\sigma_k(\mathbf{M})/2} \\
\|\mathbf{V}\| &\leq \sqrt{2\|\mathbf{M}\|}, & \sigma_{\min}(\mathbf{Y}^\top \mathbf{V}) &\geq \sqrt{\sigma_k(\mathbf{M})/2}
\end{aligned}$$

Proof. For spectral norm of \mathbf{U} , we have:

$$\|\mathbf{U}\|^2 = \|\mathbf{D}\| = \|\mathbf{UV}^\top\| \leq \|\mathbf{M}\| + \|\mathbf{M} - \mathbf{UV}^\top\| \leq \|\mathbf{M}\| + \|\mathbf{M} - \mathbf{UV}^\top\|_F \leq 2\|\mathbf{M}\|$$

For the minimum singular value of $\mathbf{U}^\top \mathbf{U}$, we have:

$$\begin{aligned}
\sigma_{\min}(\mathbf{U}^\top \mathbf{U}) &= \sigma_k(\mathbf{D}) = \sigma_k(\mathbf{UV}^\top) \geq \sigma_k(\mathbf{M}) - \|\mathbf{M} - \mathbf{UV}^\top\| \\
&\geq \sigma_k(\mathbf{M}) - \|\mathbf{M} - \mathbf{UV}^\top\|_F \geq \frac{9}{10}\sigma_k(\mathbf{M})
\end{aligned}$$

By symmetry, the same holds for \mathbf{V} . On the other hand, we have:

$$\begin{aligned}
\frac{1}{10}\sigma_k(\mathbf{M}) &\geq \|\mathbf{M} - \mathbf{UV}^\top\|_F \geq \|\mathcal{P}_{\mathbf{X}^\perp}(\mathbf{M} - \mathbf{UV}^\top)\|_F = \|\mathcal{P}_{\mathbf{X}^\perp} \mathbf{UV}^\top\|_F = \|\mathcal{P}_{\mathbf{X}^\perp} \mathbf{W}_U \mathbf{D}\|_F \\
&\geq \|\mathcal{P}_{\mathbf{X}^\perp} \mathbf{W}_U \mathbf{D}\| \geq \frac{9}{10}\sigma_k(\mathbf{M}) \|\mathbf{X}^\perp \mathbf{W}_U\|
\end{aligned}$$

Let the principal angle between \mathbf{X} and \mathbf{W}_U to be θ . This gives $\sin^2 \theta = \|\mathbf{X}^\perp \mathbf{W}_U\|^2 \leq \frac{1}{9}$. Thus $\cos^2 \theta = \sigma_{\min}^2(\mathbf{X}^\top \mathbf{W}_U) \geq \frac{8}{9}$. Therefore:

$$\sigma_{\min}^2(\mathbf{X}^\top \mathbf{U}) \geq \sigma_{\min}^2(\mathbf{X}^\top \mathbf{W}_U) \sigma_{\min}(\mathbf{U}^\top \mathbf{U}) \geq \sigma_k(\mathbf{M})/2$$

□

C.2 Proof of Theorem C.1

Now, we are ready for our key theorem. By Lemma C.2, Lemma C.3, and Lemma C.4, we already know the function has good property locally in the region $\mathcal{D} = \{(\mathbf{U}, \mathbf{V}) \mid \|\mathbf{M} - \mathbf{U}\mathbf{V}^\top\|_F \leq \frac{1}{10}\sigma_k(\mathbf{M})\}$ which alludes linear convergence. Similar to the symmetric PSD case, the work remains is to prove that once we initialize inside this region, our algorithm will guarantee \mathbf{U}, \mathbf{V} never leave this region with high probability even with relatively large stepsize. Again, we also need to control the incoherence of $\mathbf{U}_t, \mathbf{V}_t$ over all iterates additionally to achieve tight sample complexity and near optimal runtime.

Following is our formal proof.

Proof of Theorem C.1. For simplicity of notation, we assume $d = d_1 = d_2$, and do not distinguish d_1 and d_2 . However, it is easy to check our proof never use the property \mathbf{M} is square matrix. The proof easily extends to $d_1 \neq d_2$ case by replacing d in the proof with suitable d_1, d_2 .

Define event $\mathfrak{E}_t = \{\forall \tau \leq t, f(\mathbf{U}_\tau, \mathbf{V}_\tau) \leq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2, \max_i g_i(\mathbf{U}_\tau, \mathbf{V}_\tau) \leq \frac{100\mu\kappa\kappa^2}{d}, \max_j h_j(\mathbf{U}_\tau, \mathbf{V}_\tau) \leq \frac{100\mu\kappa\kappa^2}{d}\}$. Theorem C.1 is equivalent to prove event \mathfrak{E}_T happens with high probability. The proof achieves this by constructing two supermartingales for $f(\mathbf{U}_t, \mathbf{V}_t)1_{\mathfrak{E}_t}, g_i(\mathbf{U}_t, \mathbf{V}_t)1_{\mathfrak{E}_t}$ and $h_i(\mathbf{U}_t, \mathbf{V}_t)1_{\mathfrak{E}_t}$ (where $1_{(\cdot)}$ denote indicator function), applies concentration argument.

The proofs also follow similar structure as symmetric PSD case:

1. The constructions of supermartingales
2. Their probability 1 bound and variance bound in order to apply Azuma-Bernstein inequality
3. Final combination of concentration results to conclude the proof

Then let filtration $\mathfrak{F}_t = \sigma\{SG(\mathbf{U}_0, \mathbf{V}_0), \dots, SG(\mathbf{U}_{t-1}, \mathbf{V}_{t-1})\}$ where $\sigma\{\cdot\}$ denotes the sigma field. Also let event, note $\mathfrak{E}_t \subset \mathfrak{F}_t$. Also $\mathfrak{E}_{t+1} \subset \mathfrak{E}_t$, and thus $1_{\mathfrak{E}_{t+1}} \leq 1_{\mathfrak{E}_t}$.

By Lemma C.4, we immediately know that conditioned on \mathfrak{E}_t , we have $\|\mathbf{U}_t\| \leq \sqrt{2}, \|\mathbf{V}_t\| \leq \sqrt{2}, \sigma_{\min}(\mathbf{X}^\top \mathbf{U}_t) \geq 1/\sqrt{2\kappa}, \sigma_{\min}(\mathbf{Y}^\top \mathbf{V}_t) \geq 1/\sqrt{2\kappa}$. We will use this fact throughout the proof.

For simplicity, when it's clear from the context, we denote:

$$\begin{pmatrix} \Delta_{\mathbf{U}} \\ \Delta_{\mathbf{V}} \end{pmatrix} = -\eta SG(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) = \begin{pmatrix} \mathbf{U}_{t+1} \\ \mathbf{V}_{t+1} \end{pmatrix} - \begin{pmatrix} \tilde{\mathbf{U}}_t \\ \tilde{\mathbf{V}}_t \end{pmatrix}$$

Construction of supermartingale G : First, since potential function $g_l(\mathbf{U}, \mathbf{V})$ is forth-order polynomial, we can expand:

$$\begin{aligned} g_l(\tilde{\mathbf{U}}_{t+1}, \tilde{\mathbf{V}}_{t+1}) &= g_l(\mathbf{U}_{t+1}, \mathbf{V}_{t+1}) = g_l(\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}}, \tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) \\ &= \mathbf{e}_l^\top (\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}}) (\tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}})^\top (\tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) (\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}})^\top \mathbf{e}_l \\ &= g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + 2\mathbf{e}_l^\top \Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t \mathbf{e}_l + R_2 \\ &= g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + R_1 \end{aligned}$$

Where we denote R_1 as the sum of first order terms and higher order terms (all second/third/forth order terms), and R_2 as the sum of second order terms and higher order terms.

We now give a proposition about properties of R_1 and R_2 which involves a lot calculation, and postpone its proof in the end of this section.

Proposition C.5. *With above notations, we have following inequalities hold true.*

$$\begin{aligned}\mathbb{E}[R_2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^2 O(\mu^2 k^2 \kappa^4) \mathbf{1}_{\mathfrak{E}_t} \\ |R_1| \mathbf{1}_{\mathfrak{E}_t} &\leq \eta O(\mu^2 k^2 \kappa^5) \mathbf{1}_{\mathfrak{E}_t} \quad w.p \ 1 \\ \mathbb{E}[R_1^2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^2 O\left(\frac{\mu^3 k^3 \kappa^6}{d}\right) \mathbf{1}_{\mathfrak{E}_t}\end{aligned}$$

Then by taking conditional expectation, we have:

$$\begin{aligned}\mathbb{E}[g_l(\tilde{\mathbf{U}}_{t+1}, \tilde{\mathbf{V}}_{t+1}) \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &= \mathbb{E}[g_l(\mathbf{U}_{t+1}, \mathbf{V}_{t+1}) \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \\ &\leq \mathbb{E}[g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + 2\mathbf{e}_l^\top \Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t \mathbf{e}_l + R_2] \mathbf{1}_{\mathfrak{E}_t}\end{aligned}$$

The first order term can be calculated as:

$$\begin{aligned}&[-\mathbb{E}2\mathbf{e}_l^\top \Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t \mathbf{e}_l] \mathbf{1}_{\mathfrak{E}_t} \\ &= [-4\mathbf{e}_l^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l - 4\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}^\top) \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l] \mathbf{1}_{\mathfrak{E}_t} \\ &= [-4\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 4\mathbf{e}_l^\top \mathbf{M} \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l - 4\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}^\top) \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l] \mathbf{1}_{\mathfrak{E}_t} \\ &\leq [-4[\sigma_{\min}(\tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t) \|\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top\|^2 + \|\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top\| \|\tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top\| \|\mathbf{e}_l^\top \mathbf{M}\| \\ &\quad + \|\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top\| \|\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}^\top\|_F \|\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top\|]] \mathbf{1}_{\mathfrak{E}_t} \\ &\leq [-\frac{2}{\kappa} g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + 80 \frac{\mu k \kappa}{d} + \frac{4}{10\kappa} g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t)] \mathbf{1}_{\mathfrak{E}_t} \\ &\leq [-\frac{1}{\kappa} g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + 80 \frac{\mu k \kappa}{d}] \mathbf{1}_{\mathfrak{E}_t}\end{aligned}$$

In second last inequality, we use key observation:

$$\|\mathbf{e}_k^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top\| = \|\mathbf{e}_k^\top \mathbf{W}_{\mathbf{U}} \mathbf{D} \mathbf{W}_{\mathbf{V}}^\top\| = \|\mathbf{e}_k^\top \mathbf{W}_{\mathbf{U}} \mathbf{D} \mathbf{W}_{\mathbf{U}}^\top\| = \|\mathbf{e}_k^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top\|$$

By Proposition C.5, we know $\mathbb{E}[R_2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(\mu^2 k^2 \kappa^4) \mathbf{1}_{\mathfrak{E}_t}$. Combine both facts and recall $\eta < \frac{c}{\mu d k \kappa^3 \log d}$, we have:

$$\begin{aligned}\mathbb{E}[g_i(\tilde{\mathbf{U}}_{t+1}, \tilde{\mathbf{V}}_{t+1}) \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq [(1 - \frac{\eta}{\kappa}) g_i(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + \frac{80\eta\mu k \kappa}{d} + O(\eta^2 \mu^2 k^2 \kappa^4)] \mathbf{1}_{\mathfrak{E}_t} \\ &\leq [(1 - \frac{\eta}{\kappa}) g_i(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + \frac{90\eta\mu k \kappa}{d}] \mathbf{1}_{\mathfrak{E}_t}\end{aligned}$$

The last inequality is achieved by choosing c small enough.

Let $G_{it} = (1 - \frac{\eta}{\kappa})^{-t} (g_i(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_{t-1}} - 90 \frac{\mu k \kappa^2}{d})$. This gives:

$$\mathbb{E}[G_{i(t+1)} | \mathfrak{F}_t] \leq (1 - \frac{\eta}{\kappa})^{-t} (g_i(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_t} - 90 \frac{\mu k \kappa^2}{d}) \leq G_{it}$$

The right inequality is true since $\mathbf{1}_{\mathfrak{E}_t} \leq \mathbf{1}_{\mathfrak{E}_{t-1}}$. This implies G_{it} is supermartingale.

Probability 1 bound for G : We also know:

$$G_{i(t+1)} - \mathbb{E}[G_{i(t+1)}|\mathfrak{F}_t] = (1 - \frac{\eta}{\kappa})^{-(t+1)}[R_1 - \mathbb{E}R_1]1_{\mathfrak{E}_t}$$

By Proposition C.5, we know with probability 1 that $|R_1|1_{\mathfrak{E}_t} \leq \eta O(\mu^2 k^2 \kappa^5)1_{\mathfrak{E}_t}$. This gives with probability 1:

$$|G_{it} - \mathbb{E}[G_{it}|\mathfrak{F}_{t-1}]| \leq (1 - \frac{\eta}{\kappa})^{-t} \eta O(\mu^2 k^2 \kappa^5)1_{\mathfrak{E}_{t-1}} \quad (17)$$

Variance bound for G : We also know

$$\text{Var}(G_{i(t+1)}|\mathfrak{F}_t) = (1 - \frac{\eta}{\kappa})^{-2(t+1)} [\mathbb{E}R_1^2 1_{\mathfrak{E}_t} - (\mathbb{E}R_1 1_{\mathfrak{E}_t})^2] \leq \mathbb{E}[R_1^2 1_{\mathfrak{E}_t}|\mathfrak{F}_t]$$

By Proposition C.5, we know that $\mathbb{E}[R_1^2 1_{\mathfrak{E}_t}|\mathfrak{F}_t] \leq \eta^2 O(\frac{\mu^3 k^3 \kappa^6}{d})1_{\mathfrak{E}_t}$. This gives:

$$\text{Var}(G_{it}|\mathfrak{F}_{t-1}) \leq (1 - \frac{\eta}{\kappa})^{-2t} \eta^2 O(\frac{\mu^3 k^3 \kappa^6}{d})1_{\mathfrak{E}_{t-1}} \quad (18)$$

Bernstein's inequality for G : Let $\sigma^2 = \sum_{\tau=1}^t \text{Var}(G_{i\tau}|\mathfrak{F}_{\tau-1})$, and R satisfies, with probability 1 that $|G_{i\tau} - \mathbb{E}[G_{i\tau}|\mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Then By standard Bernstein concentration inequality, we know:

$$P(G_{it} \geq G_{i0} + s) \leq \exp(\frac{s^2/2}{\sigma^2 + Rs/3})$$

Since $G_{i0} = g_i(\tilde{\mathbf{U}}_0, \tilde{\mathbf{V}}_0) - 90\frac{\mu k \kappa^2}{d}$, let $s' = O(1)(1 - \frac{\eta}{\kappa})^t [\sqrt{\sigma^2 \log d} + R \log d]$, we know

$$P\left(g_i(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t)1_{\mathfrak{E}_{t-1}} \geq 90\frac{\mu k \kappa^2}{d} + (1 - \frac{\eta}{\kappa})^t (g_i(\tilde{\mathbf{U}}_0, \tilde{\mathbf{V}}_0) - 90\frac{\mu k \kappa^2}{d}) + s'\right) \leq \frac{1}{3d^{11}}$$

By Eq.(17), we know $R = (1 - \frac{\eta}{\kappa})^{-t} \eta O(\mu^2 k^2 \kappa^5)$ satisfies that $|G_{i\tau} - \mathbb{E}[G_{i\tau}|\mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Also by Eq. (18), we have:

$$(1 - \frac{\eta}{\kappa})^t \sqrt{\sigma^2 \log d} \leq \eta O\left(\sqrt{\frac{\mu^3 k^3 \kappa^6 \log d}{d}}\right) \sqrt{\sum_{\tau=1}^t (1 - \frac{\eta}{\kappa})^{2t-2\tau}} \leq \sqrt{\eta} O\left(\sqrt{\frac{\mu^3 k^3 \kappa^7 \log d}{d}}\right)$$

by $\eta < \frac{c}{\mu d k \kappa^3 \log d}$ and choosing c to be small enough, we have:

$$s' = \sqrt{\eta} O\left(\sqrt{\frac{\mu^3 k^3 \kappa^7 \log d}{d}}\right) + \eta O(\mu^2 k^2 \kappa^5 \log d) \leq 10\frac{\mu k \kappa^2}{d}$$

Since initialization gives $\max_i g_i(\mathbf{U}_0, \mathbf{V}_0) \leq \frac{10\mu k \kappa^2}{d}$, therefore:

$$P(g_i(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t)1_{\mathfrak{E}_{t-1}} \geq 100\frac{\mu k \kappa^2}{d}) \leq \frac{1}{3d^{11}}$$

That is equivalent to:

$$P(\mathfrak{E}_{t-1} \cap \{g_i(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \geq 100\frac{\mu k \kappa^2}{d}\}) \leq \frac{1}{3d^{11}} \quad (19)$$

By symmetry, we can also have corresponding result for $h_j(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t)$.

Construction of supermartingale \mathbf{F} : Similarly, we also need to construct a martingale for $f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t)$. Again, we can write f as forth order polynomial:

$$\begin{aligned} f(\tilde{\mathbf{U}}_{t+1}, \tilde{\mathbf{V}}_{t+1}) &= f(\mathbf{U}_{t+1}, \mathbf{V}_{t+1}) = f(\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}}, \tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) \\ &= \text{tr} \left([(\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}})(\tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) - \mathbf{M}][(\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}})(\tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) - \mathbf{M}]^\top \right) \\ &= f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + 2\text{tr}(\Delta_{\mathbf{U}}\tilde{\mathbf{V}}_t^\top(\tilde{\mathbf{U}}_t\tilde{\mathbf{V}}_t^\top - \mathbf{M})^\top) + 2\text{tr}(\Delta_{\mathbf{V}}\tilde{\mathbf{U}}_t^\top(\tilde{\mathbf{U}}_t\tilde{\mathbf{V}}_t^\top - \mathbf{M})) + Q_2 \\ &= f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + Q_1 \end{aligned}$$

Where we denote Q_1 as the sum of first order terms and higher order terms (all second/third/forth order terms), and Q_2 as the sum of second order terms and higher order terms.

We also now give a proposition about properties of Q_1 and Q_2 which involves a lot calculation, and postpone its proof in the end of this section.

Proposition C.6. *With above notations, we have following inequalities hold true.*

$$\begin{aligned} \mathbb{E}[Q_2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^2 O(\mu dk \kappa^2) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_t} \\ |Q_1| \mathbf{1}_{\mathfrak{E}_t} &\leq \eta O(\mu dk \kappa^3) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_t} \quad w.p \ 1 \\ \mathbb{E}[Q_1^2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^2 O(\mu dk \kappa^2) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_t} \end{aligned}$$

Then by Proposition C.6, we know $\mathbb{E}[Q_2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(\mu dk \kappa^2) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_t}$. By taking conditional expectation, we have:

$$\begin{aligned} &\mathbb{E}[f(\mathbf{U}_{t+1}) \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \\ &\leq [f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) - \mathbb{E}\langle \nabla f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t), \eta SG(\mathbf{U}_t) \rangle + \mathbb{E}Q_2] \mathbf{1}_{\mathfrak{E}_t} \\ &= [f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) - \eta \left\| \nabla f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \right\|_F^2 + \mathbb{E}Q_2] \mathbf{1}_{\mathfrak{E}_t} \\ &\leq \left[\left(1 - \frac{2\eta}{\kappa}\right) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + \eta^2 O(\mu dk \kappa^2) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \right] \mathbf{1}_{\mathfrak{E}_t} \\ &\leq \left(1 - \frac{\eta}{\kappa}\right) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_t} \end{aligned}$$

Let $F_t = \left(1 - \frac{\eta}{\kappa}\right)^{-t} f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_{t-1}}$, we know F_t is also a supermartingale.

Probability 1 bound: We also know

$$F_{t+1} - \mathbb{E}[F_{t+1} | \mathfrak{F}_t] = \left(1 - \frac{\eta}{\kappa}\right)^{-(t+1)} [Q_1 - \mathbb{E}Q_1] \mathbf{1}_{\mathfrak{E}_t}$$

By Proposition C.6, we know with probability 1 that $|Q_1| \mathbf{1}_{\mathfrak{E}_t} \leq \eta O(\mu dk \kappa^3) f(\mathbf{U}_t, \mathbf{V}_t) \mathbf{1}_{\mathfrak{E}_t}$. This gives with probability 1:

$$|F_t - \mathbb{E}F_t| \leq \left(1 - \frac{\eta}{\kappa}\right)^{-t} \eta O(\mu dk \kappa^3) f(\mathbf{U}_{t-1}) \mathbf{1}_{\mathfrak{E}_{t-1}} \leq \left(1 - \frac{\eta}{\kappa}\right)^{-t} \left(1 - \frac{\eta}{2\kappa}\right)^t \eta O(\mu dk \kappa) \mathbf{1}_{\mathfrak{E}_{t-1}} \quad (20)$$

Variance bound: We also know

$$\text{Var}(F_{t+1} | \mathfrak{F}_t) = \left(1 - \frac{\eta}{\kappa}\right)^{-2(t+1)} [\mathbb{E}Q_1^2 \mathbf{1}_{\mathfrak{E}_t} - (\mathbb{E}Q_1 \mathbf{1}_{\mathfrak{E}_t})^2] \leq \left(1 - \frac{\eta}{\kappa}\right)^{-2(t+1)} \mathbb{E}[Q_1^2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t]$$

By Proposition C.6, we know that $\mathbb{E}[Q_1^2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(\mu dk \kappa^2) f^2(\mathbf{U}_t, \mathbf{V}_t) \mathbf{1}_{\mathfrak{E}_t}$. This gives:

$$\text{Var}(F_t | \mathfrak{F}_{t-1}) \leq \left(1 - \frac{\eta}{\kappa}\right)^{-2t} \eta^2 O(\mu dk \kappa^2) f^2(\mathbf{U}_{t-1}) \mathbf{1}_{\mathfrak{E}_{t-1}} \leq \left(1 - \frac{\eta}{\kappa}\right)^{-2t} \left(1 - \frac{\eta}{2\kappa}\right)^{2t} \eta^2 O\left(\frac{\mu dk}{\kappa^2}\right) \mathbf{1}_{\mathfrak{E}_{t-1}} \quad (21)$$

Bernstein's inequality: Let $\sigma^2 = \sum_{\tau=1}^t \text{Var}(F_\tau | \mathfrak{F}_{\tau-1})$, and R satisfies, with probability 1 that $|F_\tau - \mathbb{E}[F_\tau | \mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Then By standard Bernstein concentration inequality, we know:

$$P(F_t \geq F_0 + s) \leq \exp\left(-\frac{s^2/2}{\sigma^2 + Rs/3}\right)$$

Let $s' = O(1)(1 - \frac{\eta}{\kappa})^t [\sqrt{\sigma^2 \log d} + R \log d]$, this gives:

$$P(f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_{t-1}} \geq (1 - \frac{\eta}{\kappa})^t f(\mathbf{U}_0) + s') \leq \frac{1}{3d^{10}}$$

By Eq.(20), we know $R = (1 - \frac{\eta}{\kappa})^{-t} (1 - \frac{\eta}{2\kappa})^t \eta O(\mu d k \kappa)$ satisfies that $|F_\tau - \mathbb{E}[F_\tau | \mathfrak{F}_{\tau-1}]| \leq R$, $\tau = 1, \dots, t$. Also by Eq. (21), we have:

$$\begin{aligned} (1 - \frac{\eta}{\kappa})^t \sqrt{\sigma^2 \log d} &\leq \eta O\left(\sqrt{\frac{\mu d k \log d}{\kappa^2}}\right) \sqrt{\sum_{\tau=1}^t (1 - \frac{\eta}{\kappa})^{2t-2\tau} (1 - \frac{\eta}{2\kappa})^{2\tau}} \\ &\leq (1 - \frac{\eta}{2\kappa})^t \eta O\left(\sqrt{\frac{\mu d k \log d}{\kappa^2}}\right) \sqrt{\sum_{\tau=1}^t (1 - \frac{\eta}{\kappa})^{2t-2\tau} (1 - \frac{\eta}{2\kappa})^{2\tau-2t}} \leq (1 - \frac{\eta}{2\kappa})^t \sqrt{\eta} O\left(\sqrt{\frac{\mu d k \log d}{\kappa}}\right) \end{aligned}$$

by $\eta < \frac{c}{\mu d k \kappa^3 \log d}$ and choosing c to be small enough, we have:

$$s' = (1 - \frac{\eta}{2\kappa})^t [\sqrt{\eta} O\left(\sqrt{\frac{\mu d k \kappa \log d}{\kappa}}\right) + \eta O(\mu d k \kappa)] \leq (1 - \frac{\eta}{2\kappa})^t \left(\frac{1}{20\kappa}\right)^2$$

Since $F_0 = f(\mathbf{U}_0) \leq (\frac{1}{20\kappa})^2$, therefore:

$$P(f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \mathbf{1}_{\mathfrak{E}_{t-1}} \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2) \leq \frac{1}{3d^{10}}$$

That is equivalent to:

$$P(\mathfrak{E}_{t-1} \cap \{f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2\}) \leq \frac{1}{3d^{10}} \quad (22)$$

Probability for event \mathfrak{E}_T : Finally, combining the concentration result for martingale G (Eq.(19)) and martingale F (Eq.(22)), we conclude:

$$\begin{aligned} &P(\mathfrak{E}_{t-1} \cap \bar{\mathfrak{E}}_t) \\ &= P\left[\mathfrak{E}_{t-1} \cap \left([\cup_i \{g_i(\mathbf{U}_t, \mathbf{V}_t) \geq 100 \frac{\mu k \kappa^2}{d}\}] \cup [\cup_j \{h_j(\mathbf{U}_t, \mathbf{V}_t) \geq 100 \frac{\mu k \kappa^2}{d}\}] \cup \{f(\mathbf{U}_t) \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2\}\right)\right] \\ &\leq 2 \sum_{i=1}^d P(\mathfrak{E}_{t-1} \cap \{g_i(\mathbf{U}_t, \mathbf{V}_t) \geq 100 \frac{\mu k \kappa^2}{d}\}) + P(\mathfrak{E}_{t-1} \cap \{f(\mathbf{U}_t, \mathbf{V}_t) \geq (1 - \frac{\eta}{2\kappa})^t (\frac{1}{10\kappa})^2\}) \\ &\leq \frac{1}{d^{10}} \end{aligned}$$

Since

$$P(\bar{\mathfrak{E}}_T) = \sum_{t=1}^T P(\mathfrak{E}_{t-1} \cap \bar{\mathfrak{E}}_t) \leq \frac{T}{d^{10}}$$

We finishes the proof. \square

Finally we give proof for Proposition C.5 and Proposition C.6. The proof mostly consists of expanding every term and careful calculations.

Proof of Proposition C.5. For simplicity of notation, we hide the term $1_{\mathfrak{E}_t}$ in all following equations. Reader should always think every term in this proof multiplied by $1_{\mathfrak{E}_t}$. Recall that:

$$\begin{aligned} SG(\mathbf{U}, \mathbf{V}) &= 2d^2(\mathbf{U}\mathbf{V}^\top - \mathbf{M})_{ij} \begin{pmatrix} \mathbf{e}_i \mathbf{e}_j^\top \mathbf{V} \\ \mathbf{e}_j \mathbf{e}_i^\top \mathbf{U} \end{pmatrix} \\ \begin{pmatrix} \Delta \mathbf{U} \\ \Delta \mathbf{V} \end{pmatrix} &= -\eta SG(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) = \begin{pmatrix} \mathbf{U}_{t+1} \\ \mathbf{V}_{t+1} \end{pmatrix} - \begin{pmatrix} \tilde{\mathbf{U}}_t \\ \tilde{\mathbf{V}}_t \end{pmatrix} \end{aligned}$$

We first prove first three inequality. Recall that:

$$\begin{aligned} g_l(\tilde{\mathbf{U}}_{t+1}, \tilde{\mathbf{V}}_{t+1}) &= g_l(\mathbf{U}_{t+1}, \mathbf{V}_{t+1}) = g_l(\tilde{\mathbf{U}}_t + \Delta \mathbf{U}, \tilde{\mathbf{V}}_t + \Delta \mathbf{V}) \\ &= \mathbf{e}_l^\top (\tilde{\mathbf{U}}_t + \Delta \mathbf{U}) (\tilde{\mathbf{V}}_t + \Delta \mathbf{V})^\top (\tilde{\mathbf{V}}_t + \Delta \mathbf{V}) (\tilde{\mathbf{U}}_t + \Delta \mathbf{U})^\top \mathbf{e}_l \\ &= g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + 2\mathbf{e}_l^\top \Delta \mathbf{U} \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \Delta \mathbf{V} \tilde{\mathbf{U}}_t \mathbf{e}_l + R_2 \\ &= g_l(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + R_1 \end{aligned}$$

By expanding the polynomial, we can write out the first order term:

$$\begin{aligned} R_1 - R_2 &= 2\mathbf{e}_l^\top \Delta \mathbf{U} \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \Delta \mathbf{V} \tilde{\mathbf{U}}_t \mathbf{e}_l \\ &= -4\eta d^2 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij} \left(\delta_{il} \mathbf{e}_j^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{lj} (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{il} \right) \end{aligned}$$

Second order term:

$$\begin{aligned} R_2 - R_3 &= \mathbf{e}_l^\top \Delta \mathbf{U} \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \Delta \mathbf{U}^\top \mathbf{e}_l + \mathbf{e}_l^\top \tilde{\mathbf{U}}_t \Delta \mathbf{V}^\top \Delta \mathbf{V} \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \Delta \mathbf{U} \tilde{\mathbf{V}}_t^\top \Delta \mathbf{V} \tilde{\mathbf{U}}_t^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \Delta \mathbf{U} \Delta \mathbf{V}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \mathbf{e}_l \\ &= 4\eta^2 d^4 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^2 \\ &\quad \cdot \left(\delta_{il} \left\| \mathbf{e}_j^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top \right\|^2 + (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{li}^2 + 2\delta_{il} (\tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top)_{jj} (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{ii} + 2\delta_{il} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij}^2 \right) \end{aligned}$$

Third order term:

$$\begin{aligned} R_4 - R_3 &= 2\mathbf{e}_l^\top \Delta \mathbf{U} \tilde{\mathbf{V}}_t^\top \Delta \mathbf{V} \Delta \mathbf{U}^\top \mathbf{e}_l + 2\mathbf{e}_l^\top \Delta \mathbf{U} \Delta \mathbf{V}^\top \Delta \mathbf{V} \tilde{\mathbf{U}}_t^\top \mathbf{e}_l \\ &= -16\eta^3 d^6 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^3 \delta_{il} \left((\tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top)_{jj} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij} + (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{ii} \right) \end{aligned}$$

Fourth order term:

$$R_4 = \mathbf{e}_l^\top \Delta \mathbf{U} \Delta \mathbf{V}^\top \Delta \mathbf{V} \Delta \mathbf{U}^\top \mathbf{e}_l = 16\eta^4 d^8 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^4 \delta_{il} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij}^2$$

For the ease of proof, we denote $\chi = \frac{\mu k \kappa^2}{d}$, then we know conditioned on event \mathfrak{E}_t , we have: $\max_i \left\| \mathbf{e}_i^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \right\|^2 \leq O(\chi)$, and $\max_j \left\| \mathbf{e}_j^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \right\|^2 \leq O(\chi)$. Some key inequality we need to use in the proof are listed here:

$$\left\| \mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \right\| = \left\| \mathbf{e}_l^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top \right\| \quad \text{and} \quad \left\| \mathbf{e}_l^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{U}}_t^\top \right\| = \left\| \mathbf{e}_l^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top \right\| \quad (23)$$

and

$$|(\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t)_{ij}| \leq \left\| \mathbf{e}_i^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \right\| \leq O(\sqrt{\chi}) \quad (24)$$

The same also holds true for:

$$|(\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{ii}| \leq O(\sqrt{\chi}) \quad \text{and} \quad |(\tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top)_{jj}| \leq O(\sqrt{\chi}) \quad (25)$$

Another fact we frequently used is:

$$\frac{1}{2} \left\| \mathbf{e}_i^\top \mathbf{U} \mathbf{V}^\top \right\|^2 \leq \left\| \mathbf{e}_i^\top \mathbf{U} \right\|^2 \leq 2\kappa \left\| \mathbf{e}_i^\top \mathbf{U} \mathbf{V}^\top \right\|^2$$

This gives:

$$\left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty \leq \max_k \left\| \mathbf{e}_k \tilde{\mathbf{U}}_t \right\| \max_k \left\| \mathbf{e}_k \tilde{\mathbf{V}}_t \right\| + \max_k \|\mathbf{e}_k \mathbf{X}\| \max_k \|\mathbf{e}_k \mathbf{Y}\| \|\mathbf{S}\| \leq O(\chi \kappa)$$

and recall we choose $\eta < \frac{c}{\mu d k \kappa^3 \log d}$, where c is some universal constant, then we have:

$$\eta d^2 \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty = O(\eta d^2 \chi \kappa) \leq O(1) \quad (26)$$

With equation (23), (24), (25), (26), now we are ready to prove Lemma.

For the first inequality $\mathbb{E}[R_2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(\mu^2 k^2 \kappa^4) \mathbf{1}_{\mathfrak{E}_t}$:

$$\mathbb{E}[R_2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \mathbb{E}[|R_2 - R_3| \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] + \mathbb{E}[|R_3 - R_4| \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] + \mathbb{E}[|R_4| \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t]$$

For each term, we can bound as:

$$\begin{aligned} \mathbb{E}[|R_2 - R_3| \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^2 O(d^2) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^2 \left(\delta_{il} O(\chi) + (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{li}^2 \right) \\ &\leq \eta^2 O(d^2) \max_l \left\| \mathbf{e}_l^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \right\|^2 \sum_i \left(\delta_{il} O(\chi) + (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{li}^2 \right) \leq \eta^2 O(d^2 \chi^2) \\ \mathbb{E}[|R_3 - R_4| \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^3 O(d^4) \sum_{ij} |\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}|_{ij}^3 \delta_{il} O(\chi) \\ &\leq \eta^3 O(d^4) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty \left\| \mathbf{e}_l^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \right\|^2 O(\chi) \leq \eta^2 O(d^2 \chi^2) \\ \mathbb{E}[|R_4| \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^4 O(d^6) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^4 \delta_{il} O(\chi) \\ &\leq \eta^4 O(d^6) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^2 \left\| \mathbf{e}_l^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \right\|^2 O(\chi) \leq \eta^2 O(d^2 \chi^2) \end{aligned}$$

This gives in sum that

$$\mathbb{E}[R_2 \mathbf{1}_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(d^2 \chi^2) \mathbf{1}_{\mathfrak{E}_t} = \eta^2 O(\mu d k \kappa^2) f^2(\mathbf{U}_t) \mathbf{1}_{\mathfrak{E}_t}$$

For the second inequality $|R_1|1_{\mathfrak{E}_t} \leq \eta O(\mu^2 k^2 \kappa^5)1_{\mathfrak{E}_t}$ **w.p 1:**

$$|R_1|1_{\mathfrak{E}_t} \leq |R_1 - R_2|1_{\mathfrak{E}_t} + |R_2 - R_3|1_{\mathfrak{E}_t} + |R_3 - R_4|1_{\mathfrak{E}_t} + |R_4|1_{\mathfrak{E}_t}$$

For each term, we can bound as:

$$\begin{aligned} |R_1 - R_2|1_{\mathfrak{E}_t} &\leq \eta O(d^2) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty O(\chi) \leq \eta O(d^2 \chi^2 \kappa) \\ |R_2 - R_3|1_{\mathfrak{E}_t} &\leq \eta^2 O(d^4) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^2 O(\chi) \leq \eta O(d^2 \chi^2 \kappa) \\ |R_3 - R_4|1_{\mathfrak{E}_t} &\leq \eta^3 O(d^6) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^3 O(\chi) \leq \eta O(d^2 \chi^2 \kappa) \\ |R_4|1_{\mathfrak{E}_t} &\leq \eta^4 O(d^8) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^4 O(\chi) \leq \eta O(d^2 \chi^2 \kappa) \end{aligned}$$

This gives in sum that, with probability 1:

$$|R_1|1_{\mathfrak{E}_t} \leq \eta O(d^2 \chi^2 \kappa)1_{\mathfrak{E}_t} = \eta O(\mu^2 k^2 \kappa^5)1_{\mathfrak{E}_t}$$

For the third inequality $\mathbb{E}[R_1^2 1_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(\frac{\mu^3 k^3 \kappa^6}{d})1_{\mathfrak{E}_t}$:

$$\mathbb{E}R_1^2 1_{\mathfrak{E}_t} \leq 4 [\mathbb{E}(R_1 - R_2)^2 1_{\mathfrak{E}_t} + \mathbb{E}(R_2 - R_3)^2 1_{\mathfrak{E}_t} + \mathbb{E}(R_3 - R_4)^2 1_{\mathfrak{E}_t} + \mathbb{E}R_4^2 1_{\mathfrak{E}_t}]$$

For each term, we can bound as:

$$\begin{aligned} \mathbb{E}(R_1 - R_2)^2 1_{\mathfrak{E}_t} &\leq \eta^2 O(d^2) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^2 \left(\delta_{il} O(\chi^2) + (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{lj} (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{il}^2 \right) \\ &\leq \eta^2 O(d^2) \max_{l'} \left\| \mathbf{e}_{l'}^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \right\|^2 \sum_i \left(\delta_{il} O(\chi^2) + (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{lj} (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{il}^2 \right) \leq \eta^2 O(d^2 \chi^3) \\ \mathbb{E}(R_2 - R_3)^2 1_{\mathfrak{E}_t} &\leq \eta^4 O(d^6) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^4 \left(\delta_{il} O(\chi^2) + (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{li}^4 \right) \\ &\leq \eta^4 O(d^6) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^2 \max_{l'} \left\| \mathbf{e}_{l'}^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \right\|^2 \sum_i \left(\delta_{il} O(\chi^2) + (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{li}^4 \right) \leq \eta^2 O(d^2 \chi^3) \\ \mathbb{E}(R_3 - R_4)^2 1_{\mathfrak{E}_t} &\leq \eta^6 O(d^{10}) \sum_{ij} |\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}|_{ij}^6 \delta_{il} O(\chi^2) \\ &\leq \eta^6 O(d^{10}) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^4 \left\| \mathbf{e}_l^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \right\|^2 O(\chi^2) \leq \eta^2 O(d^2 \chi^3) \\ \mathbb{E}R_4^2 1_{\mathfrak{E}_t} &\leq \eta^8 O(d^{14}) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^8 \delta_{il} O(\chi^2) \\ &\leq \eta^8 O(d^{14}) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^6 \left\| \mathbf{e}_l^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \right\|^2 O(\chi^2) \leq \eta^2 O(d^2 \chi^3) \end{aligned}$$

This gives in sum that:

$$\mathbb{E}R_1^2 1_{\mathfrak{E}_t} \leq \eta^2 O(d^2 \chi^3)1_{\mathfrak{E}_t} = \eta^2 O(\frac{\mu^3 k^3 \kappa^6}{d})1_{\mathfrak{E}_t}$$

This finishes the proof. □

Proof of Proposition C.6. Similarly to the proof of Proposition C.5, we hide the term $1_{\mathfrak{E}_t}$ in all following equations. Reader should always think every term in this proof multiplied by $1_{\mathfrak{E}_t}$. Recall that:

$$\begin{aligned}
f(\tilde{\mathbf{U}}_{t+1}, \tilde{\mathbf{V}}_{t+1}) &= f(\mathbf{U}_{t+1}, \mathbf{V}_{t+1}) = f(\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}}, \tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) \\
&= \text{tr} \left([(\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}})(\tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) - \mathbf{M}] [(\tilde{\mathbf{U}}_t + \Delta_{\mathbf{U}})(\tilde{\mathbf{V}}_t + \Delta_{\mathbf{V}}) - \mathbf{M}]^\top \right) \\
&= f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + 2\text{tr}(\Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})^\top) + 2\text{tr}(\Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})) + Q_2 \\
&= f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) + Q_1
\end{aligned}$$

By expanding the polynomial, we can write out the first order term:

$$\begin{aligned}
Q_1 - Q_2 &= 2\text{tr}(\Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})^\top) + 2\text{tr}(\Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})) \\
&= -4\eta d^2 (\mathbf{U}\mathbf{V}^\top - \mathbf{M})_{ij} \left(\mathbf{e}_j^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})^\top \mathbf{e}_i + \mathbf{e}_i^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \mathbf{e}_j \right)
\end{aligned}$$

The second order term:

$$\begin{aligned}
Q_2 - Q_3 &= -\text{tr}(\Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top \tilde{\mathbf{V}}_t \Delta_{\mathbf{U}}^\top) + \text{tr}(\tilde{\mathbf{U}}_t \Delta_{\mathbf{V}}^\top \Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t^\top) + 2\text{tr}(\Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top \Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t^\top) + 2\text{tr}(\Delta_{\mathbf{U}} \Delta_{\mathbf{V}}^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})^\top) \\
&= 4\eta^2 d^4 (\mathbf{U}\mathbf{V}^\top - \mathbf{M})_{ij}^2 \\
&\quad \cdot \left(\left\| \mathbf{e}_j^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top \right\|^2 + \left\| \mathbf{e}_i^\top \tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top \right\|^2 + (\tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top)_{jj} (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{ii} + (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij} \right)
\end{aligned}$$

The third order term:

$$\begin{aligned}
Q_3 - Q_4 &= 2\text{tr}(\Delta_{\mathbf{U}} \tilde{\mathbf{V}}_t^\top \Delta_{\mathbf{V}} \Delta_{\mathbf{U}}^\top) + 2\text{tr}(\Delta_{\mathbf{U}} \Delta_{\mathbf{V}}^\top \Delta_{\mathbf{V}} \tilde{\mathbf{U}}_t^\top) \\
&= -16\eta^3 d^6 (\mathbf{U}\mathbf{V}^\top - \mathbf{M})_{ij}^3 \left((\tilde{\mathbf{V}}_t \tilde{\mathbf{V}}_t^\top)_{jj} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij} + (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{U}}_t^\top)_{ii} \right)
\end{aligned}$$

The fourth order term:

$$Q_4 = \text{tr}(\Delta_{\mathbf{U}} \Delta_{\mathbf{V}}^\top \Delta_{\mathbf{V}} \Delta_{\mathbf{U}}^\top) = 16\eta^4 d^8 (\mathbf{U}\mathbf{V}^\top - \mathbf{M})_{ij}^4 (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top)_{ij}^2$$

Again, in addition to equation (23), (23), (24), (25), we also need following inequality:

$$\begin{aligned}
&\left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty 1_{\mathfrak{E}_t} = \max_{ij} |\text{tr}(\mathbf{e}_i^\top (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \mathbf{e}_j)| 1_{\mathfrak{E}_t} \\
&= \max_{ij} |\text{tr}(\mathbf{e}_i^\top (\mathcal{P}_{\mathbf{X}} + \mathcal{P}_{\mathbf{X}_\perp})(\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \mathbf{e}_j)| 1_{\mathfrak{E}_t} \\
&\leq \max_{ij} |\text{tr}(\mathbf{e}_i^\top \mathcal{P}_{\mathbf{X}}(\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}) \mathbf{e}_j)| 1_{\mathfrak{E}_t} + \max_{ij} |\text{tr}(\mathbf{e}_i^\top \mathcal{P}_{\mathbf{X}_\perp} \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top \mathbf{e}_j)| 1_{\mathfrak{E}_t} \\
&\leq \max_i \left\| \mathbf{e}_i^\top \mathbf{X} \right\| \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_F 1_{\mathfrak{E}_t} + \max_j \left\| \mathbf{e}_j^\top \mathbf{W}_{\mathbf{V}} \right\| \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_F 1_{\mathfrak{E}_t} \\
&\leq O(\kappa \sqrt{\chi}) \sqrt{f(\mathbf{U}_t)} \tag{27}
\end{aligned}$$

Now we are ready to prove Lemma.

For the first inequality $\mathbb{E}[Q_2 1_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(\mu d k \kappa^2) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t}$:

$$\mathbb{E}[Q_2 1_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \mathbb{E}[|Q_2 - Q_3| 1_{\mathfrak{E}_t} | \mathfrak{F}_t] + \mathbb{E}[|Q_3 - Q_4| 1_{\mathfrak{E}_t} | \mathfrak{F}_t] + \mathbb{E}[|Q_4| 1_{\mathfrak{E}_t} | \mathfrak{F}_t]$$

For each term, we can bound as:

$$\begin{aligned} \mathbb{E}[|Q_2 - Q_3| 1_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^2 O(d^2) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^2 O(\chi) = \eta^2 O(d^2 \chi) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\ \mathbb{E}[|Q_3 - Q_4| 1_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^3 O(d^4) \sum_{ij} |\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}|_{ij}^3 O(\chi) \\ &\leq \eta^3 O(d^4 \chi) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \leq \eta^2 O(d^2 \chi) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\ \mathbb{E}[|Q_4| 1_{\mathfrak{E}_t} | \mathfrak{F}_t] &\leq \eta^4 O(d^6) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^4 O(\chi) \\ &\leq \eta^4 O(d^6 \chi) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^2 f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \leq \eta^2 O(d^2 \chi) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \end{aligned}$$

This gives in sum that

$$\mathbb{E}[Q_2 1_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(d^2 \chi) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t} = \eta^2 O(\mu d k \kappa^2) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t}$$

For the second inequality $|Q_1| 1_{\mathfrak{E}_t} \leq \eta O(\mu d k \kappa^3) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t}$ **w.p 1**:

$$|Q_1| 1_{\mathfrak{E}_t} \leq |Q_1 - Q_2| 1_{\mathfrak{E}_t} + |Q_2 - Q_3| 1_{\mathfrak{E}_t} + |Q_3 - Q_4| 1_{\mathfrak{E}_t} + |Q_4| 1_{\mathfrak{E}_t}$$

For each term, we can bound as:

$$\begin{aligned} |Q_1 - Q_2| 1_{\mathfrak{E}_t} &\leq \eta O(d^2) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_F O(\sqrt{\chi}) \leq \eta O(d^2 \chi \kappa) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\ |Q_2 - Q_3| 1_{\mathfrak{E}_t} &\leq \eta^2 O(d^4) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^2 O(\chi) \leq \eta^2 O(d^4 \chi^2 \kappa^2) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) = \eta O(d^2 \chi \kappa) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\ |Q_3 - Q_4| 1_{\mathfrak{E}_t} &\leq \eta^3 O(d^6) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^3 O(\chi) \leq \eta O(d^2 \chi \kappa) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\ |Q_4| 1_{\mathfrak{E}_t} &\leq \eta^4 O(d^8) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^4 O(\chi) \leq \eta O(d^2 \chi \kappa) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \end{aligned}$$

This gives in sum that, with probability 1:

$$|Q_1| 1_{\mathfrak{E}_t} \leq \eta O(d^2 \chi \kappa) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t} = \eta O(\mu d k \kappa^3) f(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t}$$

For the third inequality $\mathbb{E}[Q_1^2 1_{\mathfrak{E}_t} | \mathfrak{F}_t] \leq \eta^2 O(\mu d k \kappa^2) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t}$:

$$\mathbb{E}Q_1^2 1_{\mathfrak{E}_t} \leq 4 [\mathbb{E}(Q_1 - Q_2)^2 1_{\mathfrak{E}_t} + \mathbb{E}(Q_2 - Q_3)^2 1_{\mathfrak{E}_t} + \mathbb{E}(Q_3 - Q_4)^2 1_{\mathfrak{E}_t} + \mathbb{E}Q_4^2 1_{\mathfrak{E}_t}]$$

For each term, we can bound as:

$$\begin{aligned}
\mathbb{E}(Q_1 - Q_2)^2 1_{\mathfrak{E}_t} &\leq \eta^2 O(d^2) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^2 \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_F^2 O(\chi) \leq \eta^2 O(d^2 \chi) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\
\mathbb{E}(Q_2 - Q_3)^2 1_{\mathfrak{E}_t} &\leq \eta^4 O(d^6) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^4 O(\chi^2) \\
&\leq \eta^4 O(d^6 \chi^2) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^2 \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_F^2 \leq \eta^4 O(d^6 \chi^3 \kappa^2) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\
&\leq \eta^2 O(d^2 \chi) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\
\mathbb{E}(Q_3 - Q_4)^2 1_{\mathfrak{E}_t} &\leq \eta^6 O(d^{10}) \sum_{ij} |\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M}|_{ij}^6 O(\chi^2) \\
&\leq \eta^6 O(d^{10} \chi^2) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^4 \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_F^2 \leq \eta^2 O(d^2 \chi) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) \\
\mathbb{E}Q_4^2 1_{\mathfrak{E}_t} &\leq \eta^8 O(d^{14}) \sum_{ij} (\tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M})_{ij}^8 \delta_{il} O(\chi^2) \\
&\leq \eta^8 O(d^{14} \chi^2) \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_\infty^6 \left\| \tilde{\mathbf{U}}_t \tilde{\mathbf{V}}_t^\top - \mathbf{M} \right\|_F^2 \leq \eta^2 O(d^2 \chi) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t)
\end{aligned}$$

This gives in sum that:

$$\mathbb{E}Q_1^2 1_{\mathfrak{E}_t} \leq \eta^2 O(d^2 \chi) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t} = \eta^2 O(\mu d k \kappa^2) f^2(\tilde{\mathbf{U}}_t, \tilde{\mathbf{V}}_t) 1_{\mathfrak{E}_t}$$

This finishes the proof. □