

Fast Global Convergence for Low-rank Matrix Recovery via Riemannian Gradient Descent with Random Initialization

THOMAS Y. HOU, ZHENZHEN LI, ZIYUN ZHANG

January 1, 2021

Abstract

In this paper, we propose a new global analysis framework for a class of low-rank matrix recovery problems on the Riemannian manifold. We analyze the global behavior for the Riemannian optimization with random initialization. We use the Riemannian gradient descent algorithm to minimize a least squares loss function, and study the asymptotic behavior as well as the exact convergence rate. We reveal a previously unknown geometric property of the low-rank matrix manifold, which is the existence of spurious critical points for the simple least squares function on the manifold. We show that under some assumptions, the Riemannian gradient descent starting from a random initialization with high probability avoids these spurious critical points and only converges to the ground truth in nearly linear convergence rate, i.e. $\mathcal{O}(\log(\frac{1}{\epsilon}) + \log(n))$ iterations to reach an ϵ -accurate solution. We use two applications as examples for our global analysis. The first one is a rank-1 matrix recovery problem. The second one is the Gaussian phase retrieval problem. The second example only satisfies the weak isometry property, but has behavior similar to that of the first one except for an extra saddle set. Our convergence guarantee is nearly optimal and almost dimension-free, which fully explains the numerical observations. The global analysis can be potentially extended to other data problems with random measurement structures and empirical least squares loss functions.

1 Introduction

Low-rank matrix recovery problems have been extensively studied in machine learning, signal processing, imaging science, advanced statistics, information theory and quantum mechanics, etc. Related problems include but are not limited to matrix factorization, matrix sensing, matrix completion, phase retrieval, robust PCA, density matrix detection, subspace clustering, to name a few.

Earlier approaches on these problems include convex relaxation methods [11, 12, 13, 14, 40], which convexify the problems by replacing rank minimization with nuclear norm minimization, and achieve exact recovery with provable guarantee. More recent works have shifted focus to non-convex methods due to their lighter computational cost. Specifically, those methods mainly depend on the Burer-Monteiro factorization (e.g. parameterize a rank- r matrix $Z = (U, V) = UV^T$, with $U, V \in \mathbb{F}^{n \times r}$) and the global landscape analysis of the corresponding non-convex objective function [5, 6, 22, 23, 24, 25, 34, 42]. For example, for some specific machine learning problems, it has been shown that the landscape does not have spurious local

minima. This property, combined with the convergence results in non-convex optimization [23, 29], leads to convergence guarantee for second-order stationary points. We refer the reader to the excellent survey papers [17, 26, 28, 40] on low-rank matrix recovery.

Despite such progress, many questions remain unanswered. In particular, the convergence rate results obtained by those non-convex methods are largely sub-linear, in sharp contrast to numerical observations indicating fast and nearly linear convergence. Furthermore, the Burer-Monteiro factorization [8, 9] may result in uncountably many artificial critical points. Lastly, individual problems are usually handled case by case without a full understanding of the intrinsic mechanism behind these problems.

To understand these questions, we propose a unified global analysis framework from the viewpoint of a Riemannian manifold. Instead of factorizing the low-rank matrix, we impose the low-rank constraint by restricting the domain to the low-rank matrix manifold $\mathcal{M}_r := \{Z : \text{rank}(Z) = r, Z \in \mathbb{F}^{n_1 \times n_2}\}$, $\mathbb{F} = \mathbb{R}$ or \mathbb{C} . We consider the following optimization problem over this low-rank manifold:

$$\min_{Z \in \mathcal{M}_r} f(Z) = \frac{1}{2} \|T(Z) - y\|_2^2, \quad (1)$$

where $T : \mathcal{M}_r \rightarrow \mathbb{R}^m$ is a linear operator, $T(Z) = \frac{1}{\sqrt{m}} (\langle A_1, Z \rangle, \dots, \langle A_m, Z \rangle)^\top$, and $y \in \mathbb{R}^m$ with $y_j = \frac{1}{\sqrt{m}} \langle A_j, X \rangle$. The formulation is general and covers many different low-rank matrix recovery problems, as shown in the following examples.

Example. We give a few specific examples of the operator T and measurements $\{A_j\}_{j=1}^m$.

- 1) Matrix sensing: $T : \mathcal{M}_r \rightarrow \mathbb{R}^m$, where $\{A_j\}_{j=1}^m \subset \mathbb{F}^{n_1 \times n_2}$ have entries drawn i.i.d from $\mathcal{N}(0, 1)$, if $\mathbb{F} = \mathbb{R}$; and $\frac{\sqrt{2}}{2} \mathcal{N}(0, 1) + i \frac{\sqrt{2}}{2} \mathcal{N}(0, 1)$, if $\mathbb{F} = \mathbb{C}$.
- 2) Matrix completion: $T : \mathcal{M}_r \rightarrow \mathbb{R}^m$, where $\{A_j\}_{j=1}^m \subset \mathbb{F}^{n_1 \times n_2}$ are generated by a uniform sampling of indices $\Omega \subset [n_1] \times [n_2]$ of a $n_1 \times n_2$ matrix. The matrix A_j is the indicator of the j -th sampled entry, i.e. with value 1 in the sampled indices and 0's in the other indices.
- 3) Gaussian phase retrieval: $T : \mathcal{M}_1 \rightarrow \mathbb{R}^m$, where $\{A_j\}_{j=1}^m \subset \mathbb{F}^{n \times n}$ are also rank-1 matrices. In the real case, $A_j = a_j a_j^\top$ with $a_j \in \mathbb{R}^n$ and their entries are drawn i.i.d from $\mathcal{N}(0, 1)$; in the complex case, $A_j = a_j a_j^*$ with $a_j \in \mathbb{C}^n$ and their entries are drawn i.i.d from $\frac{\sqrt{2}}{2} \mathcal{N}(0, 1) + i \frac{\sqrt{2}}{2} \mathcal{N}(0, 1)$.

All the above examples can be considered as random sensing of a low-rank matrix, where T is a linear operator and A_j 's are drawn from some random distribution. There are also works on the number of measurements and their distributions in order to guarantee successful recovery. Despite the difference in the problem setting in these models and the distribution of A_j 's, their population problems share some common properties. For this reason, we will focus on the global convergence behavior of the population problem in this paper. More specifically, the population loss of matrix sensing and matrix completion is $\mathbb{E}f(Z) = c \|Z - X\|_F^2$ for some positive constant $c > 0$, while that of the phase retrieval problem is $\mathbb{E}f(Z) = c \|Z - X\|_F^2 + \frac{1}{2} (\|Z\|_F - \|X\|_F)^2$ with $c = 1$ or $\frac{1}{2}$ (see Theorem 2.3). Thus we study the

asymptotic behavior and the convergence rate of a sequence generated by randomly initialized¹ Riemannian gradient descent (the implementation of projected gradient descent defined in (6), with details in Appendix A.2) minimizing

$$\min_{Z \in \mathcal{M}_r} F_1(Z) = \frac{1}{2} \|Z - X\|_F^2, \quad (2)$$

$$\text{or } \min_{Z \in \mathcal{M}_r} F_2(Z) = \frac{\theta}{2} (\|Z\|_F - \|X\|_F)^2 + \|Z - X\|_F^2. \quad (3)$$

In other words, we analyze the isometry case $F_1(z) = \|Z - X\|_F^2$, and the weak isometry case $C_1 \|Z - X\|_F^2 \leq F_2(Z) \leq C_2 \|Z - X\|_F^2$, where $C_2 > C_1 > 0$. Many low-rank matrix recovery problems belong to the former category, while the phase retrieval problem falls into the latter category (for more details see Section 2.2).

The main results of our global analysis framework are as follows.

Theorem (Informal version of Theorem 2.1). *With high probability no less than $1 - \frac{1}{\text{poly}(n)}$, the sequence generated by a randomly initialized Riemannian gradient descent needs $\mathcal{O}(\log n + \log \frac{1}{\epsilon})$ iterations to reach an ϵ -accurate solution of X in minimizing (2).²*

Theorem (Informal version of Theorem 2.3). *For a class of specific weak isometry problems (3), with high probability no less than $1 - \frac{1}{\text{poly}(n)}$, the sequence generated by a random initialized Riemannian gradient descent needs $\mathcal{O}(\log n + \log \frac{1}{\epsilon})$ iterations to reach an ϵ -accurate solution of X .*

The above results provide a partial explanation for the mechanism behind the nearly linear convergence rate of vanilla first-order methods. The $\mathcal{O}(\log n)$ term in the number of iterations is mainly due to the fact that these problems have benign saddle points or saddle-like spurious points, and the sequence can escape these saddle points or saddle-like points in $\mathcal{O}(\log n)$ iterations. Under the setting of our problems, the vanilla Riemannian first-order scheme with random initialization converges to the second-order stationary points in a nearly linear rate essentially independent of the dimensionality of the problem with high probability.


Numerical illustration. In Figure 1, we give some representative numerical results obtained using the Riemannian gradient descent (defined (6) with details in Appendix A.2) with random initialization on the manifold to minimize the least squares loss functions of three problems. We observe nearly linear convergence in all three experiments. It is consistent with the theoretical results stated above. Note that the nearly linear convergence rate here is in sharp contrast with the sub-linear theoretical guarantees from previous works. Our goal is thus to bridge this gap between theory and practice.

1.1 Comparison and challenges

Burer-Monteiro parametrization. Many low-rank matrix recovery problems are based on the Burer-Monteiro parameterization [6, 8, 9, 25, 42], which parameterizes a low-rank matrix $Z \in$

¹Randomly initialized means the initialization is drawn from the general random distribution (defined in Definition 3.12).

² $\Omega(N)$ means that there exist constants $C \geq c > 0$ such that $c \cdot N \leq \Omega(N) \leq C \cdot N$, while $\mathcal{O}(N)$ means that there exist constant $C > 0$ such that $\mathcal{O}(N) \leq C \cdot N$.



(a) *Least squares function* with $F_1(Z) = \frac{1}{2} \|Z - X\|_F^2$, $n=200$, $r=10$.

(b) *Matrix sensing* with $f(Z) = \frac{1}{2} \|T(Z) - y\|_2^2$, $n=100$, $r=5$, $m=2500$.

(c) *Phase retrieval* with $f(Z) = \frac{1}{2} \|T(Z) - y\|_2^2$, $n=100$, $m=1000$.

Figure 1: The PGD with random initialization. Each log-error band stands for the results from 100 independent experiments.

$\mathbb{F}^{n \times n}$ by UU^H , with $U \in \mathbb{F}^{n \times r}$.³ It focuses on the least squares loss function as follows:

$$\min_{U \in \mathbb{F}^{n \times r}} \frac{1}{2} \|T(UU^H) - y\|_2^2. \quad (4)$$

In contrast, we study a different formulation as follows

$$\min_{Z \in \mathcal{M}_r} \frac{1}{2} \|T(Z) - y\|_2^2, \quad (5)$$

and perform analysis directly in the manifold domain of Z instead of the Euclidean domain of U . Problem (5) has a few advantages over Problem (4). Firstly, the Burer-Monteiro factorization could result in many more artificial critical points due to the reformulation. To see this, if $X = UU^H$ is a critical point, then $X = (UP)(UP)^H$ (the exponent H denotes the Hermitian adjoint) with any unitary $P \in \mathbb{F}^{r \times r}$ is also a critical point. Such duplication of critical points can cause troubles both in the asymptotic behavior and the convergence rate [21, 27]. Secondly, in many machine learning problems, the low-dimensional structure does not always translate into an appropriate factorization. Thus the global analysis on the low-rank matrix manifold provides a first step towards a broader extension to other applications in data science. Lastly, the objective function of (1) is quadratic, while that of (4) is 4th-degree in terms of U . The former is more desirable in analysis.

Rank-1 versus rank-r. There are plenty of works that explore the asymptotic landscape and exact convergence rate for rank-1 problems [16, 47]. Our results differ from these previous works in that we provide a unified framework of analysis that applies to the general rank-r problem (Theorem 2.1), which is much more challenging than the rank-1 problem. To see the major technical challenge for general rank-r problems, note that while for the rank-1 problem, the core matrix is an 1×1 matrix (a scalar), for general rank-r problem it becomes an $r \times r$ matrix. The closed form solutions of some quantities (e.g. the angles between the column spaces) are no longer available, which adds considerable difficulty to the convergence analysis.

³We assume that X and Z are SPSD (symmetric positive semidefinite). For a treatment of the asymmetric case, we refer to [24].

Convergence rate. Earlier landscape analysis on the low-rank matrix recovery [6, 22, 24, 25, 34, 42], combined with the convergence guarantee for the non-convex optimization [23, 29], indicates polynomial convergence towards the second-order stationary point. More recently, the authors in [16] achieved nearly linear convergence for the rank-1 phase retrieval problem.

We point out that the result of [16] is consistent with our global analysis results (Section 2.2). In addition, we stress that the nearly linear convergence is common among other low-rank matrix recovery problems, at least from the manifold optimization perspective. Our work explores the following aspects: (1) whether the nearly optimal and fast convergence rate can be proved for the general rank- r matrix recovery; (2) how a weak isometry property affects the results; and (3) what is the common mechanism behind many different kinds of low-rank matrix recovery problems.

1.2 Insights and discussions

The power of randomness. First-order schemes are widely used in large-scale computation due to their light computational cost. Under certain assumptions, first-order schemes are shown to have fast local convergence to the neighborhood of stationary points [36, 37]. A common problem with first-order schemes is that undesirable critical points including saddle points and local maxima could occur. Due to the lack of geometry information around critical points, first-order schemes with bad initialization could be trapped around the undesirable critical points instead of converging to the local minima. However, when augmented with randomness, the first-order schemes work well and have some provable guarantees. Below we discuss two major ways of incorporating randomness, namely the randomly perturbed first-order schemes and the randomly initialized first-order schemes.

- *Perturbed first-order schemes.* There are a few studies on the convergence of perturbed first-order schemes towards second-order stationary points both in the Euclidean and the Riemannian settings, see [18, 21, 23, 29, 43]. These results show that general global convergence rate is polynomial and almost dimension-free. Whereas the intermittent perturbations help perturbed schemes escape the saddles better than non-perturbed first-order schemes in the worst case [21], they also prevent a very accurate approximation of the ground truth without further (and sometimes complicated) modifications.
- *Randomly initialized first-order schemes.* Though it has been proved that randomly initialized gradient descent asymptotically escapes saddles and only converges to the local minima [27, 32, 33, 39], its convergence rate is much less clear. In the worst case, when the initialization is close to the stable manifold of saddle points, the convergence towards the local minima slows down substantially. Indeed, the authors of the previous work [21] show that, in the worst case, the randomly initialized gradient descent can take exponential time to escape from the saddles. Despite such worst case scenario, the optimal efficiency of saddle escape behavior in a more general sense remains unclear. A recent answer to this question is given by the authors of [16], who show that for the rank-1 phase retrieval problem, gradient descent with random initialization has a nearly linear and almost dimension-free convergence rate, improving upon the previous polynomial

convergence rate. This motivates us to study the mechanism behind the fast convergence rate and establish similar results for general rank- r matrix recovery problems.

We point out that an alternative way to avoid saddle points or local maxima is to use Hessian-based schemes and make use of second-order geometry information around the critical points. But the computational cost of such methods is expensive and its implementation can be complicated.

Population loss versus finite-sample. Previous works [16] have approached the heuristic treatment of the finite-sample loss function by splitting it into the population loss function plus the deviation from the population. We adopt a similar viewpoint on the Riemannian manifold.

For example, using the Riemannian gradient descent with random initialization to minimize

$$f(Z) = \frac{1}{2} \|T(Z) - T(X)\|_2^2, \quad Z \in \mathcal{M}_r,$$

one can rewrite the empirical least-square loss function and its gradient into

$$\begin{aligned} f(Z) &= F(Z) + \frac{1}{m} \sum_{j=1}^m \theta_j, & \theta_j &= \langle A_j, Z - X \rangle^2 - \mathbb{E} \left(\langle A_j, Z - X \rangle^2 \right); \\ \nabla f(Z) &= \nabla F(Z) + \frac{1}{m} \sum_{j=1}^m \eta_j, & \eta_j &= \nabla \theta_j. \end{aligned}$$

Here, $F(Z) := \mathbb{E}(f(Z))$, and θ_j 's and η_j 's are i.i.d mean-zero random variables (deviation terms). By the central limit theorem, heuristically, with high probability we have

$$\frac{1}{m} \sum_{j=1}^m \eta_j \lesssim \frac{\text{poly}(\log m)}{\sqrt{m}}.$$

By requiring $m \gtrsim n \cdot \text{poly}(\log n)$, one can control the deviations $\sum_{j=1}^m \eta_j$. Then the minimization of the empirical loss function $f(Z)$ is close to that of the population loss function $F(Z)$. There are many methods to control the randomness, for example the truncation methods [15, 34]. In this way, we can understand the mechanism behind many random sensing problems by analyzing the global convergence of the population dynamics.

1.3 Organization of this paper

The rest of the paper is organized as follows. In Section 2, we present the main results of this paper, namely the framework of global analysis for three low-rank matrix recovery problems. We show the nearly linear convergence guarantee for the randomly initialized Riemannian gradient descent towards the ground truth, when minimizing different kinds of the population least squares loss functions. Specifically, Section 2.1 gives brief preliminary information, Section 2.2 states the main results, and Section 2.3 highlights the main ideas of the proof. In Section 3, we lay out a full description of the proof strategy for the first main result in four parts. In particular, Section 3.1 describes the trajectory behavior, which serves as the key intermediate steps

for proving the first main result. Section 3.2 introduces the Łojasiewicz convergence tool as a fundamental convergence guarantee. Section 3.3 is devoted to the geometry of the low-rank matrix manifold, where we classify all the critical points of the simple least squares function and study the spurious regions around spurious critical points. Some technical lemmas are given in Section 3.4. Section 4 contains the proofs of the second and third main results, which are similar but considerably simpler than that of the first main result. Finally, in Section 5, we make some concluding remarks. In Appendix A, we present the manifold setting and the light-computational optimization technique. All technical proofs are deferred to Appendix B, and some auxiliary lemmas are provided in Appendix C.

2 Main results

In this section, we introduce the main results of this paper. The optimization problems we are looking at are (2) and (3), namely we minimize either $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$ or $F_2(Z) = \frac{\theta}{2}(\|Z\|_F - \|X\|_F)^2 + \|Z - X\|_F^2$ on the low-rank matrix manifold \mathcal{M}_r . The analysis is based on the case of symmetric positive semi-definite (SPSD), and one can refer to [24] for the treatment of asymmetric case.

2.1 Preliminaries

We first introduce the necessary notations, optimization techniques and assumptions for the statement of our main results.

Notations. Denote \mathcal{M}_s as the fixed rank manifold $\{Z \in \mathbb{F}^{n_1 \times n_2} : \text{rank}(Z) = s\}$, and $\overline{\mathcal{M}_s}$ as its closure. In the symmetric (Hermitian) case, denote $n = n_1 = n_2$. Throughout the paper, $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$. Let $(\cdot)^H$ denote the transpose (adjoint), i.e. $(\cdot)^H = (\cdot)^\top$ when $\mathbb{F} = \mathbb{R}$; and $(\cdot)^H = (\cdot)^*$ when $\mathbb{F} = \mathbb{C}$. The matrix X always denotes the ground truth matrix with $\text{rank}(X) = r$, while Z^* can be any fixed point (or accumulating point) of an algorithm. For integer $s > 0$, denote $[s] = \{1, 2, \dots, s\}$. For any $j \in [r]$, let $\setminus j := [r] \setminus \{j\}$. Denote $X = UDU^H$. For any $j \in [r]$, let $d_j := D(j, j)$, and $U_j := U(:, j)$, where d_j 's are in the descending order. Let $\sigma_j(\cdot)$ denote the j -th largest singular value or eigenvalue of a matrix. Unless otherwise specified, the vector norm we use is $\|\cdot\|_2$ and the matrix norm is $\|\cdot\|_F$. Unless otherwise specified, we use $0 < c, C < \infty$ to denote any absolute constant independent of n in our statement that may vary in different contexts. The symbol $\Omega(N)$ means that there exist constants $C \geq c > 0$ such that $c \cdot N \leq \Omega(N) \leq C \cdot N$, and $\mathcal{O}(N)$ means that there exist a constant $C > 0$ such that $\mathcal{O}(N) \leq C \cdot N$. In this paper, we focus on the large n regime, and other quantities including r and $\sigma_r(X)$ will be treated as constants and will be ignored in the $\mathcal{O}(\cdot)$ and $\Omega(\cdot)$. The symbol $\text{poly}(n)$ stands for a nonnegative quantity upper bounded by $C \cdot n^k$ for some $C > 0$ and $k \in \mathbb{N}_+$.

The optimization technique. We use a Riemannian gradient descent method called the Projected Gradient Descent (the PGD):

$$Z_{k+1} = \mathcal{R}\left(Z_k - \alpha_k P_{T_{Z_k}}(\nabla f(Z_k))\right). \quad (6)$$

Here $P_{T_{Z_k}}$ is the projection onto the tangent space of \mathcal{M} at point Z_k , α_k is the k -th stepsize, and $\mathcal{R} : T_Z \rightarrow \mathcal{M}$ is a retraction operator. This can be viewed as the vanilla first-order scheme on the Riemannian manifold. A more detailed description of the method and the respective operations are deferred to Appendix A.

Assumptions. The following assumptions are necessary for all of the main results.

Assumption 1: Assume $\alpha > 0$ is a small constant such that the discretized system can be well approximated by the continuous system. In the following, we assume that the first-order approximation can provide an accurate approximation for sufficiently small α , and higher order terms are Lipschitz continuous.

Assumption 2: Assume that the singular values of Z are always simple and do not cross one another along the whole gradient flow or gradient descent trajectory.

Remark. Gradient flow of Z with non-crossing singular values is common in practice. In fact, it is *generic*⁴ as is proved in [20]. However, gradient flows with crossing singular values are also observed in some experiments. Assumption 2 is mainly for the purpose of technical analysis, as crossing singular values introduces extra difficulty. It could be overcome by considering subspaces of singular vectors as a whole.

2.2 Statement of main results

We are now ready to state the three main results of this paper.

Main result 1: Global convergence for least squares on \mathcal{M}_r

We will prove that the randomly initialized Riemannian gradient descent (the PGD) for $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$ with high probability escapes the spurious critical points \mathcal{S}_{Z^*} (to be defined in Section 3.3) and converges to the global minimum X . Moreover, with high probability the convergence rate is nearly linear and almost dimension-free, i.e., for $k \gtrsim \mathcal{O}(\log n)$, we have $\|Z_k - X\|_F \leq e^{-ck} \|Z_0 - X\|_F$, with the constant $c > 0$ uniformly bounded from below. Specifically, we have the following theorem.

Theorem 2.1. *Let $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$, and $X = UDU^H$ has distinct eigenvalues. Let $\{Z_k\}_k$ be the sequence generated by the Riemannian gradient descent (the PGD) initialized at Z_0 which is drawn from the general random distribution (as defined in Definition 3.12). Then, with high probability no less than $1 - \frac{1}{\text{poly}(n)}$, we have the following convergence results:*

- 1) *The initial matrix Z_0 falls into the branch such that $\lim_{k \rightarrow \infty} Z_k = X$. Further, there exists constant stepsize $\alpha > 0$ small enough and $\mathcal{T} = \mathcal{O}(\log n)$ such that:*

$$\|Z_k - X\|_F^2 \leq e^{-ck}, \text{ for all } k \geq \mathcal{T}.$$

The positive constant c depends on $\sigma_r(X)$ and α . In other words, it takes no more than $\mathcal{O}(\log \frac{1}{\epsilon} + \log n)$ iterations to get an ϵ -accurate solution (i.e. to achieve $\|Z - X\| \leq \epsilon \|X\|_F$) via the randomly initialized Riemannian gradient descent (the PGD).

⁴A property of a topological space is called *generic* if it holds for a subset of the space which is of the second Baire category.

2) With probability 0, Z_0 falls into the other branches such that $\lim_{k \rightarrow \infty} Z_k = Z^* \in \mathcal{S}_{Z^*}$ (defined in Lemma 3.7).

We will prove Theorem 2.1 by using Theorem 3.1, Theorem 3.2 and Theorem 3.3, and will provide a detailed proof strategy in Section 3.

Stationary points. For $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$, the stationary points of the PGD are either the ground truth X , or

$$\mathcal{S}_{Z^*} := \{Z^* : Z^* = U_1 D_1 V_1^H, \text{ where } U = (U_1, U_2), D = \text{diag}\{D_1, D_2\} \text{ and } V = (V_1, V_2), Z^* \neq X\}.$$

See also Lemma 3.7 for more details of this result. If X has distinct eigenvalues, then $|\mathcal{S}_{Z^*}| = 2^r - 1$. In particular, every point in \mathcal{S}_{Z^*} has a property similar to that of a saddle point (we call it a saddle-like property), except that the Hessian has a $-\infty$ direction, which indicates that the curvature is singular at such point. The neighborhood of \mathcal{S}_{Z^*} is crucial because the Riemannian gradient becomes degenerate.

Attraction by spurious stationary points. The fundamental convergence tool to be presented in Theorem 3.4 ensures linear convergence towards X when minimizing $F_1(Z)$ via the Riemannian gradient descent (the PGD) in most part of the manifold. However, when the sequence gets close to \mathcal{S}_{Z^*} , in the local region $\mathcal{B}(\mathcal{S}_{Z^*}, \delta)$, the convergence slows down and the sequence is attracted to \mathcal{S}_{Z^*} . Although the worst initialization leads to non-escape from \mathcal{S}_{Z^*} (or, say ∞ number of iterations to escape), we estimate the upper bound on the number of iterations needed to escape \mathcal{S}_{Z^*} with high probability.

Estimation of the iterations trapped by \mathcal{S}_{Z^*} . From Lemma 3.13, the “angle” (the product of the two column vector matrices) between the randomly initialized column space and the ground truth column space is of order $\Omega(\frac{1}{n})$. If the angle remains $\Omega(\frac{1}{n})$ when the sequence enters the spurious region $\mathcal{B}(\mathcal{S}_{Z^*}, \delta)$ (defined in Lemma 3.9) and it grows exponentially fast, then it takes $\mathcal{O}(\log n)$ iterations for the angle to become $\Omega(1)$. By Lemma 3.10, this indicates successful escape from the spurious regions. To further understand why it takes $\mathcal{O}(\log n)$ iterations to escape the spurious regions, we use a toy example (Example 3.18) to show that for a general strict saddle, exponential escape behavior is also present. However, the Z^* differs from such a general strict saddle in that the Hessian of Z^* has $-\infty$ directions. Still, we can bound the number of iterations needed to escape from $\mathcal{B}(\mathcal{S}_{Z^*}, \delta)$ by throwing away a small probability measure of points. Such analysis is based on Lemma 3.17, Lemma 3.19 and Lemma 3.21.

Main result 2: Global convergence for least squares on \mathcal{M}_1

The second main result can be regarded as a special case of Main result 1 restricted to $r = 1$. However, the analysis is much easier to follow. We list this as one of the main results because the trajectory behavior of this case is very similar to that of phase retrieval in the third main result. When $r = 1$, by Lemma 3.7, we have $\mathcal{S}_{Z^*} = \{0\}$. In addition, we can directly write out the closed form gradient descent or gradient flow. The result is stated in the following theorem.

Theorem 2.2. (Rank-1 isometry case) Assume X satisfies $\text{rank}(X) = 1$, $\|X\|_F = 1$. We consider $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$ with $Z \in \mathcal{M}_1$. Let $\{Z_k\}$ be the sequence generated by the Riemannian gradient descent (the PGD) initialized at Z_0 which is drawn from the general random distribution. Denote $Z = zz^H$ and $X = xx^H$, $h = \|Z\|_F$ and $\rho = \frac{\langle X, Z \rangle}{\|X\|_F \|Z\|_F}$. Then, we have

1) The evolution dynamics can be described by the following discrete system:

$$\begin{aligned} h_{k+1} &= h_k + \alpha(-h_k + \rho_k) + \mathcal{O}(\alpha^2), \\ \rho_{k+1} &= \rho_k + 2\alpha \frac{\rho_k}{h_k} (1 - \rho_k) + \mathcal{O}(\alpha^2). \end{aligned}$$

2) Further, there exists $\alpha > 0$ small enough such that with high probability no less than $1 - \frac{1}{\text{poly}(n)}$, the Riemannian gradient descent (the PGD) only converges to the equilibrium with $\rho^* = h^* = 1$, meaning $\lim_{k \rightarrow \infty} Z_k = X$. Moreover, it takes $\mathcal{O}(\log \frac{1}{\epsilon} + \log n)$ iterations to generate an ϵ -accurate solution, i.e. to achieve $\|Z - X\|_F \leq \epsilon \|X\|_F$.

Theorem 2.2 is in preparation for the next main result on the global convergence for the special case of phase retrieval. As we will see below, the population loss function of phase retrieval differs from F_1 in that it only satisfies a weak isometry property. A detailed proof of Theorem 2.2 and its connection with Theorem 2.3 can be found in Section 4. The proof is similar to that of Theorem 2.1 but much simpler.

Main result 3: Global convergence for the population phase retrieval problem

Isometry properties weaker than the RIP are also common in various real-world applications. An example is the phase retrieval problem, whose loss function is given below:

$$f(Z) = \frac{1}{2} \|T(Z) - y\|_2^2 = \frac{1}{2m} \sum_{j=1}^m \langle A_j, Z - X \rangle^2.$$

Here, $X = xx^H$ is the ground truth, $y = T(X)$, $Z = zz^H$, and $A_j = a_j a_j^H$, where a_j 's are i.i.d drawn from $\mathcal{N}(0, I_n)$ or $\frac{1}{\sqrt{2}} (\mathcal{N}(0, I_n) + i \cdot \mathcal{N}(0, I_n))$. To simplify the analysis, we only establish the result for the population loss function here. We focus on studying the problem on the Riemannian manifold and revealing its connection to the rank-1 isometry case (Theorem 2.2). Our proof complements that of [16], which establishes a complete proof for random measurements from a different viewpoint.

Theorem 2.3. *For the Gaussian phase retrieval problem (7), we have the following results*

- 1) Let $F_2(Z) := \mathbb{E} f(Z)$ be the population loss of (7). Then, we have $F_2(Z) = \frac{\theta}{2} (\|Z\|_F - \|X\|_F)^2 + c \cdot \|Z - X\|_F^2$, where $\theta = 1$, and $c = 1$ when $\mathbb{F} = \mathbb{R}$, or $c = \frac{1}{2}$ when $\mathbb{F} = \mathbb{C}$.
- 2) Consider $F_2(Z) = \frac{\theta}{2} (\|Z\|_F - \|X\|_F)^2 + \|Z - X\|_F^2$, with $0 < \theta < \Omega(1)$. Let $\{Z_k\}$ be the sequence generated by the Riemannian gradient descent (the PGD) initialized at Z_0 which is drawn from the general random distribution. Denote $Z = zz^H$, $X = xx^H$, $h = \|Z\|_F$ and $\rho = \frac{\langle X, Z \rangle}{\|X\|_F \|Z\|_F}$. Then, the evolution dynamics can be described by the following discrete system:

$$\begin{aligned} h_{k+1} &= h_k + \alpha(\theta - (2 + \theta)h_k + 2\rho_k) + \mathcal{O}(\alpha^2), \\ \rho_{k+1} &= \rho_k + \alpha \frac{4\rho_k}{h_k} (1 - \rho_k) + \mathcal{O}(\alpha^2). \end{aligned}$$

- 3) Further, there exists $\alpha > 0$ small enough such that with high probability no less than $1 - \frac{1}{\text{poly}(n)}$, the Riemannian gradient descent (the PGD) converges to the equilibrium with $\rho^* = h^* = 1$ only, meaning $\lim_{k \rightarrow \infty} Z_k = X$. Moreover, it takes $\mathcal{O}(\log \frac{1}{\epsilon} + \log n)$ iterations to generate an ϵ -accurate solution, i.e. $\|Z - X\|_F \leq \epsilon \|X\|_F$.

The proof of Theorem 2.3 is built upon Theorem 2.2, as F_2 satisfies a weaker isometry property than that of F_1 . The detailed proof is given in Section 4.

2.3 Sketch of proof

Here we highlight a few high level ideas of our proofs for the main theorems stated in the previous subsection. For a more comprehensive and detailed presentation of the proof strategy and intermediate results, the reader may refer to Section 3 and Section 4.

Fundamental convergence guarantee by the Łojasiewicz inequality. The Łojasiewicz inequality [1, 3, 7, 41] has long been studied as a fundamental tool for convergence analysis. It is especially useful in proving the linear convergence rate of first-order optimization methods. In Section 3.2, we derive a version of this tool tailored for the PGD method, stated in Theorem 3.4. Using this theorem, the task of checking the convergence rate of the PGD for a specific problem is reduced to checking Conditions (D) and (L) for the objective function. We then observe that for Problems (2) and (3), these conditions are satisfied except for some small regions on the manifold. Such regions are later dubbed *spurious regions*, see e.g. the blue balls in Figure 2. This leads to the next important result on the geometry of the low-rank matrix manifold.

Geometry of spurious regions over \mathcal{M}_r . The low-rank matrix manifold \mathcal{M}_r is a non-closed set. This fact poses special challenges to our analysis. We find out that for the simple least squares loss function $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$, apart from the ground truth solution $Z = X$, there are a few *spurious critical points* in $\overline{\mathcal{M}_r} \setminus \mathcal{M}_r$, which we denote as \mathcal{S}_{Z^*} . In the neighborhood of these spurious critical points, there are certain *spurious regions* where Condition (L) could fail. In other words, outside the spurious regions, linear convergence rate is guaranteed; while inside the spurious regions, the convergence rate slows down. It then becomes important to characterize how the PGD may escape these regions and converge to the ground truth. A full description of the spurious critical points and spurious regions, along with examples and illustrations, is given in Section 3.3.

Three-stage description of the trajectory behavior. It now remains to study how the trajectory of the randomly initialized gradient descent sequence escapes the spurious regions and converges to the ground truth on the manifold. In Section 3.1, we divide the whole trajectory into three stages. In the worst case, the sequence is dragged towards $Z^* = 0$ in the first stage, then escapes $Z^* = 0$ and approaches some other $Z^* \in \mathcal{S}_{Z^*} \setminus \{0\}$ in the second stage, and finally escapes such Z^* in the third stage. We show that by throwing out a small probability measure, the total number of iterations needed to reach the ϵ -neighborhood of X is bounded by $\mathcal{O}(\log(\frac{1}{\epsilon}) + \log n)$, as stated in the main theorems. This worst case is however rarely observed in numerical experiments; usually, the sequence generated by the PGD simply avoids all spurious critical points and converges straight to the ground truth. Proofs of the three-stage behavior are stated in

Theorem 3.1, Theorem 3.2, Theorem 3.3, and rely on a series of technical lemmas detailed in Section 3.4.




Figure 2: Illustration of the trajectory of the PGD on \mathcal{M}_r .

3 Riemannian global analysis and proof of Theorem 2.1

This section is primarily devoted to the proof of Theorem 2.1. This theorem is the central result of this paper and fully describes the trajectory behavior of the PGD for the least squares loss function on the rank- r matrix manifold \mathcal{M}_r . The proofs of Theorem 2.2 and 2.3 in Section 4 mainly follow that of Theorem 2.1, which serves as a general framework for the global analysis. Their proofs also depend on the technical results in Sections 3.2 and 3.4 but differ in some specific aspects.

We introduce some more notations to be used here. For any $\alpha > 0$, define $\phi_\alpha : \mathcal{M}_r \rightarrow \mathcal{M}_r$ as the iteration function of the PGD with stepsize α , such that $Z_{k+1} = \phi_\alpha(Z_k)$. Let $Z_k = V_{z,k} \Sigma_{z,k} V_{z,k}^H$ and $Z_{k+1} = V_{z,k+1} \Sigma_{z,k+1} V_{z,k+1}^H$ be the eigenvalue decomposition of Z_k and Z_{k+1} respectively. For any $\alpha > 0$ which is a constant stepsize, define $\phi_{\alpha,j,\Sigma_{z,k},V_{z,k}(\cdot,\setminus j)}(\cdot)$ as the iteration function of the column space, such that $\phi_{\alpha,j,\Sigma_{z,k},V_{z,k}(\cdot,\setminus j)}(V_{z,k}(\cdot,j)) = V_{z,k+1}(\cdot,j)$ for $j \in [r]$. In other words, the function maps the j -th column of $V_{z,k}$ (the eigenspace of Z_k) to the j -th column of $V_{z,k+1}$ (the eigenspace of Z_{k+1}). For any $Z \in \mathcal{M}_r$, $\mathcal{B}(Z, \delta) := \{Z' : \|Z - Z'\|_F \leq \delta, Z' \in \mathcal{M}_r\}$ is the neighborhood, while $\mathcal{B}(\cdot, \delta)$ defined in Lemma 3.9 is the “spurious region”. The reason why we need the parameters Σ_z and $V_{z,\setminus j}$ in the definition of $\phi_{\alpha,j,\Sigma_{z,k},V_{z,k}(\cdot,\setminus j)}(\cdot)$ is to ensure that the mapping is well-defined.

3.1 Key intermediate steps for Theorem 2.1

To prove Theorem 2.1, we divide the trajectory of the PGD on the manifold into three stages as follows.

1. **Stage 1:** The iterative sequence starts from a random initialization point. For iteration $k \lesssim \frac{1}{\alpha \delta^2} \log \frac{1}{\delta}$, with high probability exceeding $1 - \frac{1}{\text{poly}(n)}$, the sequence either enters the δ -local region of X or reaches the local region of $Z^* = 0$.
2. **Stage 2:** If the sequence enters the δ -local region of X , it will converge linearly to the target X , so it further takes $\mathcal{O}(\log \frac{1}{\epsilon})$ iterations to generate an ϵ -accurate solution. On the other

hand, if the sequence reaches the δ -neighborhood of $Z^* = 0$, it takes $\mathcal{O}(\log n)$ iterations to escape the local region, and enters the stage 3.

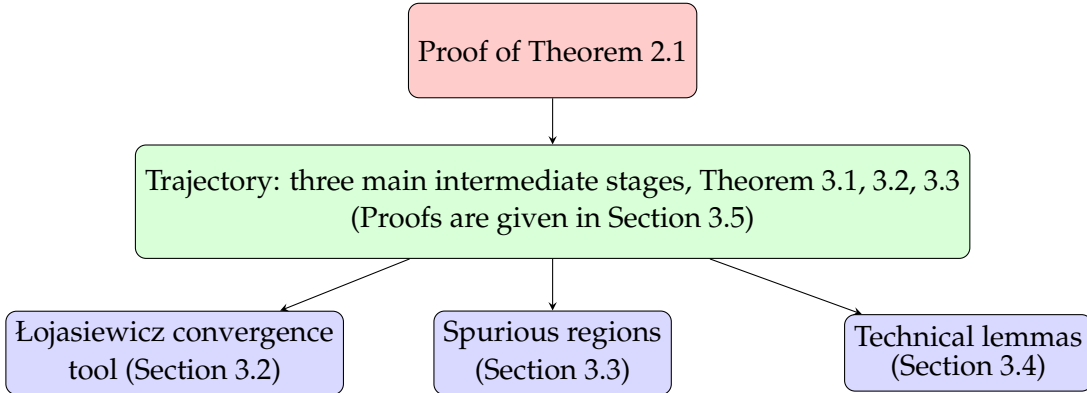
3. **Stage 3:** The sequence either enters the δ -local region of X without getting close to any $Z^* \in \mathcal{S}_{Z^*} \setminus \{0\}$, or, reaches a δ -spurious region of other $Z^* \in \mathcal{S}_{Z^*} \setminus \{0\}$. If it enters the δ -local region of X , it takes $\mathcal{O}(\log \frac{1}{\epsilon})$ iterations to generate an ϵ -accurate solution. On the other hand, if it reaches a δ -spurious region of some $Z^* \in \mathcal{S}_{Z^*} \setminus \{0\}$, then with high probability after $\mathcal{O}(\log n)$ iterations the sequence will escape the δ -spurious region and take additional $\mathcal{O}(\log \frac{1}{\epsilon})$ iterations to generate an ϵ -accurate solution.

The following three theorems state the above three stages in mathematical terms respectively. Their proofs are given in Section 3.5.

Theorem 3.1. *Under the setting of Theorem 2.1, there exists stepsize $\alpha > 0$ and a small enough constant $\delta > 0$ such that with high probability exceeding $1 - \frac{1}{\text{poly}(n)}$, for all $k \leq K_0 = \mathcal{O}(1)$, we have $\phi^k(Z_0) \in \mathcal{B}(X, \delta) \cup \mathcal{B}(0, \delta) \setminus \cup_{Z^* \in \mathcal{S}_{Z^*} \setminus \{0\}} \mathcal{B}(Z^*, \delta)$. Here, the constant in the fail probability depends on α , δ and K_0 .*

Theorem 3.2. *Under the setting of Theorem 2.1, suppose K_0 is a positive integer such that $(\cup_{i=1}^{K_0-1} Z_i) \cap \mathcal{B}(0, \delta) = \emptyset$ and $Z_{K_0} \in \mathcal{B}(0, \delta)$. Then, there exists stepsize $\alpha > 0$, a small enough constant $\delta > 0$ and an absolute constant $C_1 = \mathcal{O}(1) > 0$ such that with high probability exceeding $1 - \frac{1}{\text{poly}(n)}$, we have $\cup_{i=K_1}^{\infty} Z_i \cap \mathcal{B}(0, \delta) = \emptyset$. Here, $K_1 = K_0 + C_1 \log n$.*

Theorem 3.3. *Under the setting of Theorem 2.1, for any $Z^* \in \mathcal{S}_{Z^*} \setminus \{0\}$, suppose there exists a positive integer K_1 such that $\{\cup_{i=1}^{K_1-1} Z_i\} \cap \mathcal{B}(\mathcal{S}_{Z^*} \setminus \{0\}, \delta) = \emptyset$ and $Z_{K_1} \in \mathcal{B}(Z^*, \delta)$. Then, there exists stepsize $\alpha > 0$, a small enough constant $\delta > 0$ and $C_2 = \mathcal{O}(1)$ such that with high probability exceeding $1 - \frac{1}{\text{poly}(n)}$, we have $\cup_{i=K_2}^{\infty} Z_i \cap \mathcal{B}(Z^*, \delta) = \emptyset$. Here, $K_2 = K_1 + C_2 \log n$.*



Proof of Theorem 2.1. Let the stepsize $\alpha > 0$ and $0 < \delta < \frac{\sigma_r(X)}{2}$ be small enough constants that meet the requirements in Theorem 3.1, Theorem 3.2 and Theorem 3.3. Then, by Lemma 3.22, we have linear convergence in the local neighborhood of X . In addition, by Lemma 3.7 and Lemma 3.9, such linear convergence is true in the majority of the manifold except for the spurious regions. In other words, if $Z_k \in \mathcal{M}_r \setminus \mathcal{B}(\mathcal{S}_{Z^*}, \delta)$, then $\|Z_{k+1} - X\|_F \leq e^{-c} \|Z_k - X\|_F$. Here, $c = -\log(1 - \Omega(\alpha\delta^2))$. If for all $k \leq K_0 = \mathcal{O}(\frac{1}{\alpha\delta^2} \log \frac{1}{\delta})$, $Z_k \in \mathcal{M}_r \setminus \mathcal{B}(\mathcal{S}_{Z^*}, \delta)$, then Z_{K_0} enters the

neighborhood of the ground truth X , i.e. $Z_{K_0} \in B(X, \delta)$. By Lemma 3.22, we know that it takes additional $\mathcal{O}(\log \frac{1}{\epsilon})$ iterations to generate an ϵ -accurate solution. On the other hand, if there exists $k \leq K_0 = \mathcal{O}(\frac{1}{\alpha \delta^2} \log \frac{1}{\delta})$ such that Z_k enters the spurious region of Z^* , i.e. $Z_k \in \mathcal{B}(\mathcal{S}_{Z^*}, \delta)$, then we know from Theorem 3.1 that with high probability exceeding $1 - \frac{1}{\text{poly}(n)}$, we have $Z^* = 0$. By Theorem 3.2, it then takes $\mathcal{O}(\log n)$ iterations to escape $\mathcal{B}(0, \delta)$, with fail probability controlled by $\frac{1}{\text{poly}(n)}$. Next, the sequence either reaches $B(X, \delta)$ in $\mathcal{O}(\frac{1}{\alpha \delta^2} \log \frac{1}{\delta})$ iterations, or enters another spurious region $\mathcal{B}(\mathcal{S}_{Z^*}, \delta)$ at the time $K_1 = \mathcal{O}(\log n)$. In the former case, it further takes $\mathcal{O}(\log \frac{1}{\epsilon})$ iterations to generate an ϵ -accurate solution; in the latter case, by Theorem 3.3, it takes $\mathcal{O}(\log n)$ iterations to escape $\mathcal{B}(\mathcal{S}_{Z^*}, \delta)$, with the fail probability controlled by $\frac{1}{\text{poly}(n)}$. Combining the above results, with high probability exceeding $1 - \frac{1}{\text{poly}(n)}$, it takes $\mathcal{O}(\log n + \log \frac{1}{\epsilon})$ iterations for the randomly initialized Riemannian gradient descent (the PGD) to generate an ϵ -accurate solution, i.e. $\|Z - X\|_F < \epsilon \|X\|_F$. \square

The proofs of Theorems 3.1-3.3 will follow the roadmap laid out in Section 2.3. In particular, from Section 3.2 to 3.4, each subsection will introduce a group of technical results corresponding to a main idea in Section 2.3. Specifically, they are the fundamental convergence tool by the Łojasiewicz inequality, the geometry of spurious regions on \mathcal{M}_r , and the dynamics of the trajectory behavior. We then use those technical results to prove Theorems 3.1-3.3 in Section 3.5.

3.2 Fundamental convergence guarantee of the PGD

The *Łojasiewicz inequality* is a powerful tool for analyzing the convergence rate of gradient-based methods, which is named after S. Łojasiewicz [31, 35]. Previous works have used the Łojasiewicz inequality to prove the convergence rate in many Euclidean optimization problems as well as Riemannian optimization problems, e.g. [1, 2, 3, 4, 7, 41, 46], just to name a few.

The following theorem serves as a primary tool to determine the convergence rate of the PGD when minimizing a differentiable function $f(\cdot) : \mathcal{M} \rightarrow \mathbb{R}$. We assume that $\{Z_k\}$ generated by the PGD is bounded for the rest of the paper.

Theorem 3.4. *Let \mathcal{M} be a Riemannian manifold, $f(\cdot) : \mathcal{M} \rightarrow \mathbb{R}$ be a differentiable loss function to be minimized, $\{Z_k\}$ ($k = 1, 2, \dots$) be a sequence generated by the PGD iterations (8). Assume that the following conditions hold:*

- 1) (Descent Inequality) *There exists $C_d > 0$ such that*

$$f_k - f_{k+1} \geq C_d \|P_{T_{Z_k}}(\nabla f(Z_k))\| \|Z_{k+1} - Z_k\|; \quad (\text{D})$$

- 2) (Łojasiewicz Gradient Inequality) *There exists $0 < C_l < +\infty$ such that*

$$|f(Z_k) - f(Z^*)|^{1-\omega} \leq C_l \|P_{T_{Z_k}}(\nabla f(Z_k))\|, \quad (\text{L})$$

with $0 < \omega \leq \frac{1}{2}$. Here, Z^ is the accumulating point for $\{Z_k\}$.*

Then, if the learning rate α satisfies $0 < \alpha < \frac{2C_l^2}{C_d}$, the sequence $\{Z_k\}$ converges to their accumulating point Z^* with the following convergence rate:

$$\|Z_k - Z^*\| \leq \begin{cases} e^{-ck} & , \text{if } \omega = \frac{1}{2}; \\ k^{-\frac{\omega}{1-2\omega}} & , \text{if } 0 < \omega < \frac{1}{2}. \end{cases}$$

When $\omega = \frac{1}{2}$, linear convergence rate can be guaranteed with $c = -\log(1 - \frac{\alpha C_d}{2C_l^2}) > 0$. Here $\|\cdot\|$ is an arbitrary norm under which there is a first order retraction on the manifold.

Remark. A few remarks are in order.

- 1) This theorem explores the convergence rate of the PGD to an accumulating point. This convergence rate depends on the property of the function f (reflected in the exponent ω), and the constants $C_d > 0$, $C_l > 0$ and the learning rate $\alpha > 0$.
- 2) With Theorem 3.4, the task of checking the convergence rate is reduced to checking conditions (D) and (L) and determining the respective constants.
- 3) This theorem only requires Z^* to be an accumulating point, but Z^* is not necessarily a local minimum. It can also be other types of critical points such as saddle points.
- 4) From this result, we can see that the convergence rate is faster with a larger stepsize α , or a steeper Riemannian gradient $\|P_{T_{Z_k}}(\nabla f(Z_k))\|_F$ (which makes C_l smaller).
- 5) An extension of this theorem has been proposed in [48], which allows for mixed norms in (D) and (L).

The two conditions stated in Theorem 3.4 are satisfied in the majority of the manifold \mathcal{M}_r for the loss functions that we consider, including the neighborhood of the ground truth X . Thus fast convergence rate in the majority of the manifold is easy to derive. In the rest of the manifold, though, Condition (L) could fail and the convergence rate could deteriorate. As we will see in the next section, these regions are the *spurious regions* near the *spurious critical points* on the manifold. Special analysis is needed to study how the PGD escapes these spurious regions.

We also mention that another version of Theorem 3.4 is given in Section 4, specially tailored for the objective functions with weak isometry, which will be used in our convergence study of the PGD for phase retrieval.

3.3 Geometry of the spurious regions on the low-rank matrix manifold

In this subsection, we look into the critical points of $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$, which consists of the ground truth X , and the set of spurious critical points denoted as \mathcal{S}_{Z^*} . We study the spurious regions $\mathcal{B}(\mathcal{S}_{Z^*}, \delta)$ near those spurious critical points, where Condition (L) of Theorem 3.4 is violated and special treatment is needed.

3.3.1 Spurious fixed points on $\overline{\mathcal{M}_r}$

Definition 3.5 (Converging set). We define the converging set of a fixed point $Z \in \mathcal{M}_r$ as follows:

$$\mathcal{C}_{Z^*} = \{Z : Z \in \mathcal{M}_r, \text{ and for } \{Z_k\}_{k=1}^{\infty} \text{ generated by the PGD initialized at } Z, \lim_{k \rightarrow \infty} Z_k = Z^*\}.$$

Definition 3.6 (Stable/Spurious fixed points). For a manifold \mathcal{M} with a Lebesgue measure, we define spurious fixed points as the fixed points of an iterative algorithm whose converging sets have zero measure, and stable fixed points as those whose converging sets have positive measure.

Lemma 3.7. Consider the PGD applied to $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$. Assume $X = UDV^H$ is a singular value decomposition of X , where $D \in \mathbb{R}^{r \times r}$ is a non-singular diagonal matrix and $U \in \mathbb{F}^{n_1 \times r}$, $V \in \mathbb{F}^{n_2 \times r}$. Then,

1) There are two types of fixed points: one is the ground truth $Z = X$, and the other consists of the set

$$\mathcal{S}_{Z^*} := \{Z^* : Z^* = U_1 D_1 V_1^H, \text{ where } U = (U_1, U_2), D = \text{diag}\{D_1, D_2\} \text{ and } V = (V_1, V_2), Z^* \neq X\}.$$

2) Specifically, if X has distinct singular values⁵, then \mathcal{S}_{Z^*} has cardinality $|\mathcal{S}_{Z^*}| = 2^r - 1$. Assume that $X = \sum_{i=1}^r d_i u_i v_i^H$, then $\mathcal{S}_{Z^*} = \{Z^* = \sum_{i=1}^r d_i \eta_i u_i v_i^H, \text{ with } \eta \in \{0, 1\}^r \text{ and } \eta \neq (1, 1, \dots, 1)^H\}$.

The set \mathcal{S}_{Z^*} consists of $2^r - 1$ points (including $Z^* = 0$) if the singular values of X are distinct, or contains some submanifolds of \mathcal{M}_r , if at least one singular value has multiplicity more than one. In Theorem 2.1, one can show that with high probability, \mathcal{S}_{Z^*} is the set of spurious fixed points of the PGD, while $Z = X$ is a stable fixed point. We call them ‘‘spurious fixed points’’ because they are not stable and when the tangent cone is taken into consideration they are not true fixed points. Moreover, with high probability the sequence generated by the randomly initialized PGD does not converge to any one of these spurious fixed points. Below is an example.

Example 3.8. Assume that $n = 3, r = 2$. Let

$$X = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad Z_0 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then the $\{Z_k\}_{k=0}^{\infty}$ generated by the PGD and their limit point are given by

$$Z_k = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & (1 - \alpha)^k \end{pmatrix}, \quad Z^* := \lim_{k \rightarrow \infty} Z_k = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

We see that Z^* is a spurious critical point. Note that even though each Z_k is in \mathcal{M}_2 , their limit is in \mathcal{M}_1 . Figure 3 is a visualization of the gradient $\|P_{T_Z}(Z - X)\|_F$ in the neighborhood of a spurious Z^* . We can see that the gradient is essentially singular near Z^* . There is only one direction in which the sequence converges to Z^* . Along other directions, the Riemannian gradient remains large and the PGD will likely slip away. We point out that such property of these spurious fixed points is very similar to that of strict saddle points in many non-convex optimization problems, although they are not exactly the same.

⁵Which means that the eigenvalues of X all have algebraic multiplicity equal to 1.




Figure 3: Magnitude of the gradient in the neighborhood of a spurious fixed point

Example 3.8 demonstrates the importance of considering $\overline{\mathcal{M}_r}$ instead of \mathcal{M}_r itself, and gives us a healthy warning that we cannot assume the PGD always stays in the interior of \mathcal{M}_r and converges to a minimum.

However, numerical evidence shows that for the majority of cases, the PGD almost surely escapes the spurious critical points \mathcal{S}_{Z^*} and converges to the global minimum. In fact, we will prove that with high probability the set of initial points that converge to \mathcal{S}_{Z^*} is a measure-zero lower-dimensional submanifold on \mathcal{M}_r .

3.3.2 Spurious regions in the neighborhood of spurious fixed points

As mentioned in Section 3.2, there are some regions on \mathcal{M}_r that violate the conditions for fast convergence guarantee in Theorem 3.4. In this subsection, we take a detailed look at those regions.

To ensure that Condition (L) holds with $\omega = \frac{1}{2}$, what we essentially need is

$$\|P_{T_Z}(Z - X)\|_F \geq C_L, \quad C_L > 0.$$

However, for some special Z which occupies a small part of the whole domain, the Riemannian gradient $P_{T_Z}(Z - X)$ becomes so small that this lower bound is violated. We use *spurious regions* to refer to the regions where Z violates this lower bound.

The following results show that the spurious regions are in the neighborhood of the spurious fixed points.

Lemma 3.9 (Spurious regions). *Assume that $X = UDV^H$ is the singular value decomposition of X . Let $0 < \delta < \sigma_r(X)/2$. Then the spurious regions can be characterized as follows*

$$\begin{aligned} \mathcal{B}(\mathcal{S}_{Z^*}, \delta) &:= \{Z : \|P_{T_Z}(Z - X)\|_F \leq \delta\} \cap \mathcal{M}_r \cap \left\{Z : \|Z - X\|_F > \frac{\sigma_r(X)}{2}\right\} \\ &= \left\{Z : Z = \begin{pmatrix} B & \tilde{B} \end{pmatrix} \begin{pmatrix} D_1 + E_1 & 0 \\ 0 & E_2 \end{pmatrix} \begin{pmatrix} C^H \\ \tilde{C}^H \end{pmatrix}, \right. \\ &\quad \left. \|E_1\|, \|E_2\|, \|P_B - P_{U_1}\|, \|P_{\tilde{B}} - P_{\tilde{U}}\|, \|P_C - P_{V_1}\|, \|P_{\tilde{C}} - P_{\tilde{V}}\| \leq \mathcal{O}(\delta)\right\} \\ &= \cup_{Z^* \in \mathcal{S}_{Z^*}} \mathcal{B}(Z^*, \delta) \end{aligned}$$

where $U = (U_1, U_2)$ and $V = (V_1, V_2)$ are the $(s, r - s)$ dimensional splitting for some $0 < s < r$; $\tilde{U}, \tilde{V}: \tilde{U} \subset \text{col}(U)^\perp, \tilde{V} \subset \text{col}(V)^\perp, \tilde{U}^H \tilde{U} = \tilde{V}^H \tilde{V} = I_{r-s}$, and $D = \text{diag}(D_1, D_2)$ where D_1 and D_2 are diagonal matrices.

Remark. The intuition behind Lemma 3.9 is that an s -dimensional principal part of Z is “almost aligned” with an s -dimensional principal part of X , and their singular values are close to each other; while the other $(r - s)$ -dimensional part of Z is “almost perpendicular” to the other $(r - s)$ -dimensional part of X , and the singular values of that part of Z is very small.

Lemma 3.10. *For the case when the matrices are symmetric positive semi-definite (SPSD), there exist eigenvalue decompositions $X = UDU^H$, $Z = V_z \Sigma_z V_z^H$ and $U = V_z R + \tilde{V}_z S$, with $\tilde{V}_z \in \text{Col}(V_z)^\perp$. If $Z \in \mathcal{B}(\mathcal{S}_{Z^*}, \delta)$, where $\text{rank}(Z^*) = s$, then we have:*

$$R = \begin{pmatrix} I_s + E_1 & E_2 \\ E_3 & E_4 \end{pmatrix},$$

with $\|E_1\|^{\frac{1}{2}}, \|E_2\|, \|E_3\|$, and $\|E_4\| \leq \mathcal{O}(\delta)$, $\text{rank}(Z^*) = s$, $0 < s < 1$, $S \in \mathbb{F}^{r \times r}$.

Remark. In the case where the eigenvalues of X are not distinct, simply let U_1 be the basis of the best subspace that V_z can capture.

3.3.3 Convergence rate outside the spurious regions

We now show that as long as the spurious regions are excluded, we can establish the linear convergence rate of the PGD, using Conditions (D) and (L) in Theorem 3.4.




Figure 4: Convergence guarantee on the majority of \mathcal{M}_r

Lemma 3.11. *Let $F_1(Z) = \frac{1}{2} \|Z - X\|_F^2$, $X, Z \in \mathcal{M}_r$. If the Z_k 's generated by the PGD remain in a bounded subset of \mathcal{M}_r , and stay in the set $\{Z : \|P_{T_Z}(Z - X)\|_F \geq C_L\} \cup \{Z : \|Z - X\|_F \leq \frac{\sigma_r(X)}{2}\}$, we have:*

- 1) $\|P_{T_Z}(\nabla f(Z_k))\|_F \geq C_1 \|Z_k - X\|_F$, for all k . That is, condition (L) holds with $\omega = \frac{1}{2}$;
- 2) There exists some $C_2 > 0$ such that

$$f_k - f_{k+1} \geq C_2 \|P_{T_{Z_k}}(\nabla f(Z_k))\|_F \|Z_{k+1} - Z_k\|_F.$$

Thus, by Theorem 3.4, there exists $\alpha > 0$ such that the sequence $\{Z_k\}$ generated by the PGD (8) converges to X in a linear convergence rate:

$$\|Z_k - X\|_F \leq e^{-ck}.$$

Here, $c = -\log(1 - \frac{\alpha C_1^2 C_2}{2}) > 0$ is an absolute constant.

Remark. One can take $C_1 = \frac{\sigma_r(X)}{2}$, then c only depends on the condition number of X .

Lemma 3.11 ensures linear convergence to X in the majority of the manifold outside of the spurious regions. The rest of the analysis thus evolves around how the trajectory of the PGD escapes the spurious regions.

3.4 Technical Lemmas

In Section 3.1, we have outlined how the trajectory of the PGD is divided into three stages, corresponding to Theorem 3.1, Theorem 3.2 and Theorem 3.3. In this subsection, we introduce a few technical lemmas needed for the proofs of Theorems 3.1-3.3.

We first define a general random distribution on the low-rank manifold.

Definition 3.12 (General random distribution). *Z is said to be drawn from a general random initialization, if $Z = V_1 \Sigma V_2^H$ where V_1 and V_2 are drawn from a uniform distribution on the Stiefel manifold $\mathbb{V}_r(\mathbb{R}^n)$, and the entries of Σ are drawn independently from a uniform distribution over $[C_1, C_2]$ with $C_2 > C_1 \geq 0$.*

Remark. 1) The simplest example of a general random distribution is the following rank-1 Gaussian sampling. One can construct $Z_0 = cu_0 u_0^H \in \mathbb{F}^{n \times n}$, where $c > 0$ is a constant and u_0 is drawn from $\frac{1}{\sqrt{n}} \mathcal{N}(0, 1)^n$ for $\mathbb{F} = \mathbb{R}$, or $\frac{1}{\sqrt{2n}} \mathcal{N}(0, 1)^n + i \frac{1}{\sqrt{2n}} \mathcal{N}(0, 1)^n$ for $\mathbb{F} = \mathbb{C}$. It is equivalent to $Z_0 = \rho v_0 v_0^H$ with v_0 drawn from $\mathbb{V}_1(\mathbb{R}^n)$ and ρ drawn from $\Omega(\frac{1}{n}) \chi^2(n)$ or $\Omega(\frac{1}{2n}) \chi^2(2n)$.⁶

- 2) In practical computation, a uniform distribution on the Stiefel manifold can be easily constructed as follows. For $\mathbb{F} = \mathbb{R}$, let G be drawn from $\frac{1}{\sqrt{n}} \mathcal{N}(0, 1)^{n \times r}$, and construct $V = G(G^H G)^{-\frac{1}{2}}$. For $\mathbb{F} = \mathbb{C}$, change the law of G to $\frac{1}{\sqrt{2n}} (\mathcal{N}(0, 1) + i \cdot \mathcal{N}(0, 1))^{n \times r}$.
- 3) The density functions of the sampling laws of V and Σ can be extended from constants to more general ones. For example, it suffices to require that $\rho(V) \in [c, C]$ for any $V \in \mathbb{V}_r(\mathbb{R}^n)$, where $C > c > 0$. This allows more flexibility in the initialization.

For any given $s \in [r]$, we have that the marginal distribution of $V_s := V(:, s)$ is an uniform distribution on S^{n-1} . Therefore, for any given $u : \|u\|_2 = 1$, we have $\|V_s^H u\| \gtrsim \frac{1}{\sqrt{n}}$ with high probability, if we ignore the log-factor. Specifically, we have the following lemma.

Lemma 3.13. *Assume $W = (W_1, W_2, \dots, W_r) \sim \text{Uniform}(\mathbb{V}_r(\mathbb{R}^n))$. For any $u_0 \in \mathbb{F}^n$ such that $\|u_0\|_2 = 1$, we have:*

⁶Defined as $\chi^2(s) := \sum_{j=1}^s n_j^2$ where $n_j, j = 1, 2, \dots, s$ are i.i.d standard normal variables.

- 1) $\mathbb{E}(\|u_0^H W\|_2^2) = \frac{r}{n}$;
- 2) $\text{Prob}(|u_0^H W_i|^2 \geq \Omega(\frac{1}{n \log n})) \geq 1 - e^{-\Omega(n)} - \mathcal{O}(\frac{1}{\log n}), i \in [r]$;
- 3) $\text{Prob}(|u_0^H W_i|^2 \geq \Omega(\frac{\log n}{n})) \leq \frac{1}{\text{poly}(n)} + e^{-\Omega(n)}, i \in [r]$;
- 4) $\text{Prob}(|u_0^H W_i|^2 \geq \Omega(\frac{1}{n^p})) \geq 1 - e^{-\Omega(n)} - \mathcal{O}(\frac{1}{n^p}), i \in [r]$.

Lemma 3.13 shows that the general random distribution with high probability captures weak information of order $\Omega(\frac{1}{n})$ of a given column space if we ignore some log factors.

Definition 3.14 (Intrinsic/Spurious branch). *We say that a point $Z_0 \in \mathcal{M}$ belongs to an intrinsic branch if the limit point $\lim_{k \rightarrow \infty} Z_k$ is a stable fixed point, where $\{Z_k\}$ is generated by the PGD with initialization at Z_0 . We say that Z_0 belongs to a spurious branch if the limit point of $\{Z_k\}$ is a spurious fixed point.*

In the following, we introduce a few technical results that describe the dynamics of the singular values and column spaces of the sequence.

Lemma 3.15. *Consider the gradient flow of $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$. Assume $Z = V_z \Sigma_z V_z^H$ and $X = UDU^H$ are the eigenvalue decompositions of the SPSD (symmetric positive semidefinite) matrices Z and X respectively. Denote $U = V_z R + \tilde{V}_z Q$ with $\tilde{V}_z \in \text{Col}(V_z)^\perp$. We have:*

$$\frac{d}{dt} \sum_{j=1}^r \|D^{\frac{1}{2}} R^H e_j\|^2 = \sum_{j=1}^r \frac{2}{\Sigma_z(j, j)} \|QDR^H e_j\|^2.$$

Furthermore, denote the spectra of $R^H R$ and RR^H by $\Sigma_{RR} = \text{diag}\{r_1, r_2, \dots, r_r\}$, we have the following inequality:

$$\frac{d}{dt} \left(\sum_{j=1}^r r_j \right) \gtrsim \sum_{j=1}^r (1 - r_j) r_j. \quad (7)$$

Remark. For all $j \in [r]$, we have $0 \leq r_j \leq 1$, where the r_j 's describe the angle between the column spaces of Z and X . From (7), we can conclude that $\sum_{j=1}^r r_j$ is non-decreasing. The increase of $\sum_{j=1}^r r_j$ only slows down if for all $j \in [r]$, r_j is close to either 0 or 1. Note that $\sum_{j=1}^r r_j$ is an indicator of whether the sequence is close to a spurious critical point. For example, when $Z \in \mathcal{B}(0, \delta)$ and $\sum_{j=1}^r r_j = \Omega(\delta) > 0$, we have $\frac{d}{dt} (\sum_{j=1}^r r_j) \geq C \cdot \sum_{j=1}^r r_j$ for some constant C , i.e. the sum increases exponentially fast. The sequence eventually leaves $\mathcal{B}(0, \delta)$ and never comes back. The other spurious critical points $Z^* \neq 0$ can be treated similarly, see the proof of Theorem 3.3 in Section 3.5.

Definition 3.16. *Assume $g : \mathcal{M} \rightarrow \mathcal{N}$ is a differentiable mapping. The derivative of g at $Z \in \mathcal{M}$, denoted as $Dg(Z)$, is a linear operator $T_Z \rightarrow T_{g(Z)}$. For any $\xi \in T_Z$, $Dg(Z)\xi = \frac{d(g \circ \gamma)}{dt}(0)$ with $\gamma(0) = Z$ and $\frac{d}{dt}\gamma(0) = \xi \in T_Z$. (Here, T_Z and $T_{g(Z)}$ are tangent spaces at Z and $g(Z)$ respectively.)*

Lemma 3.17. For stepsize $\alpha > 0$, let $Z_+ := \phi_\alpha(Z) = \mathcal{R}(Z - \alpha P_{T_Z}(Z - X))$. Let $Z = V_z \Sigma_z V_z^H$, $Z_+ = V_{z,+} \Sigma_{z,+} V_{z,+}^H$, and $X = UDU^H$ be the eigenvalue decompositions of the SPSD (Symmetric positive semidefinite) matrices Z , Z_+ and X respectively. Denote $\phi_{\alpha,i,\Sigma_z,V_z(\cdot,\cdot)} : \mathbb{V}_1(\mathbb{R}^n) \rightarrow \mathbb{V}_1(\mathbb{R}^n)$ as the iteration function of the column space vector, which maps $V_z(\cdot,i)$ to $V_{z,+}(\cdot,i)$. Then, for any Σ_z and $V_z(\cdot,i)$, we have:

$$\langle D(\phi_{\alpha,i,\Sigma_z,V_z(\cdot,\cdot)} - I)|_{V_z(V)}, V \rangle = \frac{\alpha}{\sigma_i} V^H \tilde{V}_z Q D Q^H \tilde{V}_z^H V + \mathcal{O}(\alpha),$$

for all $V \in T_{\mathbb{V}_1(\mathbb{R}^n)}(V_z(\cdot,i))$ and $V \perp V_z(\cdot,i)$.

Example 3.18. To illustrate how the sequence escapes strict saddle points, we look at a toy example in the Euclidean space. For a quadratic function $f(Z)$ to be determined later, let $\phi(Z) = Z - \alpha \nabla f(Z)$, with Z^* being a fixed point with negative curvature such that $D(\phi - I)|_{Z^*} = -\alpha \nabla^2 f(Z)$. For any Z in a small neighborhood of Z^* , we have the following

$$\begin{aligned} \phi(Z) &= \phi(Z^*) + Z - Z^* + D(\phi - I)|_{Z^*}(Z - Z^*) \\ &= Z + D(\phi - I)|_{Z^*}(Z - Z^*). \end{aligned}$$

Here, we assume that

$$\nabla^2 f(Z) = P \begin{pmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} P^H,$$

with $P \in O(n)$. Denote $P = (p_1 \ p_2 \ p_3)$, assume $Z - Z^* = \eta_1 p_1 + \eta_2 p_2 + \eta_3 p_3$ and $\phi(Z) - Z^* = \eta_{1,+} p_1 + \eta_{2,+} p_2 + \eta_{3,+} p_3$. Now we consider the following

$$\begin{aligned} \langle \phi(Z) - Z^*, p_1 \rangle &= \langle Z - Z^*, p_1 \rangle - \alpha \langle P \begin{pmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} P^H (Z - Z^*), p_1 \rangle \\ &= \eta_1 - 2\alpha \eta_1. \end{aligned}$$

This implies $\eta_{1,+} = (1 - 2\alpha)\eta_1$. Similarly, we have $\eta_{2,+} = (1 + \alpha)\eta_2$, and $\eta_{3,+} = \eta_3$. In other words, η_2 grows exponentially. If we assume that Z satisfies $\eta_2 \geq \rho_0$, then it takes $\mathcal{O}(\log \frac{1}{\rho_0})$ iterations to escape a constant neighborhood of Z^* (that is, $\|Z - Z^*\| \geq \Omega(1)$).

Lemma 3.19. Denote $g : S \rightarrow S$. Let $U \in S$ be drawn from the uniform distribution on S with probability measure denoted as μ . If for all $z \in U \subset S$, $|\det Dg| \geq c_0$, where $c_0 > 0$ is an absolute constant, then we have

$$\mu(g^{-1}U) \leq \frac{1}{c_0} \cdot \mu(U).$$

Remark. Here, the uniform distribution can be extended to the general random distribution which allows $\max_{x,y \in S} \frac{\mu(y)}{\mu(x)} \leq C$ where $C \geq 1$ is an absolute constant. In this case, $\frac{1}{c_0}$ only involves an additional multiplicative factor.

Example 3.20. Here, we use a toy example to show how Lemma 3.19 works. Let $g(x) = \frac{1}{x}$ for all $x \in [\frac{1}{10}, 10]$, then $g(x) \in [\frac{1}{10}, 10]$. Simple computation gives $|Dg^{-1}(x)| = x^2$. Let x be a random variable drawn from the uniform distribution over $[\frac{1}{10}, 10]$. Denote the probability measure of x as μ . Now we consider the following cases:

- 1) First, consider $g^{-1}([\frac{1}{10}, 1])$. Direct computation gives $\mu(g^{-1}([\frac{1}{10}, 1])) = \mu([1, 10]) = \frac{10}{11}$. On the other hand, by Lemma 3.19, we have $|Dg^{-1}| \in [1, 100]$, which gives the bound $\frac{1}{11} \leq \mu(g^{-1}([\frac{1}{10}, 1])) \leq \frac{100}{11}$.
- 2) Similarly, consider $g^{-1}([1, 10])$. Direct computation gives $\mu(g^{-1}([1, 10])) = \mu([\frac{1}{10}, 1]) = \frac{1}{11}$. On the other hand, by Lemma 3.19, we have $|Dg^{-1}| \in [\frac{1}{100}, 1]$ which gives the upper bound $\frac{1}{110} \leq \mu(g^{-1}([1, 10])) \leq \frac{10}{11}$.

Lemma 3.21. Assume $f, g : \mathcal{M} \rightarrow \mathcal{M}$ are two mappings, f is differentiable at u , and g is differentiable at $f(u)$, then $g \circ f : \mathcal{M} \rightarrow \mathcal{M}$ is also differentiable, and $|\det D(g \circ f)(u)| = |\det D(g)(f(u))| \cdot |\det D(f)(u)|$.

Proof. Using Chain rule, for $u \in \mathcal{M}$, we have that the Jacobians satisfy

$$J_{g \circ f}(u) = J_g(v)J_f(u).$$

So, we have $|\det D(g \circ f)(u)| = |\det D(g)(f(u))| \cdot |\det D(f)(u)|$. □

Lemma 3.22 (Local convergence). Let $Z = U_z \Sigma V_z^H$ and $X = U D V^H$. Denote σ_j and d_j as the j -th largest singular value of Z and X respectively. Then

- 1) We have

$$\frac{\sigma_r^2}{\sigma_r^2 + \|X\|_F^2} \|Z - X\|_F^2 \leq \|P_{T_Z}(Z - X)\|_F^2 \leq \|Z - X\|_F^2;$$

- 2) In the local region around the ground truth $\{Z : \text{rank}(Z) = r, \|Z - X\|_F < \frac{d_r}{2}\}$, we have the lower bound of σ_r as $\sigma_r > \frac{d_r}{2} > 0$. Further, $\frac{\|P_{T_Z}(Z - X)\|_F}{\|Z - X\|_F} \gtrsim d_r$ and $\|\phi_\alpha(Z) - X\|_F \leq e^{-c} \|Z - X\|_F$ with $c = -\log(1 - \Omega(\alpha \sigma_r(X)^2))$.

3.5 Proofs of Theorem 3.1-Theorem 3.3

In this subsection, we give the detailed proofs of Theorems 3.1-3.3 using the technical results from the previous three subsections.

Proof of Theorem 3.1. Our goal is to show that with only a small probability, the sequence starting from Z_0 will reach $\mathcal{B}(Z^*, \delta)$ within K_0 iterations for $Z^* \in \mathcal{S}_{Z^*} \setminus \{0\}$. For such $Z \in \mathcal{B}(Z^*, \delta)$, by Lemma 3.10, there exists $i \in [r]$ such that $|V_{z,i}^H U_i| > 1 - \Omega(\delta)$. For such i , denote $B_i := \{W : |W^H U_i| > 1 - \Omega(\delta)\}$. We will show that the probability of the sequence initialized at Z_0 to reach $\mathcal{B}(Z^*, \delta)$ within K_0 iterations is very small. For this purpose, we estimate the upper bound for

the probability measure $\mu\left(\bigcup_{k=1}^{K_0} \phi_{\alpha,i}^{-k}(B_i)\right) := \mu\left(\bigcup_{\Sigma_z, V_z(\cdot, \setminus i)} \bigcup_{k=1}^{K_0} \phi_{\alpha,i, \Sigma_z, V_z(\cdot, \setminus i)}^{-k}(B_i)\right)$. More specifically, we have

$$\begin{aligned} \mu\left(\bigcup_{k=1}^{K_0} \phi_{\alpha,i}^{-k}(B_i)\right) &\leq \sum_{k=1}^{K_0} \max_{\Sigma_z, V_z(\cdot, \setminus i)} \mu\left(\phi_{\alpha,i, \Sigma_z, V_z(\cdot, \setminus i)}^{-k}(B_i)\right) \\ &\leq \sum_{k=1}^{K_0} \mu(B_i) \cdot \max_{\Sigma_z, V_z(\cdot, \setminus i)} \left| \det D \phi_{\alpha,i, \Sigma_z, V_z(\cdot, \setminus i)}^{-k} \right| \\ &\leq \sum_{k=1}^{K_0} C^k \cdot \mu(B_i) \\ &\lesssim \frac{1}{\text{poly}(n)}. \end{aligned}$$

The first inequality is by integrating the conditional probability measure with respect to the parameters $\Sigma_z, V_z(\cdot, \setminus i)$ and using the fact that the integral is bounded by the maximal value of the integrand, i.e. using $P(A) = \int P(A|y)p(y)dy \leq \max_y P(A|y)$. The second inequality is from Lemma 3.19 and due to the fact that Z_0 is drawn from the general random distribution. By choosing the stepsize $\alpha > 0$ small enough, we have $|\det D \phi^{-1}| \leq \frac{1}{1-\Theta(\alpha)} \leq C$ for a constant $C > 0$. Thus the third inequality can be achieved by applying Lemma 3.21 and Lemma 3.17. Since $1 - \Omega(\delta) \gtrsim \Omega(\frac{\log n}{n})$, the fourth inequality follows from Lemma 3.13 (3). Here, $\sum_{k=1}^{K_0} C^k$ can be treated as a constant since $K_0 = \Omega(1)$. \square

Proof of Theorem 3.2. For $Z^* = 0$, and $i \in [r]$, denote $H_i := \{W : Z \in \mathcal{B}(Z^*, \delta), W = V_z(\cdot, i) \text{ where } Z = V_z \Sigma_z V_z^H\}$. Consider the upper bound for $\mu(\phi_{\alpha,i}^{-k}(H_i))$:

$$\begin{aligned} \mu(\bigcup_{k=1}^{K_0} \phi_{\alpha,i}^{-k+C_1 \log n}(H_i)) &\leq \sum_{k=1}^{K_0} \max_{\Sigma_z, V_z(\cdot, \setminus i)} \mu\left(\phi_{\alpha,i, \Sigma_z, V_z(\cdot, \setminus i)}^{-k+C_1 \log n}(H_i)\right) \\ &\leq \sum_{k=1}^{K_0} \mu(H_i) \cdot \max_{\Sigma_z, V_z(\cdot, \setminus i)} \left| \det D \phi_{\alpha,i, \Sigma_z, V_z(\cdot, \setminus i)}^{-k+C_1 \log n} \right| \\ &\leq \sum_{k=1}^{K_0} K_0 \cdot \mu(H_i) \cdot C_3^k \cdot (1 - \Omega(1)\alpha)^{C_1 \log n} \\ &\lesssim \frac{1}{\text{poly}(n)}. \end{aligned}$$

The first inequality follows from $P(A) = \int P(A|y)p(y)dy \leq \max_y P(A|y)$. The second inequality is by Lemma 3.19 and due to the fact that Z_0 is drawn from the general random distribution. The third inequality is by Lemma 3.21 and Lemma 3.17. Here, due to Lemma 3.17, for $1 \leq k \leq K_0$, by requiring $\alpha = \Omega(1) > 0$ to be a small constant, at each iteration k , we have $|\det D \phi_{\alpha,i, \Sigma_k, V_k(\cdot, \setminus i)}^{-1}| \leq C_3 = \Omega(1)$. So we have $\sum_{k=1}^{K_0} C_3^k \leq C_5$ where $C_5 > 0$ is an absolute constant. At step $k \geq K_0$, by Lemma 3.17, by requiring $\delta > 0$ to be a small constant, we have $|\det D \phi_{\alpha,i}^{-1}| \leq 1 - C_4 \alpha$ with $C_4 = \Omega(1) > 0$. This holds for all $Z \in \mathcal{B}(0, \delta)$, so

$\max_{\Sigma_z, V_z(\cdot, \cdot)} |\det D\phi_{\alpha, i}^{-C_1 \log n}| \leq (1 - C_4\alpha)^{C_1 \log n}$. By requiring $C_1 = \Omega(1) \gtrsim -1/\log(1 - C_4\alpha) > 0$ large enough, we have the fourth inequality holds. Therefore, with the fail probability controlled by $\frac{1}{\text{poly}(n)}$, we have $Z_{K_1} \cap \mathcal{B}(0, \delta) = \emptyset$. Once outside $\mathcal{B}(0, \delta)$, we can use (7) in Lemma 3.15 to prove $\cup_{k=K_1}^{\infty} Z_k \cap \mathcal{B}(0, \delta) = \emptyset$, see the remark after Lemma 3.15. In conclusion, under the setting of this theorem, with high probability exceeding $1 - \frac{1}{\text{poly}(n)}$, the sequence will escape the spurious region $\mathcal{B}(0, \delta)$ after $\mathcal{O}(\log n)$ iterations. \square

Proof of Theorem 3.3. For any $Z^* = V_* \Sigma_* V_*^H \in \mathcal{S}_{Z^*}$, by Lemma 3.7, there exists $i \in [r]$ such that $\langle Z^*, U_i U_i^H \rangle = 0$. Denote $H_i := \{W : Z \in \mathcal{B}(Z^*, \delta), W = V_z(\cdot, i) \text{ where } Z = V_z \Sigma_z V_z^H\}$. Consider the upper bound for $\mu(\phi_{\alpha, i}^{-k}(H_i))$:

$$\begin{aligned} \mu(\cup_{k=1}^{K_1} \phi_{\alpha, i}^{-k+C_2 \log n}(H_i)) &\leq \sum_{k=1}^{K_1} \max_{\Sigma_z, V_z(\cdot, \cdot)} \mu\left(\phi_{\alpha, i, \Sigma_z, V_z(\cdot, \cdot)}^{-k+C_2 \log n}(H_i)\right) \\ &\leq \sum_{k=1}^{K_1} \mu(H_i) \cdot \max_{\Sigma_z, V_z(\cdot, \cdot)} \left| \det D\phi_{\alpha, i, \Sigma_z, V_z(\cdot, \cdot)}^{-k+C_2 \log n} \right| \\ &\leq K_1 \cdot \mu(H_i) \cdot C_1 \text{poly}(n) \cdot (1 - \Omega(1)\alpha)^{C_2 \log n} \\ &\lesssim \frac{1}{\text{poly}(n)}. \end{aligned}$$

The first inequality follows from $P(A) = \int P(A|y)p(y)dy \leq \max_y P(A|y)$. The second inequality is by Lemma 3.19 and due to the fact that Z_0 is drawn from the general random distribution. The third inequality is by Lemma 3.21 and Lemma 3.17. Here, due to Lemma 3.17, for $1 \leq k \leq K_1$, by requiring $\alpha = \Omega(1) > 0$ being a small constant, at each iteration k , we have $|\det D\phi_{\alpha, i, \Sigma_k, V_k(\cdot, \cdot)}^{-1}| \leq C_3 = \Omega(1)$. So we have $C_3^{K_1} \leq C_5 \text{poly}(n)$ where $C_5 > 0$ is an absolute constant. At step $k \geq K_1$, by Lemma 3.17, by requiring $\delta > 0$ to be a small constant, we have $|\det D\phi_{\alpha, i}^{-1}| \leq 1 - C_4\alpha$ with $C_4 = \Omega(1) > 0$. This holds for all $Z \in \mathcal{B}(Z^*, \delta)$, so $\max_{\Sigma_z, V_z(\cdot, \cdot)} |\det D\phi_{\alpha, i}^{-C_2 \log n}| \leq (1 - C_4\alpha)^{C_2 \log n}$. By requiring $C_2 = \Omega(1) > 0$ large enough, we have the fourth inequality holds. Therefore, with the fail probability controlled by $\frac{1}{\text{poly}(n)}$, we have $Z_{K_2} \cap \mathcal{B}(Z^*, \delta) = \emptyset$.

Once outside a spurious region $\mathcal{B}(Z^*, \delta)$, similar to the proof of Theorem 3.2, we can use (7) in Lemma 3.15 to prove $\cup_{k=K_2}^{\infty} Z_k \cap \mathcal{B}(Z^*, \delta) = \emptyset$. More specifically, we observe that the property $|\det D\phi_{\alpha, i}^{-1}| \leq 1 - \Omega(\alpha)$ holds in the region $\mathcal{B}(Z^*, \delta) \cap D_{Z^*}$ where $D_{Z^*} := \{Z : V_z(\cdot, i)^H U_i \leq 1 - \Omega(\delta)\}$. According to the first part of the proof, it can be shown that after $\mathcal{O}(\log n)$ iterations, the sequence is outside the boundary of $\mathcal{B}(Z^*, \delta) \cap D_{Z^*}$ with high probability, by requiring δ and α small enough. By Lemma C.1 and Lemma 3.15, we can prove that $\sum_{j=1}^r r_j$ grows exponentially fast. We can further prove that the sequence will never come back to $\mathcal{B}(\mathcal{S}_{Z^*}, \delta)$, using the monotone increase of $\sum_{j=1}^r r_j$. \square

4 Proofs of Theorem 2.2 and Theorem 2.3

In this section we will prove the Main result 2 (Theorem 2.2) and Main result 3 (Theorem 2.3) mentioned in Section 2. Theorem 2.2 is a special case of Theorem 2.1 restricted to $r = 1$, and its

proof is much simpler. Theorem 2.3 builds upon the previous theorem, but extends the analysis to the case of weak isometry. We provide some insights on the connections between Theorem 2.2 and Theorem 2.3 in Section 4.4.

4.1 Convergence tool for weak isometry

We first introduce Theorem 4.1, a variant of Theorem 3.4, as a fundamental tool for analyzing the convergence rate for functions with weak isometry as in Theorem 2.3. Using this theorem, we can show that if the measurement sampling operator preserves the distances of points on the manifold to the ground truth X to some extent (indicated by C_1 and C_2 in the distance-preserving condition below), and the projection operator satisfies a similar property as before, then with this $T(\cdot)$ operator, the loss function $f(Z) = \frac{1}{2}\|T(Z) - T(X)\|_F^2$ still preserves the nice properties of the original least squares loss function $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$. As a result, the sequences generated by the PGD still converge to the ground truth in a linear rate on the manifold as long as they stay outside of the spurious regions.

Theorem 4.1. *Assume $T : \mathcal{M}_r \rightarrow \mathbb{R}^m$ is a linear operator, and $f(Z) = \frac{1}{2}\|T(Z) - T(X)\|_2^2$. If the following conditions hold:*

- 1) (Distance-preserving condition) $C_1\|Z_k - X\|_F \leq \|T(Z_k) - T(X)\|_2 \leq C_2\|Z_k - X\|_F$, where $C_1, C_2 > 0$ are uniform constants for all k ;
- 2) (Critical ratio condition) $\|Z_k - X\|_F \leq C_3\|P_{T_{Z_k}}(\nabla f(Z_k))\|_F$, $C_3 > 0$ is a positive constant for all k .

Then, Conditions (D) and (L) hold with $\omega = \frac{1}{2}$. As a consequence, by Theorem 3.4, there exists a small enough $\alpha > 0$ such that the sequence $\{Z_k\}$ generated by the PGD converges to X in a linear rate: $\|Z_k - X\|_F \leq e^{-ck}$, with $c = -\log\left(1 - \Omega\left(\frac{\alpha}{C_2^2 C_3^2}\right)\right)$.

By throwing away a controllable failure probability, many random sensing applications potentially have such distance-preserving property. Some examples are mentioned in Section 1. The Restricted Isometry Property (RIP) can also be seen as a special case of this distance-preserving condition. Instead of checking the descent inequality and the Łojasiewicz inequality in Theorem 3.4, we use the above conditions as a more user-friendly version for such distance-preserving cases.

We are now ready to prove Theorems 2.2 and 2.3, using Theorem 4.1 and following a similar but simpler strategy as the one in Section 3.

4.2 Proof of Theorem 2.2

Proof. Denote $Z_+ = \phi_\alpha(Z)$ and $Z_+ = z_+ z_+^H$. Recall that $Z = z z^H$ and $X = x x^H$, $h = \|Z\|_F$ and $\rho = \frac{\langle X, Z \rangle}{\|X\|_F \|Z\|_F}$. Let $u_z = \frac{z}{\|z\|}$ and $u_{z,+} = \frac{z_+}{\|z_+\|}$. Then $\sqrt{\rho} = \langle u_z, x \rangle$. By Lemma C.3, we have

$$\begin{aligned} \frac{d}{dt} h &= u_z^H (X - Z) u_z = \rho - h, \\ \frac{d}{dt} u_z &= (I - u_z u_z^H) (X - Z) u_z \cdot \frac{1}{h} = \frac{1}{h} \cdot (\sqrt{\rho} x - \rho u_z), \\ \frac{d}{dt} \rho &= \frac{2}{h} \cdot \rho(1 - \rho). \end{aligned}$$

By **Assumption 1**, we obtain after integrating the above ODE system in time for one timestep that

$$\begin{aligned} h_{k+1} &= h_k + \alpha \cdot (-h_k + \rho_k) + \mathcal{O}(\alpha^2), \\ \rho_{k+1} &= \rho_k + 2\alpha \frac{\rho_k}{h_k} (1 - \rho_k) + \mathcal{O}(\alpha^2). \end{aligned}$$

Denote $\psi_{\alpha, h} : u_z \rightarrow u_{z,+}$ as the iteration function for u_z . For any $v \in T_{\mathbb{V}_1}(\mathbb{R}^n)$ and $t > 0$, direct computation gives that

$$\begin{aligned} \psi_\alpha(u_z) &:= u_{z,+} = u_z + \alpha \cdot \frac{\sqrt{\rho} x - \rho u_z}{h} + \mathcal{O}(\alpha^2), \\ \psi_\alpha(u_z + tv) &:= (u_z + tv)_+ \\ &= u_z + tv + \alpha \cdot \frac{\sqrt{\rho} x - \rho u_z}{h} + \alpha t \frac{\langle v, x \rangle x - 2\sqrt{\rho} \langle v, x \rangle u_z - \langle u_z, x \rangle^2 v}{h} + \mathcal{O}(t^2) + \mathcal{O}(\alpha^2). \end{aligned}$$

By **Assumption 1** and Definition 3.16, we obtain by direct calculation that

$$\langle D\psi_\alpha(u_z)(v), v \rangle = 1 + \alpha \cdot \frac{\langle v, x \rangle^2}{h} - \alpha \cdot \frac{\rho}{h} + \mathcal{O}(\alpha^2).$$

By Theorem 4.1, outside the spurious region $\mathcal{B}(0, \delta)$ (Lemma 3.9), the sequence has linear convergence rate. Thus it either enters $\mathcal{B}(0, \delta)$ or $B(X, \delta)$ in $\Omega(\frac{1}{\alpha \delta^2} \log \frac{1}{\delta})$ iterations. If the sequence enters $B(X, \delta)$, the local region of X , by Lemma 3.22 we have $\|P_{T_Z}(\nabla F_1(Z))\|_F / \|Z - X\|_F = \Omega(1)$ provided that $0 < \delta < \frac{\sigma_r(X)}{2}$. By Theorem 4.1, it further needs $\Omega(\log \frac{1}{\epsilon})$ iterations to generate an ϵ -accurate solution.

On the other hand, if the sequence enters $\mathcal{B}(0, \delta)$, similarly to the proofs of Theorem 3.1 and Theorem 3.2, one can prove that with high probability exceeding $1 - \frac{1}{\text{poly}(n)} - e^{-\Omega(n)}$, there exists $\mathcal{T} = \mathcal{O}(\log n)$ such that the sequence initialized from the general random distribution stays in $\mathcal{B}(0, \delta)$ for at most $\mathcal{T} = \mathcal{O}(\log n)$ iterations. After it escapes from $B(0, \delta)$, the sequence converges towards X in a linear rate. Combining all the above, it takes $\mathcal{O}(\log \frac{1}{\epsilon} + \log n)$ iterations to generate an ϵ -accurate solution. To prove that the sequence will never come back to the spurious region $\mathcal{B}(0, \delta)$, we can use an argument similar to the proof of Theorem 3.3. \square

4.3 Proof of Theorem 2.3

Proof. Recall that $Z = zz^H$ and $X = xx^H$, $h = \|Z\|_F$ and $\rho = \frac{\langle X, Z \rangle}{\|X\|_F \|Z\|_F} \in [0, 1]$. If $\mathbb{F} = \mathbb{R}$, the population loss of (7) is

$$\begin{aligned}\mathbb{E}f(Z) &= \frac{3}{2}\|Z\|_F^2 + \frac{3}{2}\|X\|_F^2 - \|Z\|_F\|X\|_F - 2\langle Z, X \rangle \\ &= \frac{1}{2}(\|Z\|_F - \|X\|_F)^2 + \|Z - X\|_F^2.\end{aligned}$$

Since $0 \leq (\|Z\|_F - \|X\|_F)^2 \leq \|Z - X\|_F^2$, we have

$$\|Z - X\|_F^2 \leq F(Z) \leq \frac{3}{2}\|Z - X\|_F^2.$$

If $\mathbb{F} = \mathbb{C}$, the population loss of (7) is

$$\begin{aligned}\mathbb{E}f(Z) &= \|Z\|_F^2 + \|X\|_F^2 - \|Z\|_F\|X\|_F - \langle Z, X \rangle \\ &= \frac{1}{2}(\|Z\|_F - \|X\|_F)^2 + \frac{1}{2}\|Z - X\|_F^2.\end{aligned}$$

And $\frac{1}{2}\|Z - X\|_F^2 \leq F_2(Z) \leq \|Z - X\|_F^2$.

We still denote $Z_+ = \phi_\alpha(Z)$. Assume $Z = zz^H$, and $Z_+ = z_+z_+^H$. Let $u_z = \frac{z}{\|z\|}$ and $u_{z,+} = \frac{z_+}{\|z_+\|}$. By Lemma C.3, we have

$$\begin{aligned}\frac{d}{dt}h &= u_z^H \left(2X - \left(\theta + 2 - \frac{\theta}{h} \right) Z \right) u_z = 2\rho - (2 + \theta)h + \theta, \\ \frac{d}{dt}u_z &= (I - u_z u_z^H) \left(2X + \left(\theta + 2 - \frac{\theta}{h} \right) Z \right) u_z \cdot \frac{1}{h} = \frac{2}{h} \cdot (\sqrt{\rho}x - \rho u_z), \\ \frac{d}{dt}\rho &= \frac{4}{h} \cdot \rho(1 - \rho).\end{aligned}$$

By **Assumption 1**, we have

$$\begin{aligned}h_{k+1} &= h_k + \alpha \cdot (\theta - (2 + \theta)h_k + 2\rho_k) + \mathcal{O}(\alpha^2), \\ \rho_{k+1} &= \rho_k + 4\alpha \frac{\rho_k}{h_k} (1 - \rho_k) + \mathcal{O}(\alpha^2).\end{aligned}$$

We still denote $\psi_{\alpha, h} : u_z \rightarrow u_{z,+}$ as the iteration function for u_z . Then for any $v \in T_{\mathbb{V}_1(\mathbb{R}^n)}$ and $t > 0$, direct computation gives

$$\begin{aligned}\psi_\alpha(u_z) &= u_z + 2\alpha \cdot \frac{\sqrt{\rho}x - \rho u_z}{h} + \mathcal{O}(\alpha^2), \\ \psi_\alpha(u_z + tv) &= u_z + tv + 2\alpha \cdot \frac{\sqrt{\rho}x - \rho u_z}{h} + 2\alpha t \frac{\langle v, x \rangle x - 2\sqrt{\rho} \langle v, x \rangle u_z}{h} + \mathcal{O}(\alpha^2).\end{aligned}$$

By **Assumption 1** and Definition 3.16, we obtain by direct computation that

$$\langle D\psi_\alpha(u_z)(v), v \rangle = 1 + \alpha \cdot \frac{2\langle v, x \rangle^2}{h} - \alpha \cdot \frac{2\rho}{h} + \mathcal{O}(\alpha^2).$$

Note that in this case, in addition to the spurious region $\mathcal{B}(0, \delta)$, there is another region $\mathcal{B}(Z^*, \delta) := \{Z : \rho \lesssim \delta, |\theta + 2 - \frac{\theta}{h}| \lesssim \delta\}$ where the Riemannian gradient is δ -small, as the Riemannian gradient is now

$$\nabla_{\mathcal{M}} F_2(Z) = \left(\theta + 2 - \frac{\theta}{h} \right) Z - 2P_{T_Z}(X).$$

In fact, $\|\nabla_{\mathcal{M}} F_2(Z)\| \leq \delta \iff Z \in \mathcal{B}(0, \delta) \cup \mathcal{B}(Z^*, \delta)$.

By Theorem 4.1, outside of $\mathcal{B}(0, \delta) \cup \mathcal{B}(Z^*, \delta)$, the randomly initialized sequence either enters $\mathcal{B}(0, \delta) \cup \mathcal{B}(Z^*, \delta)$ or $B(X, \delta)$ within $\Omega(\frac{1}{\alpha \delta^2} \log \frac{1}{\delta})$ iterations. If it enters local region of X as $B(X, \delta)$, by Lemma 3.22 we have $\|P_{T_Z}(\nabla F_2(Z))\|_F / \|Z - X\|_F = \Omega(1)$ provided that $0 < \delta < \frac{\sigma_r(X)}{2}$. In this case, by Theorem 4.1, it takes no more than $\Omega(\log \frac{1}{\epsilon})$ iterations to generate an ϵ -accurate solution.

On the other hand, if the sequence enters $\mathcal{B}(0, \delta) \cup \mathcal{B}(Z^*, \delta)$, then we can use an argument similar to the proofs of Theorem 3.2 and Theorem 3.3 to show that with high probability exceeding $1 - \frac{1}{\text{poly}(n)} - e^{-\Omega(n)}$, there exists $\mathcal{T} = \mathcal{O}(\log n)$ such that the sequence initialized from the general random distribution stays in $\mathcal{B}(0, \delta)$ or $\mathcal{B}(Z^*, \delta)$ for at most $\mathcal{O}(\log n)$ iterations. Afterwards, the sequence escapes from $B(0, \delta)$ and converges towards X in a linear rate. Note that the sequence will never come back to $\mathcal{B}(0, \delta) \cup \mathcal{B}(Z^*, \delta)$ once it escapes from this spurious region, as ρ is non-decreasing if $\delta \lesssim \rho \leq 1$ and $\alpha > 0$ is a small enough constant. Combining all the above, to generate an ϵ -accurate solution, i.e. $\|Z_k - X\|_F \leq \epsilon \|X\|_F$, it needs $\mathcal{O}(\log \frac{1}{\epsilon} + \log n)$ iterations. \square

4.4 Comparison of Theorem 2.2 and Theorem 2.3

Dynamics. The dynamical low-rank approximation (Lemma C.3) shows that the evolution of the column space is given by

$$\dot{U}_Z = P_{U_Z}^\perp (-\nabla F(Z)) U_Z S^{-1}.$$

For $F_1(Z) = \frac{1}{2} \|Z - X\|_F^2$, we have that $\nabla F_1(Z) = Z - X$, while for $F_2(Z) = \frac{\theta}{2} (\|Z\|_F - \|X\|_F)^2 + \|Z - X\|_F^2$ we have $\nabla F_2(Z) = (2 + \theta - \frac{\theta}{h})Z - 2X$. Although $F_2(Z)$ only satisfies the weak isometry property, direct computation shows that \dot{U}_Z for $F_1(Z)$ and $F_2(Z)$ are similar, because $P_{U_Z}^\perp$ on the right cancels out the Z terms and leaves only the X terms. That explains why the dynamics of $F_1(Z)$ is similar to that of $F_2(Z)$.

Stationary points. Theorem 2.2 is a special case of Theorem 2.1, therefore $Z^* = 0$ is the only spurious critical point and has a saddle-like property (see Section 3.3). On the other hand, for the phase retrieval problem in Theorem 2.3, it has two groups spurious critical points, which are $Z^* = 0$ and $\{Z^* : \|Z^*\|_F = \frac{\theta}{2+\theta} \|X\|_F, \langle Z, X \rangle = 0\}$. Still, the upper bound for the number of iterations that the sequence is trapped by the spurious region can be estimated in a similar way. This can be seen by comparing the proofs in Section 4.2 and Section 4.3.

Numerical illustration. To see the similarity between the evolution behavior in solving the rank-1 matrix recovery and the phase retrieval problem, we give some numerical experiments in Figure 5 and Figure 6 for a comparison. We can see that the curves of the evolution of h and ρ have similar shapes in both problems.




Figure 5: Solving the rank-1 matrix recovery by the randomly initialized PGD, with $n = 1024$, $\alpha = \frac{1}{3}$. Each band stands for the results from 100 independent experiments.

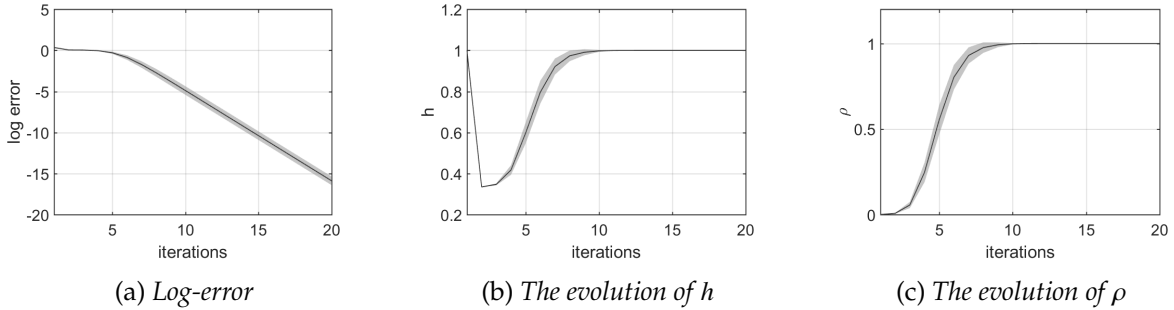


Figure 6: Solving the population phase retrieval problem by the randomly initialized PGD, with $n = 1024$, $\alpha = \frac{1}{3}$. Each band stands for the results from 100 independent experiments.

5 Conclusion and future work

In this paper, we have established a unified framework for the analysis of a class of low-rank matrix recovery problems. We have shown that using the Riemannian gradient descent (the PGD) algorithm on the low-rank matrix manifold, there is rigorous theoretical guarantee for the fast convergence rate in low-rank matrix recovery problems.

For this purpose, we first performed an extensive analysis of the low-rank matrix manifold \mathcal{M}_r itself by analyzing the simple least squares loss function $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$ where X is the ground truth. Our focus is on the symmetric positive semi-definite (SPSD) setting which is common in practice. Our results on the rank- r manifold with $r > 1$ are original and they are much more complicated than the corresponding results for the rank-1 case.

We showed that there is a ground truth and several spurious critical points on the manifold. The spurious critical points are of independent interest themselves, as they behave like strict saddle points, but their Hessians have singular eigen directions. We proved that the gradient descent or gradient flow starting from an initial guess drawn from the general random distribution converges to the ground truth with high probability. The initializations that might lead to the spurious critical points only have a small probability measure on the manifold. Improvement of this result to zero measure is left for future work.

The convergence rate towards the ground truth is nearly linear and is essentially independent of the dimensionality of the problem. The major difficulty when analyzing the convergence rate comes from estimating the upper bound for the number of iterations that the sequence is trapped by the spurious regions. Our primary tool is the iteration function of the column space derived from the dynamical low-rank approximation. We showed that with high probability, the initial angle between the column spaces of the ground truth and the random initialization point is $\tilde{\Omega}(\frac{1}{n})$, i.e. $\Omega(\frac{1}{n})$ with possible additional log-factors. The angle grows fast in spurious regions. Thus, we showed that with high probability, the sequence generated by the randomly initialized PGD escapes from the spurious regions quickly and enters a good region. When the sequence enters the good region, we then used the Łojasiewicz inequality tool to derive linear convergence.

The above analysis offers a general framework for a class of inverse problems that share a desirable structure, namely those problems whose forward problem is a linear mapping from a low-rank matrix to a vector and preserves the isometry property to some extent. The well-known RIP ensemble is a special case, but there are other applications with a weak isometry property. We analyzed the phase retrieval problem as an example of weak isometry problems, and established nearly optimal (linear) convergence rate. We focused on the population problem, i.e. the expectation of the loss function with respect to sampling. We invoked its connection with the rank-1 simple least squares problem, which can be described by scalar ODEs instead of matrix ODEs.

The global analysis for population loss functions can also be extended to finite-sample problems. The finite-sample loss function is the sum of its population loss plus some small deviations. One can control the magnitude of the deviations and show that the loss function still satisfies some weak isometry conditions. Thus the fundamental convergence guarantee by the Łojasiewicz inequality is readily applicable on the majority of the manifold. On the other hand, the geometry of the spurious regions, as well as the escape from these spurious regions, could be different from the population case. We leave the detailed analysis of finite-sample cases to our future work.

Acknowledgements. This research was in part supported by NSF Grants DMS P2259068 and DMS P2259075. We would also like to thank Prof. Jian-feng Cai for helpful suggestions.

References

- [1] Pierre-Antoine Absil, Robert Mahony, and Benjamin Andrews. Convergence of the iterates of descent methods for analytic cost functions. *SIAM Journal on Optimization*, 16(2):531–547, 2005.
- [2] Hedy Attouch and Jérôme Bolte. On the convergence of the proximal algorithm for non-smooth functions involving analytic features. *Mathematical Programming*, 116(1-2):5–16, 2009.
- [3] Hedy Attouch, Jérôme Bolte, Patrick Redont, and Antoine Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Łojasiewicz inequality. *Mathematics of Operations Research*, 35(2):438–457, 2010.
- [4] Hedy Attouch, Jérôme Bolte, and Benar Fux Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized gauss–seidel methods. *Mathematical Programming*, 137(1-2):91–129, 2013.
- [5] Yu Bai, Qijia Jiang, and Ju Sun. Subgradient descent learns orthogonal dictionaries. *arXiv preprint arXiv:1810.10702*, 2018.
- [6] Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Global optimality of local search for low rank matrix recovery. In *Advances in Neural Information Processing Systems*, pages 3873–3881, 2016.
- [7] Jérôme Bolte, Aris Daniilidis, Olivier Ley, and Laurent Mazet. Characterizations of Łojasiewicz inequalities: subgradient flows, talweg, convexity. *Transactions of the American Mathematical Society*, 362(6):3319–3363, 2010.
- [8] Samuel Burer and Renato DC Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- [9] Samuel Burer and Renato DC Monteiro. Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3):427–444, 2005.
- [10] Jian-Feng Cai and Ke Wei. Solving systems of phaseless equations via riemannian optimization with optimal sampling complexity. *arXiv preprint arXiv:1809.02773*, 2018.
- [11] Emmanuel J Candès and Yaniv Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359, 2011.
- [12] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.

- [13] Emmanuel J Candès, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [14] Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [15] Yuxin Chen and Emmanuel Candes. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Advances in Neural Information Processing Systems*, 28:739–747, 2015.
- [16] Yuxin Chen, Yuejie Chi, Jianqing Fan, and Cong Ma. Gradient descent with random initialization: fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 176(1-2):5–37, 2019.
- [17] Yuejie Chi, Yue M Lu, and Yuxin Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.
- [18] Christopher Criscitiello and Nicolas Boumal. Efficiently escaping saddle points on manifolds. In *Advances in Neural Information Processing Systems*, pages 5987–5997, 2019.
- [19] Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- [20] Luca Dieci and Timo Eirola. On smooth decompositions of matrices. *SIAM Journal on Matrix Analysis and Applications*, 20(3):800–819, 1999.
- [21] Simon S Du, Chi Jin, Jason D Lee, Michael I Jordan, Aarti Singh, and Barnabas Poczos. Gradient descent can take exponential time to escape saddle points. In *Advances in neural information processing systems*, pages 1067–1077, 2017.
- [22] Salar Fattahi and Somayeh Sojoudi. Exact guarantees on the absence of spurious local minima for non-negative robust principal component analysis. *arXiv preprint arXiv:1812.11466*, 2018.
- [23] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pages 797–842, 2015.
- [24] Rong Ge, Chi Jin, and Yi Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. *arXiv preprint arXiv:1704.00708*, 2017.
- [25] Rong Ge, Jason D Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.
- [26] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.

- [27] Thomas Y Hou, Zhenzhen Li, and Ziyun Zhang. Analysis of asymptotic escape of strict saddle sets in manifold optimization. *arXiv preprint arXiv:1911.12518*, 2019.
- [28] Jiang Hu, Xin Liu, Zai-Wen Wen, and Ya-Xiang Yuan. A brief introduction to manifold optimization. *Journal of the Operations Research Society of China*, 8(2):199–248, 2020.
- [29] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1724–1732. JMLR. org, 2017.
- [30] Othmar Koch and Christian Lubich. Dynamical low-rank approximation. *SIAM Journal on Matrix Analysis and Applications*, 29(2):434–454, 2007.
- [31] Stanislaw Łojasiewicz. Ensembles semi-analytiques, preprint 112 pp. *IHES notes*. Available at <http://perso.univ-rennes1.fr/michel.coste/Lojasiewicz.pdf>, 1965.
- [32] Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht. First-order methods almost always avoid saddle points. *arXiv preprint arXiv:1710.07406*, 2017.
- [33] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent converges to minimizers. *arXiv preprint arXiv:1602.04915*, 2016.
- [34] Zhenzhen Li, Jian-Feng Cai, and Ke Wei. Toward the optimal construction of a loss function without spurious local minima for solving quadratic equations. *IEEE Transactions on Information Theory*, 66(5):3242–3260, 2019.
- [35] Stanislaw Łojasiewicz. Une propriété topologique des sous-ensembles analytiques réels. *Les équations aux dérivées partielles*, 117:87–89, 1963.
- [36] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
- [37] Yurii Nesterov and Boris T Polyak. Cubic regularization of newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.
- [38] Donal B O’Shea and Leslie C Wilson. Limits of tangent spaces to real surfaces. *American journal of mathematics*, 126(5):951–980, 2004.
- [39] Ioannis Panageas and Georgios Piliouras. Gradient descent only converges to minimizers: Non-isolated critical points and invariant regions. *arXiv preprint arXiv:1605.00405*, 2016.
- [40] Benjamin Recht, Maryam Fazel, and Pablo A Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
- [41] Reinhold Schneider and André Uschmajew. Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality. *SIAM Journal on Optimization*, 25(1):622–646, 2015.

- [42] Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics*, 18(5):1131–1198, 2018.
- [43] Yue Sun, Nicolas Flammarion, and Maryam Fazel. Escaping from saddle points on riemannian manifolds. In *Advances in Neural Information Processing Systems*, pages 7276–7286, 2019.
- [44] Bart Vandereycken. Low-rank matrix completion by Riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, 2013.
- [45] Ke Wei, Jian-Feng Cai, Tony F Chan, and Shingyu Leung. Guarantees of Riemannian optimization for low rank matrix recovery. *SIAM Journal on Matrix Analysis and Applications*, 37(3):1198–1222, 2016.
- [46] Yangyang Xu and Wotao Yin. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences*, 6(3):1758–1789, 2013.
- [47] Richard Y Zhang, Somayeh Sojoudi, and Javad Lavaei. Sharp restricted isometry bounds for the inexistence of spurious local minima in nonconvex matrix recovery. *Journal of Machine Learning Research*, 20(114):1–34, 2019.
- [48] Ziyun Zhang. Exponential convergence of Sobolev gradient descent for a class of nonlinear eigenproblems. *arXiv preprint arXiv:1912.02135*, 2019.

A Manifold setting and optimization algorithm

A.1 Building blocks: low-rank matrix manifold

In this subsection we first present some basic structural properties of the low-rank matrix manifold as a preliminary. As mentioned in our previous work [27], the usual fixed-rank manifold may not be rigorous enough for two reasons stated below:

- 1) The fixed-rank manifold \mathcal{M}_r is not closed. Iterative optimization techniques can generate a sequence towards the ground truth, and closedness is naturally necessary for asymptotic convergence analysis.
- 2) It is possible that at some step $(Z_k - \alpha P_{T_{Z_k}}(\nabla f(Z_k)))$ happens to have rank lower than r and falls outside \mathcal{M}_r .

Therefore, for theoretical completeness we need the notion $\overline{\mathcal{M}}_r = \{Z \in \mathbb{F}^{n_1 \times n_2} : \text{rank}(Z) \leq r\}$ instead. The following are some basic definitions and essential properties on $\overline{\mathcal{M}}_r$.

Definition A.1 (Tangent Space). *Let $Z \in \mathcal{M}_r$, $Z = U\Sigma V^H$. Denote $\mathcal{U} = \text{Col}(U)$, $\mathcal{V} = \text{Col}(V)$ as the column spaces of U and V respectively. Then the tangent space of \mathcal{M}_r at Z is*

$$T_Z \mathcal{M}_r = (\mathcal{U} \otimes \mathcal{V}) \oplus (\mathcal{U} \otimes \mathcal{V}^\perp) \oplus (\mathcal{U}^\perp \otimes \mathcal{V}).$$

The projection onto the tangent space is

$$P_{T_Z} = P_{\mathcal{U}} \otimes I + I \otimes P_{\mathcal{V}} - P_{\mathcal{U}} \otimes P_{\mathcal{V}}.$$

The metric on $T\mathcal{M}$ is inherited from the metric of the embedded Euclidean space $\mathbb{F}^{n_1 \times n_2}$ which is equipped with inner product defined as $\langle A, B \rangle = \text{tr}(A^H B)$ and Frobenius norm $\|A\|_F = \sqrt{\text{tr}(A^H A)}$.

Definition A.2 (Tangent cone). *Let $Z \in \mathcal{M}_s \subset (\overline{\mathcal{M}}_r \setminus \mathcal{M}_r)$ where $s < r$, $Z = U\Sigma V^H$, $\mathcal{U} = \text{Col}(U)$, $\mathcal{V} = \text{Col}(V)$. Then the tangent cone of $\overline{\mathcal{M}}_r$ at Z is defined as follows*

$$T_Z \overline{\mathcal{M}}_r = T_Z \mathcal{M}_s \oplus \{\eta : \eta \in \mathcal{U}^\perp \otimes \mathcal{V}^\perp, \text{ and } \text{rank}(\eta) = r - s\}$$

The projection onto the tangent cone is given by

$$P_{T_Z \overline{\mathcal{M}}_r}(Y) = P_{T_Z \mathcal{M}_s}(Y) + Y_{r-s}$$

where Y_{r-s} is the best rank $(r - s)$ approximation of $Y - P_{T_Z \mathcal{M}_s}(Y)$ in the Frobenius norm.

We use P_{T_Z} for both $P_{T_Z \mathcal{M}_r}$ and $P_{T_Z \overline{\mathcal{M}}_r}$ when there is no confusion.

Remark. The low-rank matrix manifold as presented above is for general non-Hermitian case. The Hermitian case is similar except that the fixed rank manifold may have branches.

A.2 Optimization technique: the projected gradient descent (PGD)

In this section, we introduce the optimization technique on the low-rank matrix manifold, namely the projected gradient descent (PGD) with soft retraction onto the manifold. This Riemannian gradient descent technique has been studied in [10, 41, 44, 45]. For example, [45] and [10] use the Riemannian gradient descent to solve low-rank matrix recovery problems. These works also point out that the PGD enjoys light computational cost. In this paper, we will focus on the global analysis of such manifold optimization technique with random initialization.

Assume we are given a differentiable objective function $f(\cdot) : \mathcal{M} \rightarrow \mathbb{R}$ to be minimized, where \mathcal{M} can be a general Riemannian manifold. We start from a random initial guess $Z_0 \in \mathcal{M}$. Assume the sequence $\{Z_k\}_{k=0}^K$ is generated by the PGD (projected gradient descent):

$$Z_{k+1} = \mathcal{R} \left(Z_k - \alpha_k P_{T_{Z_k}} (\nabla f(Z_k)) \right). \quad (8)$$

Here $P_{T_{Z_k}}$ is the projection onto the tangent space (or tangent cone at a rank-deficient point) of \mathcal{M} at point Z_k , α_k is the k -th stepsize, and $\mathcal{R} : T_Z \rightarrow \mathcal{M}$ is a retraction defined in definition A.3. For the specific f in this paper, we have the following algorithm:

Projected Gradient Descent (PGD)

Input: Random measurements $\{A_j\}_{j=1}^m$, information $\{y_j\}_{j=1}^m$, random initialization Z_0 , maximal iteration number K_0 , tolerance ϵ .

Iteration: For $k = 1 : K_0$, do:

$$Z_{k+1} = \mathcal{R} \left(Z_k - P_{T_{Z_k}} \left(\frac{1}{m} \sum_{j=1}^m \langle Z_k - X, A_j \rangle A_j \right) \right), \text{ if } f(Z_k) < \epsilon, \text{ break.}$$

Output: Estimator $\hat{X} = Z_k$.

In the aforementioned algorithm, the retraction operation is necessary since it ensures the generated iteration points stay on the manifold \mathcal{M} . The retraction operator is defined as follows.

Definition A.3 (Retraction). Let T_Z be the tangent space (or tangent cone) of \mathcal{M} at Z . We call $\mathcal{R}_Z : T_Z \rightarrow \mathcal{M}$ a retraction, if for any $\xi \in T_Z$,

$$\lim_{\alpha \rightarrow 0^+} \frac{\|\mathcal{R}_Z(\alpha\xi) - (Z + \alpha\xi)\|_F}{\alpha} = 0. \quad (9)$$

For simplicity we may also write $\mathcal{R}_Z(\alpha\xi)$ as $\mathcal{R}(Z + \alpha\xi)$. We will refer to (9) as the first-order retraction property.

Remark. A natural retraction on \mathcal{M}_r is the following best rank- r approximation under the Frobenius norm:

$$\mathcal{R}_N(Z + \xi) = \arg \min_{Y \in \mathcal{M}_r} \|Z + \xi - Y\|_F.$$

We use \mathcal{R}_N as our \mathcal{R} in the rest of the paper. It not only satisfies the first-order retraction property, but is also second order, see also [38, 41, 44]:

$$\mathcal{R}(Z + \alpha\xi) = Z + \alpha\xi + \alpha^2\eta + \mathcal{O}(\alpha^3).$$

It is worth mentioning that there exists other manifold optimization techniques. For example, some manifold optimization methods skip the projection step $P_{T_{Z_k}}$ and compute the singular value decomposition directly. This is also called the ‘‘hard retraction’’. We choose the current projected gradient descent with ‘‘soft retraction’’ mainly due to two reasons:

- 1) Most of the operations we list are defined only for tangent bundles, e.g. first-order retraction (9) and the Riemannian Hessian. Generally, $\nabla f(Z_k)$ is not in T_{Z_k} , but $P_{T_{Z_k}}(\nabla f(Z_k))$ is.
- 2) The projected version of GD is also cheaper in terms of computation. Namely, solving $\mathcal{R}(Z_k - \alpha_k \nabla f(Z_k))$ involves calculating SVD of a $n_1 \times n_2$ matrix, while $\mathcal{R}(Z_k - \alpha_k P_{T_{Z_k}}(\nabla f(Z_k)))$ only involves that of a $2r \times 2r$ matrix, as mentioned in the previous literature [45]. Since $r \ll \min\{n_1, n_2\}$, the soft retraction is of lighter computational cost.

Light computational cost of the PGD. With the PGD, we have

$$\begin{aligned} Z_{k+1} &= \mathcal{R}(Z_k - \alpha P_{T_{Z_k}}(W_k)) \\ &= \mathcal{R}(U_k \Sigma_k V_k^H - \alpha (U_k U_k^H W_k + W_k V_k V_k^H - U_k U_k^H W_k V_k V_k^H)) \\ &= \mathcal{R}(U_k (\Sigma_k - \alpha U_k^H W_k V_k) V_k^H - \alpha U_k ((I - V_k V_k^H) W_k^H U_k)^H \\ &\quad - \alpha ((I - U_k U_k^H) W_k V_k) V_k^H) \\ &= \mathcal{R}\left(\begin{bmatrix} U_k & Q_{k,2} \end{bmatrix} \begin{bmatrix} \Sigma_k - \alpha U_k^H W_k V_k & -\alpha R_{k,1}^H \\ -\alpha R_{k,2} & 0 \end{bmatrix} \begin{bmatrix} V_{k,1}^H \\ Q_{k,1}^H \end{bmatrix} \right). \end{aligned}$$

Here $W_k = \nabla f(Z_k) \in \mathbb{F}^{n_1 \times n_2}$ and we assume $Z_k = U_k \Sigma_k V_k^H$. Assume $(I - V_k V_k^H) W_k^H U_k = Q_{k,1} R_{k,1}$ and $(I - U_k U_k^H) W_k V_k = Q_{k,2} R_{k,2}$ are the QR factorizations of the respective matrices. Notice that $Q_{k,2}^H U_k = \mathbf{0}_{r \times r}$, $Q_{k,1}^H V_k = \mathbf{0}_{r \times r}$. Therefore, to compute SVD of $Z_k - \alpha P_{T_{Z_k}}(W_k)$ it only involves solving the SVD of $\begin{bmatrix} \Sigma_k - \alpha U_k^T W_k V_k & -\alpha R_{k,1}^H \\ -\alpha R_{k,2} & 0 \end{bmatrix}$, which is only a $2r \times 2r$ matrix.

B Proofs

B.1 Proof of Theorem 3.4

Proof. Since $\{f_k\}_{k=1}^\infty$ is a monotone and lower bounded sequence, Condition (D) and continuity of $f(\cdot)$ implies convergence of $\{Z_k\}$ to some fixed point Z^* . Without loss of generality we assume that $f(Z^*) = 0$. By Conditions (D) and (L), we have

$$\begin{aligned} \|Z_{k+1} - Z_k\|_F &\leq \frac{1}{C_d \|P_{T_{Z_k}}(\nabla f(Z_k))\|_F} (f_k - f_{k+1}) \leq \frac{C_l}{C_d} f_k^{\theta-1} (f_k - f_{k+1}) \\ &\leq \frac{C_l}{C_d} \int_{f_{k+1}}^{f_k} \phi^{\theta-1} d\phi = \frac{C_l}{\theta C_d} (f_k^\theta - f_{k+1}^\theta). \end{aligned}$$

Since $\{f_k\}_{k=1}^\infty$ is a monotone and lower bounded sequence, $\{f_k\}$ is convergent. Therefore, $\{Z_k\}$ is convergent, and the limit point is Z^* .

Consider $s_k := \sum_{i=k}^\infty \|Z_{i+1} - Z_i\|_F \leq \frac{C_l}{\theta C_d} f_k^\theta$. Then we get

$$s_k^{\frac{1-\theta}{\theta}} \leq \left(\frac{C_l}{\theta C_d}\right)^{\frac{1-\theta}{\theta}} f_k^{1-\theta} \leq C_l \left(\frac{C_l}{\theta C_d}\right)^{\frac{1-\theta}{\theta}} \|P_{T_{Z_k}}(\nabla f(Z_k))\|_F.$$

Let $\xi_k = -P_{T_{Z_k}}(\nabla f(Z_k))$ and $\alpha \tilde{\xi}_k = Z_{k+1} - Z_k = \mathcal{R}(Z_k + \alpha \xi_k) - Z_k$. By the retraction property, $\alpha \tilde{\xi}_k = \alpha \xi_k + o(\alpha)$. So we have

$$s_k^{\frac{1-\theta}{\theta}} \leq C_l \left(\frac{C_l}{\theta C_d}\right)^{\frac{1-\theta}{\theta}} \|\xi_k\|_F = C_l \left(\frac{C_l}{\theta C_d}\right)^{\frac{1-\theta}{\theta}} \frac{1}{\alpha} (\|Z_{k+1} - Z_k\|_F + o(\alpha)).$$

Let $\rho = \rho(\alpha) := \frac{\alpha}{C_l} \left(\frac{\theta C_d}{C_l}\right)^{\frac{1-\theta}{\theta}}$, i.e. ρ is a constant depending only on C_l, C_d and θ . Then we have

$$\rho(\alpha) s_k^{\frac{1-\theta}{\theta}} = \|Z_{k+1} - Z_k\|_F + o(\alpha).$$

One can choose α small enough such that $0 < \rho(\alpha) < 1$ and

$$\rho(\alpha) s_k^{\frac{1-\theta}{\theta}} \leq s_k - s_{k+1}.$$

In the case of $\theta = \frac{1}{2}$, the above inequality gives $s_{k+1} \leq (1 - \rho)s_k$, which implies $\|Z_k - Z^*\|_F \leq s_k \leq (1 - \rho)^k s_0 = s_0 e^{-ck}$, where $c = -\log(1 - \rho) > 0$. This gives the linear convergence rate.

In the case of $0 < \theta < \frac{1}{2}$: Assume $s_k = c_1 k^{-p}$, then we have

$$\begin{aligned} s_{k+1} &= c_1 \frac{1}{(k+1)^p} = c_1 \frac{1}{k^p} \left(1 + \frac{1}{k}\right)^{-p} \\ &= c_1 \frac{1}{k^p} \left(1 - \frac{p}{k}\right) + O\left(\frac{1}{k^{p+2}}\right) \\ &= s_k \left(1 - p s_k^{\frac{1}{p}} c_1^{-\frac{1}{p}}\right) + O\left(\frac{1}{k^{p+2}}\right) \\ &\leq s_k \left(1 - \rho s_k^{\frac{1-2\theta}{\theta}}\right), \end{aligned}$$

where the last inequality follows from $s_{k+1} \leq s_k (1 - \rho s_k^{\frac{1-2\theta}{\theta}})$. Choose c_1 large enough, then the above inequality holds with $p = \frac{\theta}{1-2\theta}$. Thus, we obtain $\|Z_k - Z^*\|_F \leq s_k \leq c_1 k^{-\frac{\theta}{1-2\theta}}$. This implies the polynomial convergence rate. \square

B.2 Proof of Lemma 3.7

Proof. 1) It is obvious that Z is a fixed point of $Z_{k+1} = \mathcal{R}(Z_k - \alpha P_{T_{Z_k}}(\nabla f(Z_k)))$ if and only if $P_{T_Z}(\nabla F_1(Z)) = 0$. Denote $Z = U_z \Sigma_z V_z^H$ and $X = U D V^H$, then for any $\xi \in T_Z$, there exists $\Delta_1 \in \mathbb{F}^{n_1 \times r}$ and $\Delta_2 \in \mathbb{F}^{n_2 \times r}$, such that $\xi = U_z \Delta_1^H + \Delta_2 V_z^H$. Simple calculation gives

$$\begin{aligned} P_{T_Z}(\nabla F_1(Z)) = 0 &\iff \langle P_{T_Z}(Z - X), \xi \rangle = 0, \text{ for all } \xi \in T_Z \\ &\iff \langle Z - X, U_z \Delta_1^H + \Delta_2 V_z^H \rangle = 0, \text{ for all } \xi = U_z \Delta_1^H + \Delta_2 V_z^H \in T_Z \\ &\iff \text{tr}((V_z \Sigma_z - V D U^H U_z) \Delta_1^H + (\Sigma_z U_z^H - V_z^H V D U^H) \Delta_2), \text{ for all } \Delta_1 \in \mathbb{F}^{n_1 \times r}, \Delta_2 \in \mathbb{F}^{n_2 \times r} \\ &\iff V_z \Sigma_z - V D U^H U_z = \Sigma_z U_z^H - V_z^H V D U^H = 0 \\ &\iff U_z^H X = \Sigma_z V_z^H \text{ and } U_z \Sigma_z = X V_z. \end{aligned}$$

This implies $P_{U_z}(X) = P_{U_z}(Z) = Z$ and $P_{V_z}(Z^H) = Z^H = P_{V_z}(X^H)$. Assume

$$X = \begin{pmatrix} U_z & \widetilde{U}_z \end{pmatrix} \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{pmatrix} \begin{pmatrix} V_z^H \\ \widetilde{V}_z^H \end{pmatrix}.$$

Then we get $X_{11} = \Sigma_z$, $X_{12} = 0$ and $X_{21} = 0$. Therefore, we have

$$Z = U_1 D_1 V_1^H, \text{ with } U = (U_1 \ U_2), D = \text{diag} \{D_1, D_2\} \text{ and } V = (V_1 \ V_2).$$

- 2) If X has distinct singular values, then \mathcal{S} consists of the points $Z^* = \sum_{i=1}^r d_i \eta_i u_i v_i^H$, where $\eta \in \{0, 1\}^r$ and η is not $(1, 1, \dots, 1)^H \in \mathbb{R}^r$. So $|\mathcal{S}| = 2^r - 1$. □

B.3 Proof of Lemma 3.9

Proof. Assume $Z = U_z \Sigma_z V_z^H$, and let $\widetilde{U}_z, \widetilde{V}_z \in \mathbb{F}^{n \times (n-r)}$ be the orthogonal complements of U_z and V_z . We can express X in the following block form under this new basis:

$$X = \begin{pmatrix} U_z & \widetilde{U}_z \end{pmatrix} \widetilde{X} \begin{pmatrix} V_z^H \\ \widetilde{V}_z^H \end{pmatrix} = \begin{pmatrix} U_z & \widetilde{U}_z \end{pmatrix} \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{pmatrix} \begin{pmatrix} V_z^H \\ \widetilde{V}_z^H \end{pmatrix},$$

where

$$\widetilde{X} = \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{pmatrix} = Q_L \cdot D \cdot Q_R^H.$$

Then we have

$$P_{T_Z}(X) = \begin{pmatrix} U_z & \widetilde{U}_z \end{pmatrix} \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & 0 \end{pmatrix} \begin{pmatrix} V_z^H \\ \widetilde{V}_z^H \end{pmatrix}, \quad U = \begin{pmatrix} U_z & \widetilde{U}_z \end{pmatrix} Q_L, \quad V = \begin{pmatrix} V_z & \widetilde{V}_z \end{pmatrix} Q_R.$$

Assume that $\|P_{T_Z}(Z - X)\|_F \leq \delta$, then

$$\|P_{T_Z}(Z - X)\|_F = \left\| \begin{pmatrix} X_{11} - \Sigma_z & X_{12} \\ X_{21} & 0 \end{pmatrix} \right\|_F = \left\| \widetilde{X} - \begin{pmatrix} \Sigma_z & 0 \\ 0 & X_{22} \end{pmatrix} \right\|_F \leq \delta.$$

Let

$$\widetilde{S} := \begin{pmatrix} \Sigma_z & 0 \\ 0 & X_{22} \end{pmatrix} = \begin{pmatrix} I_r & 0 \\ 0 & P_L \end{pmatrix} \begin{pmatrix} \Sigma_z & 0 \\ 0 & \Sigma_{22} \end{pmatrix} \begin{pmatrix} I_r & 0 \\ 0 & P_R^H \end{pmatrix}$$

where the second equality gives the singular value decomposition of the matrix \widetilde{S} . Then $\|\widetilde{X} - \widetilde{S}\|_F \leq \delta$.

Using Lemma C.2, we have that the singular values of \widetilde{S} are δ -perturbations of those of \widetilde{X} . Note that the singular values of \widetilde{X} are the same as those of X , which are $\{d_1, \dots, d_r\} \cup \{0\}$. On the other hand, the singular values of \widetilde{S} are $\{\sigma_1, \dots, \sigma_r\} \cup \{\widetilde{\sigma}_1, \dots, \widetilde{\sigma}_{n-r}\}$, where $\{\sigma_i\}$ and $\{\widetilde{\sigma}_i\}$ are the diagonal entries of Σ_z and Σ_{22} respectively. Thus, for each σ_i , $i \in [r]$, either $|\sigma_i - d_j| = \mathcal{O}(\delta)$

for some $j \in [r]$, or $\sigma_i = \mathcal{O}(\delta)$. In other words, each σ_i either captures a singular value of X , or is close to zero.

Now let \mathcal{J} denote the set of indices of d_j 's captured by σ_i 's, and $\mathcal{J}^c = [r] \setminus \mathcal{J}$. Without loss of generality, assume $|\sigma_i - d_i| = \mathcal{O}(\delta)$ for $i \in \mathcal{J}$, i.e. their indices also match. Let $U_1 = U(:, \mathcal{J})$, $U_2 = U(:, \mathcal{J}^c)$, $V_1 = V(:, \mathcal{J})$, $V_2 = V(:, \mathcal{J}^c)$, and $D_1 = D(\mathcal{J}, \mathcal{J})$, $D_2 = D(\mathcal{J}^c, \mathcal{J}^c)$. Then $\|\Sigma_z(\mathcal{J}, \mathcal{J}) - D_1\| = \mathcal{O}(\delta)$, and $\|\Sigma_z(\mathcal{J}^c, \mathcal{J}^c)\| = \mathcal{O}(\delta)$.

By Lemma C.1, the singular subspaces of \tilde{X} and \tilde{S} corresponding to indices \mathcal{J} are δ -close. In mathematical terms, we have

$$\|P_{I(:, \mathcal{J})} - P_{Q_L(:, \mathcal{J})}\| = \mathcal{O}(\delta), \quad \|P_{I(:, \mathcal{J})} - P_{Q_R(:, \mathcal{J})}\| = \mathcal{O}(\delta),$$

where I denotes the identity matrix and P denotes the projection onto the space spanned by the column vectors. Using the relations $U = (U_z, \tilde{U}_z)Q_L$, and $V = (V_z, \tilde{V}_z)Q_R$, we have

$$\|P_{U_z(:, \mathcal{J})} - P_{U_1}\| = \mathcal{O}(\delta), \quad \|P_{V_z(:, \mathcal{J})} - P_{V_1}\| = \mathcal{O}(\delta),$$

Let $B = U_z(:, \mathcal{J})$ and $C = V_z(:, \mathcal{J})$ and we have the results regarding B and C in the lemma.

Similarly, the singular subspaces of Z corresponding to \mathcal{J}^c are δ -close to some singular subspaces perpendicular to U and V . Denote them as \tilde{U} and \tilde{V} respectively. Then we obtain

$$\|P_{U_z(:, \mathcal{J}^c)} - P_{U_1}\| = \mathcal{O}(\delta), \quad \|P_{V_z(:, \mathcal{J}^c)} - P_{V_1}\| = \mathcal{O}(\delta)$$

Let $\tilde{B} = U_z(:, \mathcal{J}^c)$ and $\tilde{C} = V_z(:, \mathcal{J}^c)$ and we have the full result. Note that by Lemma C.1, the constant in the $\mathcal{O}(\cdot)$ notation only depends on the gap between the two groups of singular values, which in our case is determined by the smallest singular value of X . \square

B.4 Proof of Lemma 3.10

Proof. Following the proof of Lemma 3.9, we have $U = U_z R + \tilde{U}_z Q = U_z(Q_1(:, [s]), Q_2(:, [r - s])) + \tilde{U}_z(Q_3(:, [s]), Q_4(:, [r - s]))$. Therefore, $R = (Q_1(:, [s]), Q_2(:, [r - s]))$, and we have

$$R = \begin{pmatrix} I_s + E_1 & E_2 \\ E_3 & E_4 \end{pmatrix},$$

with $\|E_1\|^{\frac{1}{2}}, \|E_3\|, \|E_2\|, \|E_4\| \leq \mathcal{O}(\delta)$. \square

B.5 Proof of Lemma 3.11

Proof. The proof is a direct consequence of Theorem 3.4 and Lemma 3.22. Specifically, to make use of Theorem 3.4, it suffices to check Condition (L) with $\theta = \frac{1}{2}$ and Condition (D).

As stated in the beginning of section 3.2, we assume $\{Z_k\}$ is bounded, i.e. $\|Z\|_F \leq C$. In the region $\{Z : \|P_{T_Z}(Z - X)\| \geq C_L\}$, we have $\frac{\|P_{T_Z}(Z - X)\|_F}{\|Z - X\|_F} \geq \frac{C_L}{C + \|X\|_F}$; on the other hand, in the region $\{Z : \|Z - X\|_F \leq \frac{\sigma_r(X)}{2}\}$, by Lemma 3.22, $\frac{\|P_{T_Z}(Z - X)\|_F}{\|Z - X\|_F} \geq \frac{\sigma_r(X)^2}{\sigma_r(X)^2 + 4\|X\|_F^2}$. In other words, in the region $\{Z : \|P_{T_Z}(Z - X)\| \geq C_L\} \cup \{Z : \|Z - X\|_F \leq \frac{\sigma_r(X)}{2}\}$, for $F_1(Z) = \frac{1}{2}\|Z - X\|_F^2$, we have condition (L) holds with $\theta = \frac{1}{2}$ and $C_l = \max\left\{\frac{C + \|X\|_F}{C_L}, 1 + \frac{4\|X\|_F^2}{\sigma_r(X)^2}\right\}$.

For condition (D), we consider

$$\begin{aligned}
f_k - f_{k+1} &= \frac{1}{2}\|Z_k - X\|_F^2 - \frac{1}{2}\|Z_{k+1} - X\|_F^2 \\
&= \frac{1}{2}\text{tr}((Z_k + Z_{k+1} - 2X)(Z_k - Z_{k+1})) \\
&= \text{tr}((X - Z_k)(Z_{k+1} - Z_k)) - \frac{1}{2}\|Z_{k+1} - Z_k\|_F^2.
\end{aligned}$$

Let $Z_{k+1} = Z_k + \alpha\tilde{\xi}_k$ and $\xi_n = -P_{T_{Z_k}}(\nabla f(Z_k))$, then by first-order retraction property we have $\tilde{\xi}_k = \xi_k + o(\|\xi_k\|_F)$. So the left- and right-hand side of Condition (D) are

$$\begin{aligned}
\text{LHS} &= f_k - f_{k+1} \\
&= \text{tr}(-\nabla f(Z_k)\alpha\tilde{\xi}_k) - \frac{\alpha^2}{2}\|\tilde{\xi}_k\|_F^2 \\
&= \text{tr}(-\nabla f(Z_k)(-\alpha P_{T_{Z_k}}(\nabla f(Z_k)) + o(\alpha\|\xi_k\|_F))) - \frac{\alpha^2}{2}\|\tilde{\xi}_k\|_F^2 \\
&= \alpha\|\xi_k\|_F^2 + o(\alpha\|\xi_k\|_F), \\
\text{RHS} &= C_d\|P_{T_{Z_k}}(\nabla f(Z_k))\|_F\|Z_{k+1} - Z_k\|_F \\
&= C_d\|\xi_k\|_F\|\alpha\tilde{\xi}_k\|_F \\
&= C_d\alpha\|\xi_k\|_F^2 + o(\alpha\|\xi_k\|_F^2).
\end{aligned}$$

By choosing a proper $C_d > 0$ and a small enough step size α , one can get

$$f_k - f_{k+1} \geq C_d\|P_{T_{z_k}}(\nabla f(Z_k))\|_F\|X_{k+1} - X_k\|_F,$$

which is Condition (D). The results now follow from Theorem 3.4. \square

B.6 Proof of Lemma 3.13

Proof. Since the marginal distribution of $W(:,i)$ ($i=1,2,\dots,r$) is the uniform distribution on S^{n-1} and the distribution of W is right-rotational invariant, we assume unitary $U \in \mathbb{F}^{n \times n}$ such that $Uu_0 = e_1$ and $UW = \tilde{W}$ and use the distribution of \tilde{W} to replace that of W . Now consider the marginal distribution of W_i , and write $W_i = \frac{g}{\|g\|}$ where g is drawn from $\mathcal{N}(0, I_n)$. Then, by the Bernstein-type inequality, we have the following estimation:

$$\text{Prob}\left(\left|\|g\|^2 - \mathbb{E}(\|g\|^2)\right| > t\mathbb{E}(\|g\|^2)\right) \leq 2 \exp\left\{-c \min\left\{\frac{t^2\mathbb{E}(\|g\|^2)^2}{K^2n}, \frac{t\mathbb{E}(\|g\|^2)}{K}\right\}\right\}.$$

This implies $\text{Prob}(\|g\|^2 \in (\frac{1}{2}n, \frac{3}{2}n)) \geq 1 - e^{-\Omega(n)}$. On the other hand, we have

$$\begin{aligned}
\text{Prob}\left(|g(1)|^2 - 1 > t\right) &= \text{Prob}(|g(1)| > \sqrt{1+t}) \\
&= 2 \int_{\sqrt{1+t}}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx =: \sqrt{\frac{2}{\pi}} A, \\
\text{where } A^2 &= \int_{\sqrt{1+t}}^{\infty} \int_{\sqrt{1+t}}^{\infty} e^{-\frac{x^2+y^2}{2}} dx dy = \frac{\pi}{2} e^{-\frac{1+t}{2}}.
\end{aligned}$$

Taking $t = -1 + c_1 \frac{1}{\log n}$, we have

$$\text{Prob}\left(|g(1)|^2 > c_1 \frac{1}{\log n}\right) = e^{-c_1/(4 \log n)} \geq 1 - \mathcal{O}\left(\frac{1}{\log n}\right).$$

Therefore, we have $|u_0^H W_i|^2 \gtrsim \frac{1}{n \log n}$ with the fail probability controlled by $e^{-\Omega(n)} + \mathcal{O}\left(\frac{1}{\log n}\right)$. By taking $t = c_1 \log n - 1$, with $c_1 = \Omega(1)$, we have:

$$\text{Prob}(|g(1)|^2 \geq c_1 \log n) \leq e^{-c_1 \log(n)/4} = 2 \frac{1}{n^{c_1/4}}.$$

Therefore, we have $|u_0^H W_i|^2 \lesssim \frac{\log n}{n}$ with fail probability controlled by $e^{-\Omega(n)} + \frac{1}{\text{poly}(n)}$. Taking $t = -1 + c_2/n^p$, we have

$$\frac{|g(1)|^2}{\|g\|^2} \gtrsim \frac{1}{n^{p+1}}$$

holds with fail probability controlled by $e^{-\Omega(n)} + \mathcal{O}\left(\frac{1}{n^p}\right)$. For the complex case, the proof is very similar and we omit the details here. \square

B.7 Proof of Lemma 3.15

Proof. 1) Denote $Z = U_z S_z U_z^H$ as an alternative decomposition of Z with $U_z^H U_z = I_r$ and $S_z \in \mathbb{R}^{r \times r}$. Assume $P \in \mathcal{O}(r)$ such that: $V_z = U_z P$ and $P^H S_z P = \Sigma_z$. Denote $U = U_z R_* + \widetilde{U}_z Q_* = V_z R + \widetilde{V}_z Q$. By applying the dynamical low-rank approximation in Lemma C.3, we have:

$$\begin{aligned} \frac{d}{dt} U_z &= \widetilde{U}_z Q_* D R_*^H S_z^{-1} \\ \frac{d}{dt} S_z &= R_* D R_*^H - S_z. \end{aligned}$$

Therefore, we obtain

$$\frac{d}{dt} R_* = \frac{d}{dt} (U_z^H U) = S_z^{-1} R_* D Q_*^H Q_*.$$

Using $R = V_z^H U = P^H U_z^H U = P^H R_*$, we have

$$\begin{aligned} \frac{d}{dt} R &= \frac{d}{dt} (P^H R_*) \\ &= \frac{d}{dt} (P)^H R_* + P^H \frac{d}{dt} (R_*) \\ &= \frac{d}{dt} (P)^H P R + P^H S_z^{-1} P R D Q_*^H Q_* \\ &= \left(\frac{d}{dt} (P)^H P \right) R + \Sigma_z^{-1} R D Q_*^H Q_*. \end{aligned}$$

The last equation follows from $P^H S_z^{-1} P = \Sigma_z^{-1}$ and $Q^H Q = I_r - R^H R = I_r - R^H P^H P R = Q_*^H Q_*$. Then, we have

$$\begin{aligned} \frac{d}{dt}(RDR^H) &= \Sigma_z^{-1} RDQ^H QDR^H + RDQ^H QDR^H \Sigma_z^{-1} \\ &\quad + \left(\frac{d}{dt}(P)^H P \right) RDR^H + RDR^H \left(P^H \frac{d}{dt}(P) \right). \end{aligned} \quad (10)$$

Due to the fact that $P^H P = I_r$, we have $\frac{d}{dt}(P)^H P + P^H \frac{d}{dt}(P) = 0$. Denote $M = P^H \frac{d}{dt}(P)$, then M is an antisymmetric matrix such that $M + M^H = 0$. Therefore, we have

$$\frac{d}{dt}(RDR^H) = M^H RDR^H + RDR^H M + \Sigma_z^{-1} RDQ^H QDR^H + RDQ^H QDR^H \Sigma_z^{-1}.$$

On the other hand, we have

$$\begin{aligned} \frac{d}{dt}(\Sigma_z) &= \frac{d}{dt}(P^H S_z P) \\ &= \frac{d}{dt}(P)^H S_z P + P^H \frac{d}{dt}(S_z) P + P^H S_z \frac{d}{dt}(P) \\ &= P^H \left(R_*^H D R_*^H - S_z \right) P + \left(\frac{d}{dt}(P)^H P \right) \Sigma_z + \Sigma_z \left(P^H \frac{d}{dt}(P) \right) \\ &= RDR^H - \Sigma_z + \left(\frac{d}{dt}(P_z)^H P_z \right) \Sigma_z + \Sigma_z \left(P_z^H \frac{d}{dt}(P_z) \right). \end{aligned}$$

Thus, we arrive at

$$\frac{d}{dt}(\Sigma_z(s, s)) = RDR^H(s, s) - \Sigma_z(s, s). \quad (11)$$

Since $\frac{d}{dt}(\text{offdiag}(\Sigma_z)) = 0$, we have $\text{offdiag}(RDR^H) = -M^H \Sigma_z - \Sigma_z M = M \Sigma_z + \Sigma_z M^H$.⁷

By substitution, we get

$$\begin{aligned} \frac{d}{dt}(RDR^H(j, j)) &= e_j^H (2M \Sigma_z M - \Sigma_z M M - M M \Sigma_z) e_j + \frac{2}{\sigma_j(Z)} \|QDR^H e_j\|_2^2 \\ &= 2 \sum_{k=1}^r (\sigma_j(Z) - \sigma_k(Z)) M(k, j)^2 + \frac{2}{\sigma_j(Z)} \|QDR^H e_j\|_2^2 \\ &= 2 \sum_{k \neq j} \frac{1}{\sigma_j - \sigma_k} RDR^H(k, j)^2 + \frac{2}{\sigma_j(Z)} \|QDR^H e_j\|_2^2. \end{aligned} \quad (12)$$

Note that

$$\begin{aligned} \sum_{j=1}^r \sum_{k=1}^r (\sigma_j(Z) - \sigma_k(Z)) M(k, j)^2 &= - \sum_{k=1}^r \sum_{j=1}^r (\sigma_k(Z) - \sigma_j(Z)) M(k, j)^2 \\ &= - \sum_{k=1}^r \sum_{j=1}^r (\sigma_j(Z) - \sigma_k(Z)) M(k, j)^2. \end{aligned}$$

⁷Here, $\text{offdiag}(A) := A - \text{diag}(A)$ which keeps the off-diagonal entries of A .

Thus, we obtain

$$\frac{d}{dt} \left(\sum_{j=1}^r \|D^{\frac{1}{2}} R^H e_j\|^2 \right) = \sum_{j=1}^r \frac{2}{\sigma_j(Z)} \|QDR^H e_j\|^2. \quad (13)$$

2) Denote $R = T_2 \Sigma_{RR} T_1^H$. Simple calculation gives $R^H R = T_1 \Sigma_{RR} T_1^H$ and $Q^H Q = T_1 \Sigma_{QQ} T_1^H$ with $\Sigma_{RR} + \Sigma_{QQ} = I_r$. We have

$$\begin{aligned} \sum_{j=1}^r RDR^H(j, j) &= \text{trace}(RDR^H) = \text{trace}(DR^H R) \\ &= \text{trace}(T_1^H DT_1 \Sigma_{RR}) \\ &= \sum_{j=1}^r t_j r_j. \end{aligned}$$

Here, for each $j \in [r]$, $t_j = \|D^{\frac{1}{2}} T_1 e_j\|^2 = \Omega(1)$ and $\Sigma_{RR}(j, j) = r_j$. Now, for a bounded sequence $\{Z_t\}$ we have

$$\begin{aligned} \sum_{j=1}^r \frac{2}{\sigma_j} \|QDR^H e_j\|^2 &\gtrsim \text{trace}(RDQ^H QDR^H) \\ &= \text{trace}(T_1^H DT_1 \Sigma_{QQ} T_1^H DT_1 \Sigma_{RR}). \end{aligned}$$

Denote $N = T_1^H DT_1$. We obtain the following estimate:

$$\begin{aligned} \sum_{j=1}^r \frac{2}{\sigma_j} \|QDR^H e_j\|^2 &\gtrsim \sum_{j=1}^r r_j(1 - r_j) N(j, j)^2 \\ &\gtrsim \sum_{j=1}^r r_j(1 - r_j). \end{aligned}$$

This gives

$$\frac{d}{dt} \left(\sum_{j=1}^r r_j \right) \gtrsim \sum_{j=1}^r r_j(1 - r_j).$$

□

B.8 Proof of Lemma 3.17

Proof. Assume $V_{z'} = V_z + tW$ with $W = (0, \dots, 0, V, 0, \dots, 0)$. By the evolution dynamics and **Assumption 1**, we have

$$\begin{aligned} (\phi_{\alpha, i} - I)(V_z) &= \alpha \cdot \frac{1}{\sigma_i} \widetilde{V}_z QDR^H e_i + \alpha \cdot V_z M e_i + \mathcal{O}(\alpha^2); \\ (\phi_{\alpha, i} - I)(V_{z'}) &= \alpha \cdot \frac{1}{\sigma_i} \widetilde{V}_{z'} Q'DR'^H e_i + \alpha \cdot V_{z'} M' e_i + \mathcal{O}(\alpha^2). \end{aligned}$$

Direct calculation gives $\Delta R = R' - R = V_{z'}^H U - V_z^H U = tW^H U$ and

$$\begin{aligned} \langle D(\phi_{\alpha,i} - I)|_{V_z}(V), V \rangle &= \langle \lim_{t \rightarrow 0} \frac{(\phi_{\alpha,i} - I)(V_{z'}) - (\phi_{\alpha,i} - I)(V_z)}{\|tV\|}, V \rangle + \mathcal{O}(\alpha^2) \\ &= \langle \lim_{t \rightarrow 0} \frac{\alpha \cdot \frac{1}{\sigma_i} \left(\widetilde{V}_{z'} Q' D R'^H e_i - \widetilde{V}_z Q D R^H e_i \right)}{t\|V\|}, V \rangle + \mathcal{O}(\alpha^2). \end{aligned}$$

The last equality holds because

$$\begin{aligned} \langle \lim_{t \rightarrow 0} \frac{1}{t} \cdot \alpha (V_z' M' - V_z M) e_i, V \rangle &= V^H \lim_{t \rightarrow 0} \frac{1}{t} \cdot \alpha (V_z \Delta M + tVM) e_i \\ &= M(i, i) = 0. \end{aligned}$$

Thus, we get

$$\begin{aligned} \langle D(\phi_{\alpha,i} - I)|_{V_z}(V), V \rangle &= \langle \lim_{t \rightarrow 0} \frac{\alpha \cdot \frac{1}{\sigma_i} \left(\widetilde{V}_{z'} Q' D R'^H e_i - \widetilde{V}_z Q D R^H e_i \right)}{t\|V\|}, V \rangle + \mathcal{O}(\alpha^2) \\ &= \frac{\alpha}{\sigma_i} \lim_{t \rightarrow 0} \frac{1}{t} \langle (U - V_{z'} R') D R'^H e_i - \widetilde{V}_z Q D R^H e_i, V \rangle + \mathcal{O}(\alpha^2) \\ &= \frac{\alpha}{\sigma_i} \lim_{t \rightarrow 0} \frac{1}{t} \langle U D R'^H e_i - U D R^H e_i + U D R^H e_i - \widetilde{V}_z Q D R^H e_i - V_{z'} R' D R'^H e_i, V \rangle \\ &\quad + \mathcal{O}(\alpha^2) \\ &= \frac{\alpha}{\sigma_i} \lim_{t \rightarrow 0} \frac{1}{t} \langle U D \Delta R^H e_i + (U - \widetilde{V}_z Q) D R^H e_i - V_{z'} R' D R'^H e_i, V \rangle \\ &\quad + \mathcal{O}(\alpha^2) \\ &= \frac{\alpha}{\sigma_i} \left(V^H U D U^H V - V^H V_z R D U^H V - V^H V_z W^H U D R^H e_i - V^H W R D R^H e_i \right) \\ &\quad + \mathcal{O}(\alpha^2) \\ &= \frac{\alpha}{\sigma_i} V^H \widetilde{V}_z Q D Q^H \widetilde{V}_z V - \alpha \cdot \frac{R D R^H(i, i)}{\sigma_i} + \mathcal{O}(\alpha^2). \end{aligned}$$

Note that at the beginning, $\sigma_i = \Omega(1)$, and $R D R^H(i, i) = \Omega(\frac{1}{n})$, and $\sigma_i = R D R^H(i, i) - \sigma_i$. Denote $C = \Omega(1)$ as a constant that $\frac{R D R^H(i, i)}{\sigma_i} \leq C$. Therefore, by controlling $\alpha \cdot \frac{R D R^H(i, i)}{\sigma_i} + \mathcal{O}(\alpha^2) \leq \mathcal{O}(\alpha)$, we have

$$\langle D(\phi_i - I)|_{V_z}(V), V \rangle = \frac{\alpha}{\sigma_i} V^H \widetilde{V}_z Q D Q^H \widetilde{V}_z^H V + \mathcal{O}(\alpha).$$

This completes the proof. □

B.9 Proof of Lemma 3.19

Proof. Assume $p(\cdot)$ is the density function of μ , then we have

$$\begin{aligned} \int_{g^{-1}(U)} p(x)dx &= \mu(g^{-1}(U)) \\ &= \int_U p(g^{-1}(y))|\det Dg^{-1}|dy \\ &\leq C_0 \int_U p(g^{-1}(y))dy \\ &\lesssim \mu(U). \end{aligned}$$

The second equality holds by change of variables; the first inequality holds due to the fact that for all $y \in U$, $|\det Dg^{-1}| \leq C_0 = \frac{1}{c_0}$, the second inequality holds because for all $y \in U$, we have $p(y) = p(g^{-1}(y))$ due to the uniform distribution. Then, we have the following

$$\mu(g^{-1}(U)) \leq \frac{1}{c_0}\mu(U).$$

□

B.10 Proof of Lemma 3.22

Proof. For simplicity, the following proof is based on the symmetric case where $Z = U_z \Sigma U_z^H$ and $X = UDU^H$.

1) Let $\tilde{U}_z \in \mathbb{F}^{n \times (n-r)}$ be the orthogonal complement of U_z , and

$$X = (U_z, \tilde{U}_z) \tilde{X} (U_z, \tilde{U}_z)^H = (U_z, \tilde{U}_z) \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{pmatrix} \begin{pmatrix} U_z^H \\ \tilde{U}_z^H \end{pmatrix}.$$

Then, we get

$$\begin{aligned} \|P_{T_Z}(Z - X)\|_F &= \left\| \begin{pmatrix} X_{11} - \Sigma_z & X_{12} \\ X_{21} & 0 \end{pmatrix} \right\|_F = \|X_{11} - \Sigma_z\|_F^2 + \|X_{12}\|_F^2 + \|X_{21}\|_F^2, \\ \|Z - X\|_F^2 &= \|P_{T_Z}(Z - X)\|_F^2 + \|X_{22}\|_F^2. \end{aligned}$$

We also assume that $U = U_z R + \tilde{U}_z S$, where $R \in \mathbb{F}^{r \times r}$ and $S \in \mathbb{F}^{(n-r) \times r}$. Then, we obtain

$$X_{11} = RDR^H, \quad X_{12} = RDS^H, \quad X_{21} = SDR^H, \quad X_{22} = SDS^H.$$

The goal is to find lower and upper bounds for

$$s = \frac{\|P_{T_Z}(Z - X)\|_F^2}{\|Z - X\|_F^2}.$$

It is obvious that $s \leq 1$. To identify the lower bound we consider

$$\phi = \frac{\|Z - X\|_F^2}{\|X_{22}\|_F^2}, \quad s = \frac{\phi - 1}{\phi}, \quad \phi \geq 1.$$

For fixed Σ and D , ϕ can be seen as a function of R . We express $\phi(R)$ in terms of ϕ_1 and ϕ_2 defined below:

$$\begin{aligned} \phi(R) &= \frac{\phi_1(R)}{\phi_2(R)}, \quad R \in \mathcal{D} \subset \mathbb{F}^{r \times r}, \quad \mathcal{D} = \{R : 0 \preceq R^H R \preceq I_r\}, \\ \text{where } \phi_1(R) &= \|X\|_F^2 + \|\Sigma\|_F^2 - 2\langle \Sigma, X_{11} \rangle \\ &= \|D\|_F^2 + \|\Sigma\|_F^2 - 2\text{tr}(\Sigma R D R^H), \\ \phi_2(R) &= \text{tr}\left((S D S^H)^2\right) = \text{tr}\left((D S^H S)^2\right) \\ &= \text{tr}\left((D(I - R^H R))^2\right). \end{aligned}$$

To minimize t for given D and Σ , the first-order condition is obtained by taking derivative of t over R :

$$\begin{aligned} \frac{\partial t}{\partial R} &= \frac{1}{\phi_2^2} \left(\frac{\partial \phi_1}{\partial R} \phi_2 - \frac{\partial \phi_2}{\partial R} \phi_1 \right) \\ &= \frac{1}{\phi_2^2} (-4\phi_2 \Sigma R D + 4\phi_1 R D (I - R^H R) D) \\ &= \frac{4}{\phi_2^2} (-\phi_2 \Sigma R D + \phi_1 R D S^H S D). \end{aligned}$$

Imposing $\partial t / \partial R = 0$ gives

$$\Sigma R D = \phi \cdot R D S^H S D. \quad (14)$$

We now claim that this first-order condition *cannot* be satisfied in the interior of the domain \mathcal{D} . To see this, note that the above equation (14) gives

$$\Sigma R D R^H = \phi \cdot R D S^H S D R^H, \quad \text{i.e.} \quad \langle \Sigma, X_{11} \rangle = \phi \|X_{12}\|_F^2 = \phi \|X_{21}\|_F^2.$$

Thus, we have

$$\begin{aligned} \phi &= \frac{\|Z - X\|_F^2}{\|X_{22}\|_F^2} = \frac{\|\Sigma\|_F^2 + \|X_{11}\|_F^2 - 2\langle \Sigma, X_{11} \rangle + \|X_{12}\|_F^2 + \|X_{21}\|_F^2 + \|X_{22}\|_F^2}{\|X_{22}\|_F^2} \\ &= \frac{\|\Sigma\|_F^2 + \|X_{11}\|_F^2 - (2 - 2/\phi)\langle \Sigma, X_{11} \rangle}{\|X_{22}\|_F^2} + 1 \\ &= \frac{\|X_{11} - (1 - 1/\phi)\Sigma\|_F^2}{\|X_{22}\|_F^2} + \left(\frac{2}{\phi} - \frac{1}{\phi^2}\right) \frac{\|\Sigma\|_F^2}{\|X_{22}\|_F^2} + 1 \\ &\geq \left(\frac{2}{\phi} - \frac{1}{\phi^2}\right) \frac{\|\Sigma\|_F^2}{\|X_{22}\|_F^2} + 1. \end{aligned}$$

Note that when $R \in \text{int}(\mathcal{D})$, R is full-rank. Since D, Σ are also full-rank, (14) gives $\phi^2 = \|\Sigma\|_F^2 / \|X_{22}\|_F^2$ by some simple matrix manipulation:

$$\begin{aligned}
& \Sigma R D = \phi \cdot R D S^H S D \\
\Rightarrow & R^{-1} \Sigma R = \phi \cdot D S^H S \\
\Rightarrow & \text{tr} \left((R^{-1} \Sigma R)^2 \right) = \phi^2 \text{tr} \left((D S^H S)^2 \right) \\
\Rightarrow & \text{tr} \left(\Sigma^2 \right) = \phi^2 \text{tr} \left((S D S^H)^2 \right) \\
\Rightarrow & \|\Sigma\|_F^2 = \phi^2 \|X_{22}\|_F^2.
\end{aligned}$$

Hence, we obtain

$$\phi \geq \left(\frac{2}{\phi} - \frac{1}{\phi^2} \right) \phi^2 + 1 = 2\phi - 1 + 1 = 2\phi,$$

which contradicts the fact that $\phi \geq 1$.

Therefore, either the extreme values of t are only achieved on $\partial\mathcal{D}$, or the full-rankness of D is violated. In either case, the RHS of (14) is rank-deficient. So the LHS of (14) (and thus R and X_{11}) is also rank-deficient. Therefore, we conclude that

$$\phi - 1 \geq \frac{\|X_{11} - \Sigma\|_F^2}{\|X_{22}\|_F^2} \geq \frac{\|X_{11} - \Sigma\|_F^2}{\|X\|_F^2} \geq \frac{\sigma_r^2}{\|X\|_F^2},$$

where σ_r is the smallest nonzero eigenvalue of Z . Finally, we obtain

$$s \geq \frac{\sigma_r^2}{\sigma_r^2 + \|X\|_F^2}.$$

Now we have $\|P_{T_Z}(Z - X)\|_F^2 \geq s \|Z - X\|_F^2$ for $s > 0$ as long as Z is of rank r and the smallest nonzero singular value of Z is bounded away from 0.

- 2) By Lemma C.2, when $\|Z - X\|_F < \frac{d_r}{2}$, we have $\sigma_r > \frac{d_r}{2}$. Further, by (1) we have $\frac{\|P_{T_Z}(Z - X)\|_F}{\|Z - X\|_F} \gtrsim d_r$. By Lemma 3.11, we have the locally linear convergence towards X with $c = -\log(1 - \Omega(\alpha\sigma_r(X)^2))$.

□

B.11 Proof of Theorem 4.1

Proof. The proof of Theorem 4.1 can be reduced to deducing Conditions (L) and (D) from Conditions (1) and (2) in the assumptions. Note that in the setting of $T : \mathcal{M} \rightarrow \mathbb{R}^m$, $T(Z) =$

$$\frac{1}{\sqrt{m}} \begin{pmatrix} \langle A_1, Z \rangle \\ \langle A_2, Z \rangle \\ \dots \\ \langle A_m, Z \rangle \end{pmatrix}, f(Z) = \frac{1}{2} \|T(Z) - T(X)\|_2^2, \text{ Conditions (L) and (D) can be formulated as follows.}$$

1) Condition (L): Łojasiewicz gradient inequality

$$\left(\frac{1}{2m} \sum_j \langle A_j, Z - X \rangle^2 \right)^{1-\theta} \leq C_l \|P_{T_Z} \left(\frac{1}{m} \sum_j \langle A_j, Z - X \rangle A_j \right)\|_F$$

holds with $\theta = \frac{1}{2}$.

2) Condition (D):

$$f_k - f_{k+1} \geq \sigma \|P_{T_{Z_k}} \left(\frac{1}{m} \sum_j \langle A_j, Z_k - X \rangle A_j \right)\|_F \|Z_{k+1} - Z_k\|_F.$$

The proof now goes as follows.

1) To prove condition (L): By (1) and (2) we have

$$\begin{aligned} \|P_{T_Z}(\nabla f(Z))\|_F &\geq \frac{1}{C_3} \|Z - X\|_F \\ &\geq \frac{1}{C_2 C_3} \|T(Z) - T(X)\| \\ &= \frac{\sqrt{2}}{C_2 C_3} |f(Z) - f(X)|^{\frac{1}{2}}, \end{aligned}$$

which implies that (L) holds with $\theta = \frac{1}{2}$ and $C_l = \frac{\sqrt{2}C_2C_3}{2} > 0$ is an absolute constant.

2) To prove condition (D): we first consider

$$\begin{aligned} f_k - f_{k+1} &= \frac{1}{2} \|T(Z_k) - T(X)\|_2^2 - \frac{1}{2} \|T(Z_{k+1}) - T(X)\|_2^2 \\ &= \frac{1}{2} \langle T(Z_k + Z_{k+1} - 2X), T(Z_k - Z_{k+1}) \rangle \\ &= \langle T^*T(X - Z_k), Z_{k+1} - Z_k \rangle - \frac{1}{2} \|T(Z_{k+1} - Z_k)\|_F^2. \end{aligned}$$

Assume that $Z_{k+1} = Z_k + \alpha \tilde{\xi}_k$, and $-\alpha P_{T_{Z_k}}(\nabla f(Z_k)) = \alpha \xi_k$. By first-order retraction property and condition (1), we get

$$\begin{aligned} f_k - f_{k+1} &\geq \langle -\nabla f(Z_k), \alpha \tilde{\xi}_k \rangle - \frac{C_2^2 \alpha^2}{2} \|\tilde{\xi}_k\|^2 \\ &= \langle -\nabla f(Z_k), \alpha \xi_k \rangle + o(\alpha \|\xi_k\|_F) \\ &= \langle -\nabla f(Z_k), -P_{T_{Z_k}}(\nabla f(Z_k)) \rangle + o(\alpha \|\xi_k\|_F) \\ &= \alpha \|P_{T_{Z_k}}(\nabla f(Z_k))\|_F^2 + o(\alpha \|\xi_k\|_F). \end{aligned}$$

On the other hand, we also obtain

$$\begin{aligned} \sigma \|P_{T_{Z_k}}(\nabla f(Z_k))\|_F \|Z_{k+1} - Z_k\|_F &= \sigma \|P_{T_{Z_k}}(\nabla f(Z_k))\|_F \|\alpha \tilde{\xi}_k\|_F \\ &= \sigma \alpha \|P_{T_{Z_k}}(\nabla f(Z_k))\|_F^2 + o(\alpha \|\xi_k\|_F). \end{aligned}$$

By choosing $\sigma > 0$ small enough, we have

$$f_k - f_{k+1} \geq \sigma \|P_{T_{Z_k}} \left(\frac{1}{m} \sum_j \langle A_j, Z_k - X \rangle A_j \right)\|_F \|Z_{k+1} - Z_k\|_F,$$

i.e. Condition (D) holds.

From Theorem 3.4, we conclude that projected gradient descent for the least squares loss function $f(Z) = \frac{1}{2} \|T(Z) - T(X)\|_2^2$ converges to its global minimum linearly. \square

C Auxiliary Lemmas

Lemma C.1 (The $\sin(\theta)$ Theorem, [19]). *Let A be a Hermitian operator. Assume that*

$$A = \begin{pmatrix} E_0 & E_1 \end{pmatrix} \begin{pmatrix} A_0 & 0 \\ 0 & A_1 \end{pmatrix} \begin{pmatrix} E_0^H \\ E_1^H \end{pmatrix}$$

is an invariant subspace decomposition (i.e. generalized eigenvalue decomposition) of A . Let

$$B = A + \Delta, \quad B = \begin{pmatrix} F_0 & F_1 \end{pmatrix} \begin{pmatrix} B_0 & 0 \\ 0 & B_1 \end{pmatrix} \begin{pmatrix} F_0^H \\ F_1^H \end{pmatrix}.$$

Let Θ_0 be the angle matrix between subspaces E_0 and F_0 . Define the residual as

$$R = BE_0 - E_0A_0.$$

If there is an interval $[\beta, \alpha]$ and $\delta > 0$, such that the spectrum of A_0 lies entirely in $[\beta, \alpha]$, while that of B_1 lies entirely in $(-\infty, \beta - \delta) \cup [\alpha + \delta, +\infty)$, then for every unitary-invariant norm, we have

$$\delta \|\sin \Theta_0\| \leq \|R\|.$$

In particular, this holds true for the matrix 2-norm and the Frobenius norm.

Lemma C.2. (Hoffman-Wielandt Theorem)

- 1) *Assume $Z, Z' \in \mathbb{R}^{n \times n}$ or $\mathbb{C}^{n \times n}$ are normal matrices with their corresponding ordered spectra to be $\{\lambda_j\}$ and $\{\tilde{\lambda}_j\}$. Then, we have*

$$\sqrt{\sum_{j=1}^n |\tilde{\lambda}_j - \lambda_j|^2} \leq \|Z' - Z\|_F.$$

- 2) *Assume $Z, Z' \in \mathbb{R}^{m \times n}$ or $\mathbb{C}^{m \times n}$, and denote their singular values in descending order as $\{\sigma_j\}$ and $\{\tilde{\sigma}_j\}$. Then, we have*

$$\sqrt{\sum_{j=1}^n |\tilde{\sigma}_j - \sigma_j|^2} \leq \|Z' - Z\|_F.$$

Lemma C.3. (Dynamical low-rank approximation, [30]) Assume $Z = USV^H$ with $M := P_{T_Z}(\frac{d}{dt}g(Z)) = \nabla_{\mathcal{M}_r}g(Z)$, then gradient flow satisfies:

$$\begin{aligned}\frac{d}{dt}S &= U^H M V, \\ \frac{d}{dt}U &= P_U^\perp M V S^{-1}, \\ \frac{d}{dt}V &= P_V^\perp M^H U S^{-\top}.\end{aligned}$$

Here, $P_U^\perp = I - UU^H$ and $P_V^\perp = I - VV^H$.