

Guarantees of Riemannian Optimization for Low Rank Matrix Completion

Ke Wei* Jian-Feng Cai† Tony F. Chan‡ Shingyu Leung†

March 20, 2016

Abstract

We study the Riemannian optimization methods on the embedded manifold of low rank matrices for the problem of matrix completion, which is about recovering a low rank matrix from its partial entries. Assume m entries of an $n \times n$ rank r matrix are sampled independently and uniformly with replacement. We first prove that with high probability the Riemannian gradient descent and conjugate gradient descent algorithms initialized by one step hard thresholding are guaranteed to converge linearly to the measured matrix provided

$$m \geq C_\kappa n^{1.5} r \log^{1.5}(n),$$

where C_κ is a numerical constant depending on the condition number of the measured matrix. The sampling complexity has been further improved to

$$m \geq C_\kappa n r^2 \log^2(n)$$

via the resampled Riemannian gradient descent initialization. The analysis of the new initialization procedure relies on an asymmetric restricted isometry property of the sampling operator and the curvature of the low rank matrix manifold. Numerical simulation shows that the algorithms are able to recover a low rank matrix from nearly the minimum number of measurements.

Keywords. Matrix completion, Riemannian optimization, low rank matrix manifold, tangent space, gradient descent and conjugate gradient descent methods

Mathematics Subject Classification. 15A29, 41A29, 65F10, 68Q25, 15A83, 53B21, 90C26, 65K05

1 Introduction

The problem of matrix completion attempts to recover a low rank matrix from a subset of sampled entries. This problem arises from a wide variety of practical context, such as model reduction [33], pattern recognition [18], and machine learning [4, 5]. Since the data matrix is assumed to be low rank, it is natural to seek the lowest rank matrix consistent with the observed entries by solving a rank minimization problem

$$\min_{Z \in \mathbb{R}^{n \times n}} \text{rank}(Z) \text{ subject to } \mathcal{P}_\Omega(Z) = \mathcal{P}_\Omega(X), \quad (1)$$

*Department of Mathematics, University of California at Davis, kewe@math.ucdavis.edu

†Department of Mathematics, Hong Kong University of Science and Technology, {jfc, masyleung}@ust.hk

‡Office of the President, Hong Kong University of Science and Technology, tonyfchan@ust.hk

where $X \in \mathbb{R}^{n \times n}$ is the underlying matrix to be reconstructed, Ω is a subset of indices for the known entries, and \mathcal{P}_Ω is the associated sampling operator which acquires only the entries indexed by Ω . Here we restrict our discussion to square matrices for ease of exposition, but emphasize that all the results can be extended straightforwardly to the case where the measured matrix is rectangular. Problem (1) is generally NP-hard [24] and computationally intractable. In a seminal paper of Candès and Recht [13], the authors studied nuclear norm minimization for matrix completion, where the rank objective in (1) is replaced by the nuclear norm of matrices which is the sum of the singular values. It was observed in the same paper that we cannot expect to recover all the low rank matrices from their partial known entries. For example, if there exists only one nonzero entry in a matrix, we need to see almost all of its entries to infer that it is not a zero matrix. This observation motivates the following assumption.

A0 ([13]). *Let $X \in \mathbb{R}^{n \times n}$ be a rank r matrix with the reduced singular value decomposition (SVD) $X = U\Sigma V^*$. We assume X is μ_0 -incoherent; that is, there exists an absolute numerical constant $\mu_0 > 0$ such that*

$$\|\mathcal{P}_U(e_i)\| \leq \sqrt{\frac{\mu_0 r}{n}} \quad \text{and} \quad \|\mathcal{P}_V(e_j)\| \leq \sqrt{\frac{\mu_0 r}{n}} \quad (2)$$

for $1 \leq i, j \leq n$. Here e_ℓ ($\ell = i, j$) is the ℓ -th canonical basis of \mathbb{R}^n , and \mathcal{P}_U and \mathcal{P}_V are the orthogonal projections onto the column and row spaces of X respectively.

A matrix is μ_0 -incoherent implies that its singular vectors are weakly correlated with the canonical basis. In addition, (2) is equivalent to $\|U^{(i)}\| \leq \sqrt{\mu_0 r/n}$ and $\|V^{(j)}\| \leq \sqrt{\mu_0 r/n}$, where $U^{(i)}$ and $V^{(j)}$ are the i -th and j -th rows of U and V respectively. We also require that the measured matrix cannot be too spiky.

A1 ([29]). *Let $X \in \mathbb{R}^{n \times n}$ be a rank r matrix. We assume that exists another absolute numerical constant μ_1 such that*

$$\|X\|_\infty \leq \mu_1 \sqrt{\frac{r}{n^2}} \|X\|. \quad (3)$$

This manuscript investigates recovery guarantees of a class of Riemannian gradient descent and conjugate gradient algorithms for matrix completion under the sampling with replacement model. *We first establish the local convergence of the Riemannian gradient decent algorithm and a restarted variant of the Riemannian conjugate gradient descent algorithm. Then we prove that $O(nr^2 \log^2(n))$ number of measurements are sufficient for the Riemannian optimization algorithms to converge linearly to the underlying low rank matrix when the algorithms are properly initialized.*

1.1 Notations and Organization of the Manuscript

The rest of the manuscript is organized as follows. We first summarize the notations used throughout this manuscript in the remainder of this section. In Sec. 2, we present the Riemannian gradient descent and conjugate gradient descent algorithms based on the embedded manifold of low rank matrices from the perspective of iterative hard thresholding algorithms. Then we present the theoretical guarantees of the algorithms. Empirical observations in Sec. 3 demonstrates the efficiency and robustness of the Riemannian optimization algorithms. The proofs of the main results are

presented in Sec. 4 and Sec. 5 concludes this manuscript with potential future directions. Finally, the Appendix provides proofs for the supporting technical lemmas.

Throughout this manuscript, we denote matrices by uppercase letters and vectors by lowercase letters. In particular, X denotes the $n \times n$ rank r matrix to be reconstructed, with $\sigma_{\min}(X)$ and $\sigma_{\max}(X)$ being its smallest nonzero and largest singular values respectively. The condition number of X is denoted by κ , which is defined as $\kappa = \frac{\sigma_{\max}(X)}{\sigma_{\min}(X)}$. We always assume that X obeys the conditions set out in **A0** and **A1**. For any matrix Z , the spectral norm of Z is denoted by $\|Z\|$, the Frobenius norm is denoted by $\|Z\|_F$, and the maximum magnitude of its entries is denoted by $\|Z\|_{\infty}$. We use $Z^{(i)}$ to represent the i -th row of Z . The Euclidean norm of a vector x is denoted by $\|x\|$. Operators that map matrices to matrices are denoted by calligraphic letters. In particular, \mathcal{I} denotes the identity operator. The spectral norm of a linear operator \mathcal{A} is denoted by $\|\mathcal{A}\|$.

Given a collection of indices Ω with $|\Omega| = m$ (counting multiplicity) for the observed entries, we define $p = \frac{m}{n^2}$ which is the probability of each entry being observed, and \mathcal{P}_{Ω} represents the sampling operator that maps a matrix to its component-wise product with the matrix $\sum_{(i,j) \in \Omega} e_i e_j^*$. In the main results, ε_0 is a positive numerical constant which controls the size of the contraction neighborhood of an algorithm and its value only depends on the algorithm. Finally, we use C to denote an absolute numerical constant whose value may change according to context.

2 Algorithms and Main Results

We consider the *sampling with replacement model* for Ω in which each index is sampled independently from the uniform distribution on $\{1, \dots, n\} \times \{1, \dots, n\}$. This model can be viewed as a proxy for the uniform sampling model and the failure probability under the uniform sampling model is less than or equal to the failure probability under the sampling with replacement model [38]. Though there may exist duplicates in Ω , the maximum number of duplication can be upper bounded with high probability.

Lemma 2.1 ([38]). *With probability at least $1 - n^{2-2\beta}$, the maximum number of repetitions of any entry in Ω is less than $\frac{8}{3}\beta \log(n)$ for $n \geq 9$ and $\beta > 1$.*

It follows immediately from Lem. 2.1 that with high probability we have $\|\mathcal{P}_{\Omega}\| \leq \frac{8}{3}\beta \log(n)$.

Many computationally efficient algorithms have been designed to target the matrix completion problem directly by considering the following non-convex formulation¹

$$\min_{Z \in \mathbb{R}^{n \times n}} \frac{1}{2} \langle Z - X, \mathcal{P}_{\Omega}(Z - X) \rangle \quad \text{subject to } \text{rank}(Z) = r, \quad (4)$$

see [7, 25, 30, 43, 23, 50, 44, 47, 36, 34, 35, 26, 8] and references therein for an incomplete review. Normalized iterative hard thresholding [43] (NIHT, also known as SVP [25] or IHT [21] when the stepsize is fixed) is a projected gradient descent algorithm which updates the estimate along the gradient descent direction of the objective function in (4), followed by the projection onto the set of rank r matrices via the hard thresholding operator:

$$X_{l+1} = \mathcal{H}_r(X_l + \alpha_l \mathcal{P}_{\Omega}(X - X_l)), \quad (5)$$

¹Note that the objective function in (4) is the same as $\frac{1}{2} \sum_{(i,j) \in \Omega} (Z_{ij} - X_{ij})^2$, but it might not be equal to $\frac{1}{2} \|\mathcal{P}_{\Omega}(Z - X)\|_F^2$ because in general \mathcal{P}_{Ω} is not a projection (i.e., $\mathcal{P}_{\Omega}^2 \neq \mathcal{P}_{\Omega}$) when Ω is sampled with replacement.

Algorithm 1 Riemannian Gradient Descent (RGrad)

Initialization: $X_0 = U_0 \Sigma_0 V_0^*$

for $l = 0, 1, \dots$ **do**

1. $G_l = \mathcal{P}_\Omega(X - X_l)$
2. $\alpha_l = \frac{\|\mathcal{P}_{T_l}(G_l)\|_F^2}{\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(G_l) \rangle}$
3. $W_l = X_l + \alpha_l \mathcal{P}_{T_l}(G_l)$
4. $X_{l+1} = \mathcal{H}_r(W_l)$

end for

where the hard thresholding operator \mathcal{H}_r first computes the SVD of a matrix and then sets all but the r largest singular values to zero

$$\mathcal{H}_r(Z) := P \Lambda_r Q^* \quad \text{where } Z = P \Lambda Q^* \text{ is the SVD of } Z, \text{ and } \Lambda_r(i, i) := \begin{cases} \Lambda(i, i) & i \leq r \\ 0 & i > r. \end{cases} \quad (6)$$

NIHT suffers from the slow asymptotic convergence rate of first order methods as a gradient descent algorithm. To overcome this slow convergence rate of NIHT, a conjugate gradient iterative hard thresholding (CGIHT) algorithm has been developed in [7]. In each iteration of CGIHT, the new search direction is computed as a weighted sum of the gradient descent direction and the past search direction, and the selection of the weights guarantees that the new search direction is conjugate orthogonal to the past search direction when projected onto the column space of the current estimate. In both NIHT and CGIHT, we need to compute the largest r singular values and their associated singular vectors of an $n \times n$ matrix in each iteration. Although Krylov subspace methods can be applied to compute this truncated SVD, it is still computationally expensive especially when n is large and r is moderate. The Riemannian optimization algorithms based on the embedded manifold of rank r matrices can reduce the computational cost of the truncated SVD dramatically by exploring the local structure of low rank matrices.

We first present the Riemannian gradient descent method for matrix completion in Alg. 1. Let $X_l = U_l \Sigma_l V_l^*$ be the current estimate and T_l be the tangent space of the rank r matrix manifold at X_l ; that is,

$$T_l = \{U_l Z_1^* + Z_2 V_l^* : Z_1, Z_2 \in \mathbb{R}^{n \times r}\}.$$

The new estimate X_{l+1} is obtained by updating X_l along $\mathcal{P}_{T_l}(G_l)$, the gradient descent direction projected onto the tangent space T_l , using the steepest stepsize α_l , and then followed by hard thresholding the estimate back onto the set of rank r matrices.

In the Riemannian conjugate gradient descent method (Alg. 2), the search direction is a linear combination of the projected gradient descent direction and the past search direction projected onto the tangent space of the current estimate. The orthogonalization weight β_l in Alg. 2 is selected in a way such that $\mathcal{P}_{T_l}(P_l)$ is conjugate orthogonal to $\mathcal{P}_{T_l}(P_{l-1})$, which follows from CGIHT in [7]. It is worth noting that there are other selections for β_l following from the non-linear conjugate gradient descent method in convex optimization [47]. In particular, the Riemannian conjugate gradient descent algorithm (LRGeomCG) developed in [47] for matrix completion uses the Polak-Ribière+ selection for β_l . In this manuscript, we interpret the Riemannian optimization methods directly as iterative hard thresholding algorithms with subspace projections. For the differential geometry ideas behind the algorithms, we refer the readers to [1].

Algorithm 2 Riemannian Conjugate Gradient Descent (RCG)

Initialization: $X_0 = U_0 \Sigma_0 V_0^*$, $\beta_0 = 0$ and $P_{-1} = 0$

for $l = 0, 1, \dots$ **do**

1. $G_l = \mathcal{P}_\Omega(X - X_l)$
2. $\beta_l = -\frac{\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle}{\langle \mathcal{P}_{T_l}(P_{l-1}), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle}$
3. $P_l = \mathcal{P}_{T_l}(G_l) + \beta_l \mathcal{P}_{T_l}(P_{l-1})$
4. $\alpha_l = \frac{\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle}$
5. $W_l = X_l + \alpha_l P_l$
6. $X_{l+1} = \mathcal{H}_r(W_l)$

end for

Though we still need to compute the truncated SVD of W_l to find its best rank r approximation, the truncated SVD can be computed efficiently using two QR factorizations of $n \times r$ matrices and one full SVD of a $2r \times 2r$ matrix as matrices in T_l are at most rank $2r$. To see this, notice that W_l in Algs. 1 and 2 can be rewritten as

$$W_l = X_l + \mathcal{P}_{T_l}(Z_l),$$

where $Z_l = \alpha_l G_l$ in Alg. 1 and $Z_l = \alpha_l(G_l + \beta_l P_{l-1})$ in Alg. 2. So

$$\begin{aligned} W_l &= U_l \Sigma_l V_l^* + U_l U_l^* Z_l + Z_l V_l V_l^* - U_l U_l^* Z_l V_l V_l^* \\ &= U_l (\Sigma_l + U_l^* Z_l V_l) V_l^* + U_l U_l^* Z_l (I - V_l V_l^*) + (I - U_l U_l^*) Z_l V_l V_l^* \\ &:= U_l (\Sigma_l + U_l^* Z_l V_l) V_l^* + U_l Y_1^* + Y_2 V_l^*. \end{aligned}$$

Let $Y_1 = Q_1 R_1$, $Y_2 = Q_2 R_2$ be the QR factorizations of Y_1 and Y_2 respectively. Then $U_l^* Q_2 = 0$ and $V_l^* Q_1 = 0$, and we have

$$\begin{aligned} W_l &= [U_l \quad Q_2] \begin{bmatrix} \Sigma_l + U_l^* Z_l V_l & R_1^* \\ & R_2 \\ & & 0 \end{bmatrix} \begin{bmatrix} V_l^* \\ Q_1^* \end{bmatrix} \\ &:= [U_l \quad Q_2] M_l \begin{bmatrix} V_l^* \\ Q_1^* \end{bmatrix}. \end{aligned}$$

Since $[U_l \quad Q_2]$ and $[V_l \quad Q_1]$ are both unitary matrices, the SVD of W_l can be obtained from the SVD of M_l , which is a $2r \times 2r$ matrix.

2.1 Local Convergence

We first identify a small neighborhood around the measured low rank matrix such that for any given initial guess in this neighborhood, Algs. 1 and 2 will converge linearly to the true solution. For the Riemannian gradient descent algorithm (Alg. 1), we have the following theorem.

Theorem 2.2 (Local Convergence of Riemannian Gradient Descent). *Let $X \in \mathbb{R}^{n \times n}$ be the measured rank r matrix and T be the tangent space of the rank r matrix manifold at X . Suppose*

$$\|\mathcal{P}_\Omega\| \leq \frac{8}{3} \beta \log(n), \tag{7}$$

$$\|\mathcal{P}_T - p^{-1}\mathcal{P}_T\mathcal{P}_\Omega\mathcal{P}_T\| \leq \varepsilon_0, \quad (8)$$

$$\frac{\|X_0 - X\|_F}{\sigma_{\min}(X)} \leq \frac{3p^{1/2}\varepsilon_0}{16\beta \log(n)(1 + \varepsilon_0)}, \quad (9)$$

where $\beta > 1$, and ε_0 is a positive numerical constant such that

$$\nu_g = \frac{18\varepsilon_0}{1 - 4\varepsilon_0} < 1. \quad (10)$$

Then the iterates X_l generated by Alg. 1 satisfy

$$\|X_l - X\|_F \leq \nu_g^l \|X_0 - X\|_F.$$

To state a similar result for the Riemannian conjugate gradient descent algorithm, we introduce a restarted variant of Alg. 2; that is, β_l is set 0 and restarting occurs as long as either of the following conditions is violated

$$\frac{|\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle|}{\|\mathcal{P}_{T_l}(G_l)\|_F \|\mathcal{P}_{T_l}(P_{l-1})\|_F} \leq \kappa_1, \quad \|\mathcal{P}_{T_l}(G_l)\|_F \leq \kappa_2 \|\mathcal{P}_{T_l}(P_{l-1})\|_F. \quad (11)$$

The restarting conditions are introduced not only for the sake of proof, but also to improve the robustness of the non-linear conjugate gradient descent methods [37]. The first restarting condition guarantees that the residual will be substantially orthogonal to the past search direction when projected onto the tangent space of current estimate so that the new search direction can be sufficiently gradient related. In the classical CG algorithm for linear systems, the residual is exactly orthogonal to all the past search directions. Roughly speaking, the second restarting condition implies that the projection of current residual cannot be too large when compared to the projection of the past residual since the search direction is gradient related by the first restarting condition. In our implementations, we take $\kappa_1 = 0.1$ and $\kappa_2 = 1$.

Theorem 2.3 (Local Convergence of restarted Riemannian Conjugate Gradient Descent). *Let $X \in \mathbb{R}^{n \times n}$ be the measured rank r matrix and T be the tangent space of the rank r matrix manifold at X . Suppose*

$$\|\mathcal{P}_\Omega\| \leq \frac{8}{3}\beta \log(n), \quad (12)$$

$$\|\mathcal{P}_T - p^{-1}\mathcal{P}_T\mathcal{P}_\Omega\mathcal{P}_T\| \leq \varepsilon_0, \quad (13)$$

$$\frac{\|X_0 - X\|_F}{\sigma_{\min}(X)} \leq \frac{3p^{1/2}\varepsilon_0}{16\beta \log(n)(1 + \varepsilon_0)}, \quad (14)$$

where $\beta > 1$, and ε_0 is a positive numerical constant. Assume that ε_0 obeys

$$\tau_1 + \tau_2 < 1, \quad (15)$$

where

$$\tau_1 = \frac{18\varepsilon_0 - 10\kappa_1\varepsilon_0(1 + 4\varepsilon_0)}{(1 - 4\varepsilon_0) - \kappa_1(1 + 4\varepsilon_0)} + \frac{4\kappa_2\varepsilon_0 + \kappa_1\kappa_2}{1 - 4\varepsilon_0}, \quad \tau_2 = \frac{8\kappa_2\varepsilon_0 + 2\kappa_1\kappa_2}{1 - 4\varepsilon_0}.$$

Then we have $\nu_{cg} = \frac{1}{2} \left(\tau_1 + \sqrt{\tau_1^2 + 4\tau_2} \right) < 1$ and the iterates X_l generated by Alg.2 subject to the restarting conditions in (11) satisfy

$$\|X_l - X\|_F \leq \nu_{cg}^l \|X_0 - X\|_F.$$

When $\kappa_1 = \kappa_2 = 0$, (15) is reduced to (10). On the other hand we have

$$\lim_{\varepsilon_0 \rightarrow 0} (\tau_1 + \tau_2) = 3\kappa_1\kappa_2.$$

So if $\kappa_1\kappa_2 < 1/3$, $\tau_1 + \tau_2$ can be less than one when ε_0 is small. In particular, when $\kappa_1 = 0.1$ and $\kappa_2 = 1$, a sufficient condition for $\tau_1 + \tau_2 < 1$ is $\varepsilon_0 \leq 0.01$.

The proofs of Thms. 2.2 and 2.3 are presented in Secs. 4.2 and 4.3 respectively.

2.2 Initialization and Recovery Guarantees

In Thms. 2.2 and 2.3, we list three conditions to guarantee the convergence of the algorithms. The requirement for \mathcal{P}_Ω to be bounded in (7) and (12) is just an artifact of the sampling model, and it can be satisfied with probability at least $1 - n^{2-2\beta}$ following from Lem. 2.1. The second condition (8) and (13) is a local restricted isometry property which has been established in [13] for the Bernoulli model and in [22, 38] for the sampling with replacement model, and it also plays a key role in nuclear norm minimization for matrix completion. For the sampling with replacement model, Thm. A.2 implies that as long as

$$m \geq C\beta \left(\frac{\mu_0}{\varepsilon_0^2} \right) nr \log(n),$$

$\|\mathcal{P}_T - p^{-1}\mathcal{P}_T\mathcal{P}_\Omega\mathcal{P}_T\| \leq \varepsilon_0$ with probability at least $1 - 2n^{2-2\beta}$. Thus the only issue that remains to be addressed is how to produce an initial guess that is sufficiently close to the measured matrix. We will consider two initialization strategies.

2.2.1 Initialization via One Step Hard Thresholding

A widely used initialization for matrix completion is to set $X_0 = \mathcal{H}_r(p^{-1}\mathcal{P}_\Omega(X))$. The approximation error $\|X_0 - X\|_F$ can be estimated as follows.

Lemma 2.4. *Suppose Ω with $|\Omega| = m$ is a set of indices sampled independently and uniformly with replacement. Let $X_0 = \mathcal{H}_r(p^{-1}\mathcal{P}_\Omega(X))$. Then for all $\beta > 1$*

$$\|X_0 - X\|_F \leq \sqrt{\frac{64\beta\mu_1^2nr^2\log(n)}{3m}} \|X\|$$

with probability at least $1 - 2n^{1-\beta}$ provided $m \geq 6\beta n \log(n)$.

The proof of Lem. 2.4 is presented in Sec. 4.4, which follows from a direct application of Thm. A.3 for the sampling with replacement model. A similar inequality has been established in [29] for $\widehat{X}_0 = \mathcal{H}_r(p^{-1}\widehat{\mathcal{P}_\Omega(X)})$ under the Bernoulli model using random graph theory, where $\widehat{\mathcal{P}_\Omega(X)}$ is a trimmed matrix obtained by setting all the rows and columns of $\mathcal{P}_\Omega(X)$ with too many observed

Algorithm 3 Initialization via Resampled Riemannian Gradient Descent and Trimming

Partition Ω into $L + 1$ equal groups: $\Omega_0, \dots, \Omega_L$; and the size of each group is denoted by \widehat{m} .

Set $Z_0 = \mathcal{H}_r \left(\frac{\widehat{m}}{n^2} \mathcal{P}_{\Omega_0}(X) \right)$.

for $l = 0, \dots, L - 1$ **do**

1. $\widehat{Z}_l = \text{trim}(Z_l)$

2. $Z_{l+1} = \mathcal{H}_r \left(\widehat{Z}_l + \frac{n^2}{\widehat{m}} \mathcal{P}_{\widehat{T}_l} \mathcal{P}_{\Omega_{l+1}} \left(X - \widehat{Z}_l \right) \right)$

end for

Output: $X_0 = Z_L$

Algorithm 4 trim

Input: $Z_l = U_l \Sigma_l V_l^*$

Output: $\widehat{Z}_l = A_l \Sigma_l B_l^*$, where

$$A_l^{(i)} = \frac{U_l^{(i)}}{\|U_l^{(i)}\|} \min \left(\|U_l^{(i)}\|, \sqrt{\frac{\mu_0 r}{n}} \right), \quad B_l^{(i)} = \frac{V_l^{(i)}}{\|V_l^{(i)}\|} \min \left(\|V_l^{(i)}\|, \sqrt{\frac{\mu_0 r}{n}} \right)$$

entries to zero. Moreover, an application of the standard Chernoff bound can further show that with high probability there are no rows or columns of $\mathcal{P}_{\Omega}(X)$ with too many entries [27].

It follows from Lem. 2.4 that the third condition (9) and (14) in Thms. 2.2 and 2.3 can be satisfied with probability at least $1 - 2n^{1-\beta}$ if

$$m \geq C\beta^{3/2} \left(\frac{\mu_1(1 + \varepsilon_0)}{\varepsilon_0} \right) \kappa n^{3/2} r \log^{3/2}(n).$$

Therefore we can establish the following theorem.

Theorem 2.5 (Recovery Guarantee I). *Let $X \in \mathbb{R}^{n \times n}$ be the measured rank r matrix. Suppose Ω with $|\Omega| = m$ is a set of indices sampled independently and uniformly with replacement. Let $X_0 = \mathcal{H}_r(p^{-1} \mathcal{P}_{\Omega}(X))$. Then for all $\beta > 1$, the iterates of the Riemannian gradient descent algorithm (Alg. 1) and the restarted Riemannian conjugate gradient algorithm (Alg. 2, restarting when either the inequality in (11) is violated) with $\kappa_1 \kappa_2 < 1/3$ is guaranteed to converge to X with probability at least $1 - 3n^{2-2\beta} - 2n^{1-\beta}$ provided*

$$m \geq C\beta^{3/2} \max \left\{ \frac{\mu_0}{\varepsilon_0^2}, \frac{\mu_1(1 + \varepsilon_0)}{\varepsilon_0} \right\} \kappa n^{3/2} r \log^{3/2}(n).$$

2.2.2 Initialization via Resampled Riemannian Gradient Descent and Trimming

The sampling complexity in Thm. 2.5 depends on $n^{1.5}$, rather than linearly on n . To attenuate this dependence, we consider a more delicate initialization scheme via the resampled Riemannian gradient descent followed by trimming in each iteration, see Alg. 3. The trimming procedure (Alg. 4) projects the estimate onto the set of μ_0 -incoherent matrices, while the resampling scheme breaks the dependence between the past iterate and the new sampling set. So we can establish the required isometry properties which are needed to prove the linear convergence of the iterates until

they reach an order of $p^{1/2}$ neighborhood around the measured matrix where Thms. 2.2 and 2.3 are activated. *Since the main computational cost in each iteration of Alg. 3 is $O(|\Omega_l|r)$ flops [47], the total computational cost of Alg. 3 is only slightly larger than one iteration of Alg. 1 applied on the full sampling set Ω .* The output of Alg. 3 satisfies the following property.

Lemma 2.6. *Let $X \in \mathbb{R}^{n \times n}$ be the measured rank r matrix. Then for all $\beta > 1$, with probability at least $1 - 2n^{1-\beta} - 4Ln^{2-2\beta}$,*

$$\|X_0 - X\|_F \leq \left(\frac{5}{6}\right)^L \frac{\sigma_{\min}(X)}{256\kappa^2}$$

provided $\hat{m} \geq C\beta \max\{\mu_0, \mu_1^2\} \kappa^6 nr^2 \log(n)$.

The proof of Lem. 2.6 is presented in Sec. 4.5. In Alg. 3, we use fixed stepsize $\frac{n^2}{\hat{m}}$ for ease of exposition, which can be replaced by the adaptive stepsize similar to Alg. 1. Lemma 2.6 implies that if we take

$$L \geq 6 \log \left(\frac{\beta n \log(n)}{24\varepsilon_0} \right),$$

the third condition (9) and (14) in Thms. 2.2 and 2.3 can be satisfied with probability at least $1 - 2n^{1-\beta} - 24 \log \left(\frac{\beta n \log(n)}{24\varepsilon_0} \right) n^{2-2\beta}$. So together with Lem. 2.1, Thms. A.2, 2.2 and 2.3, we can establish the following theorem.

Theorem 2.7 (Recovery Guarantee II). *Let $X \in \mathbb{R}^{n \times n}$ be the measured rank r matrix. Suppose Ω with $|\Omega| = m$ is a set of indices sampled independently and uniformly with replacement. Let X_0 be the output of Alg. 3. Then for all $\beta > 1$, the iterates of the Riemannian gradient descent algorithm (Alg. 1) and the restarted Riemannian conjugate gradient algorithm (Alg. 2, restarting when either the inequality in (11) is violated) with $\kappa_1 \kappa_2 < 1/3$ is guaranteed to converge to X with probability at least $1 - 2n^{1-\beta} - \left(24 \log \left(\frac{\beta n \log(n)}{24\varepsilon_0} \right) + 3\right) n^{2-2\beta}$ provided*

$$m \geq C\beta \max \left\{ \frac{\mu_0}{\varepsilon_0^2}, \mu_1^2 \right\} \kappa^6 nr^2 \log(n) \log \left(\frac{\beta n \log(n)}{24\varepsilon_0} \right).$$

2.3 Related Work

From the pioneer work of Recht, Fazel, and Parrilo [39, 19] and Candès and Recht [13], the low rank matrix reconstruction problem has received intensive investigations from both the theoretical and algorithmic aspects. Recht, Fazel, and Parrilo [39] studied nuclear norm minimization for the low rank matrix recovery problem where each measurement is obtained by taking inner product between the underlying low rank matrix and a dense measurement matrix. They showed that if the restricted isometry constant of the sensing operator is smaller than a numerical constant, then nuclear norm minimization is guaranteed to recover a low rank matrix exactly. Moreover, this condition can be satisfied for certain family of random measurement matrices provided the number of measurements is $O(nr)$ [12]. Many non-convex optimization algorithms have also been studied for low rank matrix recovery based on either the restricted isometry constant of the sensing operator or directly on the random measurement models, see [30, 31, 43, 25, 45, 17, 53, 40, 6, 52, 51] and

references therein. In particular, the Riemannian gradient descent and conjugate gradient descent algorithms for low rank matrix recovery have been studied in [49].

Candès and Recht [13] first studied matrix completion by nuclear norm minimization under the Bernoulli model and incoherent conditions, showing that $O(n^{1.2}r \log(n))$ measurements are sufficient for successful recovery with high probability. This result was subsequently sharpened to $O(nr \log^\alpha(n))$ in [22, 38, 14, 15]. The nuclear norm minimization problem can be solved by either semidefinite programming [46] or iterative soft thresholding algorithms [9]. A gradient descent algorithm on the Grassmannian manifold of low dimensional subspaces was studied in [29]. It was shown that $O(\kappa^6 nr^2)$ number of measurements allows for convergence of the algorithm, which was established by penalizing the incoherence of the estimate in the loss function. The proof in [29] relies on the stationary point analysis and linear convergence rate of the algorithm was not shown. This framework was further extended in [42] to gradient descent algorithms based on the product factorization of low rank matrices. Alternating minimization was studied in [28, 27] with the best known sampling complexity being $O(\kappa^8 nr \log(\frac{n}{\epsilon}))$, where the desired accuracy ϵ was introduced because the algorithm requires a fresh set of measurements in each iteration. *In contrast, Thm. 2.6 only requires resampling for the initialization stage and our algorithms can converge linearly to the desired low rank matrix (within arbitrary precision) for the fixed number of measurements.*

3 Numerical Experiments

In this section, we present empirical observations of the Riemannian gradient descent and conjugate gradient descent algorithms for random tests. We test three algorithms, namely, the Riemannian gradient descent algorithm (Alg. 1), the Riemannian conjugate gradient descent algorithm (Alg. 2), and the Riemannian conjugate gradient algorithm being restarted if one of the conditions in (11) is violated. These three algorithms are abbreviated as RGrad, RCG, and RCG restarted respectively. For RCG restarted, we take $\kappa_1 = 0.1$ and $\kappa_2 = 1$ in (11). All tested algorithms are initialized by one step hard thresholding instead of Alg. 3, as the former has already led to very good performance (see Figs. 1, 2 and 3) and preliminary numerical results didn't find much difference between these two initialization strategies for random simulations. The numerical experiments are conducted on a Mac Pro laptop with 2.5GHz quad-core Intel Core i7 CPUs and 16 GB memory and executed from Matlab 2014b.

3.1 Empirical Phase Transition

An important question in matrix completion is how many of measurements are needed in order for an algorithm to be able to reliably reconstruct a low rank matrix. We investigate the recovery abilities of the tested algorithms in the framework of phase transition, which compares the number of measurements, m , the size of an $n \times n$ matrix, n^2 , and the minimum number of measurements needed to recover an $n \times n$ rank r matrix², $(2n - r)r$, through the *undersampling* and *oversampling* ratios

$$p = \frac{m}{n^2}, \quad q = \frac{(2n - r)r}{m}.$$

The unit square $(p, q) \in [0, 1]^2$ defines a phase transition space. Given a triple (n, m, r) corresponding to a fixed pair of (p, q) , we conduct test on *ten* random instances. The test rank r matrix

²All $n \times n$ rank r matrices form a smooth embedded manifold of dimension $(2n - r)r$ in the ambient space [47].




Figure 1: Empirical phase transition curves (a) RGrad, (b) RCG and (c) RCG restarted when $n = 800$. Horizontal axis $p = m/n^2$ and vertical axis $q = (2n - r)r/m$. White denotes successful recovery in all ten random tests, and black denotes failure in all tests.

is formed as the product of two random rank r matrices; that is $X = LR$, where $L \in \mathbb{R}^{n \times r}$ and $R \in \mathbb{R}^{r \times n}$ with the entries of L and R sampled from the standard Gaussian distribution. The measurement vector $\mathcal{P}_\Omega(X)$ is obtained by sampling m entries of X uniformly at random. An algorithm is considered to have successfully recovered X if it returns a matrix X_l which satisfies

$$\frac{\|X_l - X\|_F}{\|X\|_F} \leq 10^{-2}.$$

The tests are conducted with $n = 800$ and p taking 18 equispaced values from 0.1 to 0.95.

The probabilities of successful recovery for the three tested algorithms RGrad, RCG and RCG restarted are displayed in Fig. 1. In the figure, white color indicates that the algorithm can recover all of the ten random test matrices while black color implies the algorithm fails to recover each of the randomly drawn matrices. For each tested algorithm, a clear phase transition occurs when q is greater than 0.9 for all the values of p . This implies that all the three tested algorithms are able to recover a rank r matrix from $m = C \cdot (2n - r)r$ number of measurements with C being slightly larger than 1. In addition, Fig. 1c also shows the effectiveness of our restarting conditions for the Riemannian conjugate gradient descent algorithm.

3.2 Computation Efficiency

We compare the computational efficiency of RGrad, RCG, RCG restarted by conducting random tests on matrices of 8000×8000 and rank 100. The tests are conducted with two different over-sampling ratios $1/q \in \{2, 3\}$. The algorithms are terminated when the relative residual is less than 10^{-9} . The relative residual plotted against the number of iterations and the average recovery time are presented in Fig. 2. First it can be observed that the convergence curves for RCG and RCG restarted are almost indistinguishable, differing only in one or two iterations. This again shows the effectiveness of our restarting conditions. A close look at the computational results reveals that restarting usually occurs in the first few iterations for RCG restarted. Moreover, RCG and RCG restarted are sufficiently faster than RGrad both in terms of the number of iterations and in terms of the average computation time. It takes less number of iterations for each algorithm to converge below the desired accuracy when the number of measurements increases.




Figure 2: Relative residual (mean and standard deviation over ten random tests) as function of number of iterations for $n = 8000$, $r = 100$, $1/q = 2$ (a) and $1/q = 3$ (b). The values after each algorithm are the average computational time (seconds) for convergence.




Figure 3: Performance of (a) RGrad, (b) RCG and (c) RCG restarted under different SNR.

3.3 Robustness to Additive Noise

Following the test set-up in last subsection, we further explore performance of the algorithms under the measurements with additive noise. Tests with additive noise have the sampled entries $\mathcal{P}_\Omega(X)$ corrupted by the vector

$$e = \sigma \cdot \|\mathcal{P}_\Omega(X)\|_F \cdot \frac{w}{\|w\|_2},$$

where the entries of w are i.i.d standard Gaussian random variables and σ is referred to as noise level. We conduct tests with *nine* different values of σ from 10^{-4} to 1; and for each σ , *ten* random tests are conducted. The average relative reconstruction error in dB plotted against the signal-to-noise ratio (SNR) is presented in Fig. 3 for RGrad, RCG and RCG restarted respectively. The plots clearly show the desirable linear scaling between the noise levels and the relative errors for all the three tested algorithms. The relative error decreases as the number of measurements increases.

4 Proofs

In this section, we prove the main results in Sec. 2. We first list several technical lemmas which are of independent interest. The proofs of these lemmas are presented in Appendix B.

4.1 Technical Lemmas

Lemma 4.1 (Bounds for Projections). *Let $X_l = U_l \Sigma_l V_l^*$ be a rank r matrix, and T_l be the tangent space of the rank r matrix manifold at X_l . Let $X = U \Sigma V^*$ be another r matrix, and T be the corresponding tangent space. Then*

$$\begin{aligned} \|U_l U_l^* - U U^*\| &\leq \frac{\|X_l - X\|_F}{\sigma_{\min}(X)}, & \|V_l V_l^* - V V^*\| &\leq \frac{\|X_l - X\|_F}{\sigma_{\min}(X)}, \\ \|U_l U_l^* - U U^*\|_F &\leq \frac{\sqrt{2} \|X_l - X\|_F}{\sigma_{\min}(X)}, & \|V_l V_l^* - V V^*\|_F &\leq \frac{\sqrt{2} \|X_l - X\|_F}{\sigma_{\min}(X)}, \\ \|(\mathcal{I} - \mathcal{P}_{T_l}) X\|_F &\leq \frac{\|X_l - X\|_F^2}{\sigma_{\min}(X)}, & \|\mathcal{P}_{T_l} - \mathcal{P}_T\| &\leq \frac{2 \|X_l - X\|_F}{\sigma_{\min}(X)}. \end{aligned} \quad (16)$$

Lemma 4.2 (Restricted Isometry Property in a Local Neighborhood). *Assume*

$$\|\mathcal{P}_\Omega\| \leq \frac{8}{3} \beta \log(n), \quad \|\mathcal{P}_T - p^{-1} \mathcal{P}_T \mathcal{P}_\Omega \mathcal{P}_T\| \leq \varepsilon_0, \quad \text{and} \quad \frac{\|X_l - X\|_F}{\sigma_{\min}(X)} \leq \frac{3p^{1/2} \varepsilon_0}{16\beta \log(n)(1 + \varepsilon_0)}$$

for some $0 < \varepsilon_0 < 1$ and $\beta > 1$. Then

$$\|\mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq \frac{8}{3} \beta \log(n)(1 + \varepsilon_0) p^{1/2}, \quad (17)$$

$$\|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq 4\varepsilon_0. \quad (18)$$

Lemma 4.3 (Asymmetric Restricted Isometry Property). *Let $X_l = U_l \Sigma_l V_l^*$ and $X = U \Sigma V^*$ be two fixed rank r matrices. Assume*

$$\begin{aligned} \|\mathcal{P}_{U_l}(e_i)\| &\leq \sqrt{\frac{\mu r}{n}} \quad \text{and} \quad \|\mathcal{P}_{V_l}(e_j)\| \leq \sqrt{\frac{\mu r}{n}}, \\ \|\mathcal{P}_U(e_i)\| &\leq \sqrt{\frac{\mu r}{n}} \quad \text{and} \quad \|\mathcal{P}_V(e_j)\| \leq \sqrt{\frac{\mu r}{n}} \end{aligned}$$

for $1 \leq i, j \leq n$. Suppose Ω with $|\Omega| = m$ is a set of indices sampled independently and uniformly with replacement. Then for any $\beta > 1$,

$$\left\| \frac{n^2}{m} \mathcal{P}_{T_l} \mathcal{P}_\Omega (\mathcal{P}_U - \mathcal{P}_{U_l}) - \mathcal{P}_{T_l} (\mathcal{P}_U - \mathcal{P}_{U_l}) \right\| \leq \sqrt{\frac{48\beta \mu n r \log(n)}{m}}$$

with probability at least $1 - 2n^{2-2\beta}$ provided $m \geq 15\beta \mu n r \log(n)$.

Lemma 4.4 (A Recursive Relationship). *Let ρ_1 , ρ_2 and γ be positive constants satisfying $\rho_2 \geq \rho_1$. Define*

$$\tau_1 = \rho_1 + \gamma, \quad \tau_2 = (\rho_2 - \rho_1)\gamma, \quad \text{and} \quad \nu = \frac{1}{2} \left(\tau_1 + \sqrt{\tau_1^2 + 4\tau_2} \right).$$

Let $\{c_l\}_{l \geq 0}$ be a non-negative sequence satisfying $c_1 \leq \nu c_0$ and

$$c_{l+1} \leq \rho_1 c_l + \rho_2 \sum_{j=0}^{l-1} \gamma^{l-j} c_j, \quad \forall l \geq 1.$$

Then if $\tau_1 + \tau_2 < 1$, we have $\nu < 1$ and

$$c_{l+1} \leq \nu^{l+1} c_0. \quad (19)$$

Lemma 4.5 (Chordal and Projection Distances). *Let $U_l, U \in \mathbb{R}^{n \times r}$ be two orthogonal matrices. Then there exists a $r \times r$ unitary matrix Q such that*

$$\|U_l - UQ\|_F \leq \|U_l U_l^* - U U^*\|_F.$$

4.2 Proof of Theorem 2.2

We start with a lemma which bounds the search stepsize α_l in Alg. 1.

Lemma 4.6. *Assume $\|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq 4\varepsilon_0$. Then the stepsize α_l in Alg. 1 can be bounded as*

$$\frac{1}{(1 + 4\varepsilon_0)p} \leq \alpha_l = \frac{\|\mathcal{P}_{T_l}(G_l)\|_F^2}{\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(G_l) \rangle} \leq \frac{1}{(1 - 4\varepsilon_0)p}.$$

Proof. First the assumption $\|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq 4\varepsilon_0$ implies

$$\|\mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq p \|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| + p \|\mathcal{P}_{T_l}\| \leq (1 + 4\varepsilon_0)p.$$

On one hand,

$$\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(G_l) \rangle = \langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}(G_l) \rangle \leq (1 + 4\varepsilon_0)p \|\mathcal{P}_{T_l}(G_l)\|_F^2,$$

which implies the lower bound of α_l .

On the other hand,

$$\begin{aligned} \|\mathcal{P}_{T_l}(G_l)\|_F^2 &= \langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(G_l) - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}(G_l) \rangle + \langle \mathcal{P}_{T_l}(G_l), p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}(G_l) \rangle \\ &\leq 4\varepsilon_0 \|\mathcal{P}_{T_l}(G_l)\|_F^2 + p^{-1} \langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(G_l) \rangle, \end{aligned} \quad (20)$$

from which we can obtain the upper bound of α_l by multiplying p on both sides of the above inequality followed by the rearrangement. \square

The next lemma bounds the local isometry of \mathcal{P}_Ω when the adaptive stepsize α_l is used.

Lemma 4.7. *Assume $\|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq 4\varepsilon_0$ and α_l can be bounded as in Lem. 4.6. Then the spectral norm of $\mathcal{P}_{T_l} - \alpha_l \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}$ can be bounded as*

$$\|\mathcal{P}_{T_l} - \alpha_l \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq \frac{8\varepsilon_0}{1 - 4\varepsilon_0}.$$

Proof. The lemma follows from direct calculations

$$\begin{aligned}\|\mathcal{P}_{T_l} - \alpha_l \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| &\leq \|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| + |\alpha_l - p^{-1}| \|\mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \\ &\leq 4\varepsilon_0 + \frac{4\varepsilon_0}{(1-4\varepsilon_0)p} (1+4\varepsilon_0)p \\ &= \frac{8\varepsilon_0}{1-4\varepsilon_0},\end{aligned}$$

where the second inequality follows from the assumption and Lem. 4.6. \square

Proof of Theorem 2.2. We first assume that in the l -th iteration X_l satisfies

$$\frac{\|X_l - X\|_F}{\sigma_{\min}(X)} \leq \frac{3p^{1/2}\varepsilon_0}{16\beta \log(n)(1+\varepsilon_0)} \quad (21)$$

So together with the first two assumptions (7) and (8) of Thm. 2.2, the application of Lem. 4.2 gives

$$\|\mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq \frac{8}{3}\beta \log(n)(1+\varepsilon_0)p^{1/2} \quad (22)$$

and

$$\|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| \leq 4\varepsilon_0, \quad (23)$$

which means the assumptions of Lems. 4.6 and 4.7 are satisfied.

Recall that $W_l = X_l + \alpha_l \mathcal{P}_{T_l}(G_l)$ in Alg. 1. The proof begins with the following inequality

$$\|X_{l+1} - X\|_F \leq \|X_{l+1} - W_l\|_F + \|W_l - X\|_F \leq 2\|W_l - X\|_F,$$

where the second inequality follows from the fact that X_{l+1} is the best rank r approximation of W_l . Plugging in $W_l = X_l + \alpha_l \mathcal{P}_{T_l}(G_l)$ gives

$$\begin{aligned}\|X_{l+1} - X\|_F &\leq 2\|X_l + \alpha_l \mathcal{P}_{T_l}(G_l) - X\|_F \\ &= 2\|X_l - X - \alpha_l \mathcal{P}_{T_l} \mathcal{P}_\Omega (X_l - X)\|_F \\ &\leq 2\|(\mathcal{P}_{T_l} - \alpha_l \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l})(X_l - X)\|_F \\ &\quad + 2\|(\mathcal{I} - \mathcal{P}_{T_l})(X_l - X)\|_F \\ &\quad + 2|\alpha_l| \|\mathcal{P}_{T_l} \mathcal{P}_\Omega (\mathcal{I} - \mathcal{P}_{T_l})(X_l - X)\|_F \\ &:= I_1 + I_2 + I_3.\end{aligned} \quad (24)$$

First we have

$$I_1 \leq \frac{16\varepsilon_0}{1-4\varepsilon_0} \|X_l - X\|_F$$

and

$$I_2 \leq \frac{2\|X_l - X\|_F^2}{\sigma_{\min}(X)} \leq \frac{3p^{1/2}\varepsilon_0}{8\beta \log(n)(1+\varepsilon_0)} \|X_l - X\|_F \leq \frac{\varepsilon_0}{1-4\varepsilon_0} \|X_l - X\|_F$$

which follows from Lems. 4.7 and 4.1 respectively. The third term I_3 can be bounded as follows.

$$\begin{aligned}
I_3 &\leq 2|\alpha_l| \|\mathcal{P}_\Omega \mathcal{P}_{T_l}\| \|(\mathcal{I} - \mathcal{P}_{T_l})(X)\|_F \\
&\leq \frac{2}{(1-4\varepsilon_0)p} \frac{8}{3} \beta \log(n) (1+\varepsilon_0) p^{1/2} \frac{\|X_l - X\|_F^2}{\sigma_{\min}(X)} \\
&\leq \frac{\varepsilon_0}{1-4\varepsilon_0} \|X_l - X\|_F,
\end{aligned}$$

where the second inequality follows from (22), and Lems. 4.6 and 4.1, and the third inequality follows from (21).

So inserting the bounds for I_1 , I_2 and I_3 into (24) gives

$$\|X_{l+1} - X\|_F \leq \nu_g \|X_l - X\|_F, \quad (25)$$

where

$$\nu_g = \frac{18\varepsilon_0}{1-4\varepsilon_0} < 1$$

by the assumption of Thm. 2.2.

It only remains to verify (21). By the assumption of Thm. 2.2, (21) is valid for $l = 0$. Since $\|X_l - X\|_F$ is a contractive sequence following from (25), (21) is valid for all $l \geq 0$ by induction. \square

4.3 Proof of Theorem 2.3

We first estimate α_l and β_l in Alg. 2 with the restarting conditions in (11).

Lemma 4.8. *Assume $\|\mathcal{P}_{T_l} - p^{-1}\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}\| \leq 4\varepsilon_0$. When restarting occurs, $\beta_l = 0$ and α_l can be bounded as in Lem. 4.6; otherwise*

$$|\beta_l| \leq \varepsilon_\beta \quad \text{and} \quad |\alpha_l \cdot p - 1| \leq \varepsilon_\alpha,$$

where

$$\varepsilon_\beta = \frac{4\kappa_2\varepsilon_0}{1-4\varepsilon_0} + \frac{\kappa_1\kappa_2}{1-4\varepsilon_0}, \quad \varepsilon_\alpha = \frac{4\varepsilon_0}{(1-4\varepsilon_0) - \kappa_1(1+4\varepsilon_0)}.$$

Note that the bounds above are also valid even when restarting occurs since $\kappa_1 \geq 0$ and $\kappa_2 \geq 0$.

Proof. The orthogonalization weight β_l can be bounded as follows

$$\begin{aligned}
|\beta_l| &= \left| \frac{\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle}{\langle \mathcal{P}_{T_l}(P_{l-1}), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle} \right| \\
&\leq \left| \frac{\langle \mathcal{P}_{T_l}(G_l), (\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l} - p\mathcal{P}_{T_l})(P_{l-1}) \rangle}{\langle \mathcal{P}_{T_l}(P_{l-1}), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle} \right| + \left| \frac{p \langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle}{\langle \mathcal{P}_{T_l}(P_{l-1}), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle} \right| \\
&\leq \frac{4\varepsilon_0 p}{(1-4\varepsilon_0)p} \frac{\|\mathcal{P}_{T_l}(G_l)\|_F}{\|\mathcal{P}_{T_l}(P_{l-1})\|_F} + \frac{p}{(1-4\varepsilon_0)p} \frac{|\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle|}{\|\mathcal{P}_{T_l}(P_{l-1})\|_F^2} \\
&\leq \frac{4\kappa_2\varepsilon_0}{1-4\varepsilon_0} + \frac{1}{1-4\varepsilon_0} \frac{|\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle|}{\|\mathcal{P}_{T_l}(G_l)\|_F \|\mathcal{P}_{T_l}(P_{l-1})\|_F} \frac{\|\mathcal{P}_{T_l}(G_l)\|_F}{\|\mathcal{P}_{T_l}(P_{l-1})\|_F}
\end{aligned}$$

$$\leq \frac{4\kappa_2\varepsilon_0}{1-4\varepsilon_0} + \frac{\kappa_1\kappa_2}{1-4\varepsilon_0},$$

where in the third line we used $\langle \mathcal{P}_{T_l}(P_{l-1}), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle \geq (1-4\varepsilon_0)p \|\mathcal{P}_{T_l}(P_{l-1})\|_F^2$ deducted from (20), and the last two inequalities follow from the restarting conditions.

To bound α_l , we first need to bound $\|\mathcal{P}_{T_l}(G_l)\|_F$ in terms of $\|\mathcal{P}_{T_l}(P_l)\|_F$. Note that

$$\begin{aligned} |\beta_l \langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle| &= \left| \frac{\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle}{\langle \mathcal{P}_{T_l}(P_{l-1}), \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_{l-1}) \rangle} \langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle \right| \\ &\leq \frac{(1+4\varepsilon_0)p \|\mathcal{P}_{T_l}(G_l)\|_F}{(1-4\varepsilon_0)p \|\mathcal{P}_{T_l}(P_{l-1})\|_F} |\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle| \\ &\leq \frac{1+4\varepsilon_0}{1-4\varepsilon_0} \|\mathcal{P}_{T_l}(G_l)\|_F^2 \frac{|\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle|}{\|\mathcal{P}_{T_l}(G_l)\|_F \|\mathcal{P}_{T_l}(P_{l-1})\|_F} \\ &\leq \frac{\kappa_1(1+4\varepsilon_0)}{1-4\varepsilon_0} \|\mathcal{P}_{T_l}(G_l)\|_F^2. \end{aligned}$$

So

$$\begin{aligned} |\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_{T_l}(G_l) \rangle| &= |\langle \mathcal{P}_{T_l}(G_l) + \beta_l \mathcal{P}_{T_l}(P_{l-1}), \mathcal{P}_{T_l}(G_l) \rangle| \\ &\geq \|\mathcal{P}_{T_l}(G_l)\|_F^2 - |\beta_l \langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_{l-1}) \rangle| \\ &\geq \left(1 - \frac{\kappa_1(1+4\varepsilon_0)}{1-4\varepsilon_0}\right) \|\mathcal{P}_{T_l}(G_l)\|_F^2. \end{aligned}$$

The application of the Cauchy-Schwarz inequality gives

$$\|\mathcal{P}_{T_l}(G_l)\|_F \leq \frac{1}{1 - \frac{\kappa_1(1+4\varepsilon_0)}{1-4\varepsilon_0}} \|\mathcal{P}_{T_l}(P_l)\|_F.$$

Since α_l can be rewritten as

$$\begin{aligned} \alpha_l &= \frac{\langle \mathcal{P}_{T_l}(G_l), \mathcal{P}_{T_l}(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle} \\ &= \frac{\langle \mathcal{P}_{T_l}(G_l), p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle} + \frac{\langle \mathcal{P}_{T_l}(G_l), (\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l})(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle} \\ &= \frac{\langle \mathcal{P}_{T_l}(G_l), p^{-1} \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle} + \frac{\langle \mathcal{P}_{T_l}(G_l), (\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l})(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle} \\ &= p^{-1} + \frac{\langle \mathcal{P}_{T_l}(G_l), (\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l})(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle}, \end{aligned}$$

it can be bounded as follows

$$\begin{aligned} |\alpha_l \cdot p - 1| &\leq p \left| \frac{\langle \mathcal{P}_{T_l}(G_l), (\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l})(P_l) \rangle}{\langle \mathcal{P}_{T_l}(P_l), \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}(P_l) \rangle} \right| \\ &\leq \frac{4\varepsilon_0 p}{(1-4\varepsilon_0)p} \frac{\|\mathcal{P}_{T_l}(G_l)\|_F}{\|\mathcal{P}_{T_l}(P_l)\|_F} \\ &\leq \frac{4\varepsilon_0}{(1-4\varepsilon_0) - \kappa_1(1+4\varepsilon_0)}, \end{aligned}$$

which completes the proof. \square

The following lemma bounds the local isometry of \mathcal{P}_Ω when the adaptive stepsize α_l is used.

Lemma 4.9. *Assume $\|\mathcal{P}_{T_l} - p^{-1}\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}\| \leq 4\varepsilon_0$, and α_l is bounded as in Lem. 4.8. Then the spectral norm of $\mathcal{P}_{T_l} - \alpha_l\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}$ can be bounded as*

$$\|\mathcal{P}_{T_l} - \alpha_l\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}\| \leq 4\varepsilon_0 + \varepsilon_\alpha(1 + 4\varepsilon_0).$$

Proof. Direct calculations give

$$\begin{aligned} \|\mathcal{P}_{T_l} - \alpha_l\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}\| &\leq \|\mathcal{P}_{T_l} - p^{-1}\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}\| + |\alpha_l - p^{-1}| \|\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}\| \\ &\leq 4\varepsilon_0 + |\alpha \cdot p - 1| \|p^{-1}\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l}\| \\ &\leq 4\varepsilon_0 + \varepsilon_\alpha(1 + 4\varepsilon_0), \end{aligned}$$

which completes the proof. \square

Proof of Theorem 2.3. We first assume that for all $j \leq l$

$$\frac{\|X_j - X\|_F}{\sigma_{\min}(X)} \leq \frac{3p^{1/2}\varepsilon_0}{16\beta \log(n)(1 + \varepsilon_0)}. \quad (26)$$

So together with the first two assumptions (7) and (8) of Thm. 2.3, the application of Lem. 4.2 gives

$$\|\mathcal{P}_\Omega\mathcal{P}_{T_j}\| \leq \frac{8}{3}\beta \log(n)(1 + \varepsilon_0)p^{1/2} \quad (27)$$

and

$$\|\mathcal{P}_{T_j} - p^{-1}\mathcal{P}_{T_j}\mathcal{P}_\Omega\mathcal{P}_{T_j}\| \leq 4\varepsilon_0, \quad (28)$$

which means the assumptions of Lems. 4.8 and 4.9 are satisfied for all $j \leq l$.

Analogous to (24), we have

$$\begin{aligned} \|X_{l+1} - X\|_F &\leq 2\|X_l + \alpha_l P_l - X\|_F \\ &= 2\|X_l - X - \alpha_l\mathcal{P}_{T_l}\mathcal{P}_\Omega(X_l - X) + \alpha_l\beta_l\mathcal{P}_{T_l}(P_{l-1})\|_F \\ &\leq 2\|(\mathcal{P}_{T_l} - \alpha_l\mathcal{P}_{T_l}\mathcal{P}_\Omega\mathcal{P}_{T_l})(X_l - X)\|_F \\ &\quad + 2\|(\mathcal{I} - \mathcal{P}_{T_l})(X_l - X)\|_F \\ &\quad + 2|\alpha_l| \|\mathcal{P}_{T_l}\mathcal{P}_\Omega(\mathcal{I} - \mathcal{P}_{T_l})(X_l - X)\|_F \\ &\quad + 2|\alpha_l|\beta_l \|\mathcal{P}_{T_l}(P_{l-1})\|_F \\ &:= I_4 + I_5 + I_6 + I_7. \end{aligned} \quad (29)$$

Following the argument for the proof of Thm. 2.2, the first three terms can be similarly bounded as follows

$$\begin{aligned} I_4 &\leq (8\varepsilon_0 + 2\varepsilon_\alpha(1 + 4\varepsilon_0)) \|X_l - X\|_F, \\ I_5 &\leq \frac{\varepsilon_0}{1 - 4\varepsilon_0} \|X_l - X\|_F \leq (1 + \varepsilon_\alpha)\varepsilon_0 \|X_l - X\|_F, \\ I_6 &\leq (1 + \varepsilon_\alpha)\varepsilon_0 \|X_l - X\|_F. \end{aligned}$$

To bound I_7 , first note that $\beta_l \mathcal{P}_{T_l}(P_{l-1})$ can be expressed in terms of all the previous gradients

$$\beta_l \mathcal{P}_{T_l}(P_{l-1}) = \sum_{j=0}^{l-1} \left(\prod_{i=j+1}^l \beta_i \right) \left(\prod_{k=j}^{l-1} \mathcal{P}_{T_k} \right) (G_j).$$

So

$$\begin{aligned} I_7 &\leq 2|\alpha_l| \sum_{j=0}^{l-1} \left(\prod_{i=j+1}^l \beta_i \right) \left\| \left(\prod_{k=j}^{l-1} \mathcal{P}_{T_k} \right) (G_j) \right\|_F \\ &\leq 2|\alpha_l| \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} \|\mathcal{P}_{T_j}(G_j)\|_F \\ &= 2|\alpha_l| \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} \|\mathcal{P}_{T_j} \mathcal{P}_\Omega(X_j - X)\|_F \\ &\leq 2|\alpha_l| \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} \left(\|\mathcal{P}_{T_j} \mathcal{P}_\Omega \mathcal{P}_{T_j}(X_j - X)\|_F + \|\mathcal{P}_{T_j} \mathcal{P}_\Omega(\mathcal{I} - \mathcal{P}_{T_j})(X_j - X)\|_F \right) \\ &\leq 2|\alpha_l| \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} \left(\|\mathcal{P}_{T_j} \mathcal{P}_\Omega \mathcal{P}_{T_j}(X_j - X)\|_F + \|\mathcal{P}_\Omega \mathcal{P}_{T_j}\| \frac{\|X_j - X\|_F^2}{\sigma_{\min}(X)} \right) \\ &\leq \frac{1 + \varepsilon_\alpha}{p} \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} (2(1 + 4\varepsilon_0)p \|X_j - X\|_F + p\varepsilon_0 \|X_j - X\|_F) \\ &= (1 + \varepsilon_\alpha) \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} (2(1 + 4\varepsilon_0) + \varepsilon_0) \|X_j - X\|_F, \end{aligned}$$

where the second inequality follows from the bound for β_i in Lem. 4.8 and the fact

$$\left\| \left(\prod_{k=j}^{l-1} \mathcal{P}_{T_k} \right) (G_j) \right\|_F \leq \|\mathcal{P}_{T_j}(G_j)\|_F,$$

the fourth inequality follows from Lem. 4.1, and the last inequality follows from the bound for α_l in Lem. 4.8, (26), (27), and (28).

Inserting the bounds for I_4 , I_5 , I_6 and I_7 into (29) gives

$$\begin{aligned} \|X_{l+1} - X\|_F &\leq (10\varepsilon_0 + 2\varepsilon_\alpha(1 + 5\varepsilon_0)) \|X_l - X\|_F + (1 + \varepsilon_\alpha)(2(1 + 4\varepsilon_0) + \varepsilon_0) \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} \|X_j - X\|_F \\ &\leq (10\varepsilon_0 + 2\varepsilon_\alpha(1 + 5\varepsilon_0)) \|X_l - X\|_F + (2 + 10\varepsilon_0 + 2\varepsilon_\alpha(1 + 5\varepsilon_0)) \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} \|X_j - X\|_F \\ &:= \rho_1 \|X_l - X\|_F + \rho_2 \sum_{j=0}^{l-1} \varepsilon_\beta^{l-j} \|X_j - X\|_F \end{aligned}$$

Define

$$\begin{aligned}\tau_1 &= \rho_1 + \varepsilon_\beta = \frac{18\varepsilon_0 - 10\varepsilon_0\kappa_1(1 + 4\varepsilon_0)}{(1 - 4\varepsilon_0) - \kappa_1(1 + 4\varepsilon_0)} + \frac{4\kappa_2\varepsilon_0 + \kappa_1\kappa_2}{1 - 4\varepsilon_0}, \\ \tau_2 &= (\rho_2 - \rho_1)\varepsilon_\beta = \frac{8\kappa_2\varepsilon_0 + 2\kappa_1\kappa_2}{1 - 4\varepsilon_0}, \\ \nu_{cg} &= \frac{1}{2} \left(\tau_1 + \sqrt{\tau_1^2 + \tau_2} \right).\end{aligned}$$

When $l = 0$, Alg. 2 is the same as Alg. 1, so

$$\|X_1 - X\|_F \leq \frac{18\varepsilon_0}{1 - 4\varepsilon_0} \|X_0 - X\|_F \leq \nu_{cg} \|X_0 - X\|_F,$$

following from (25). Therefore Lem. 4.4 implies if $\tau_1 + \tau_2 < 1$, we have $\nu_{cg} < 1$ and

$$\|X_{l+1} - X\|_F \leq \nu_{cg}^{l+1} \|X_0 - X\|_F.$$

Similarly, (26) is only required at $l = 0$ since we have a contractive sequence. \square

4.4 Proof of Lemma 2.4

Let $W_0 = p^{-1}\mathcal{P}_\Omega(X)$. First we have

$$\|X_0 - X\| \leq \|X_0 - W_0\| + \|W_0 - X\| \leq 2\|W_0 - X\| \leq 2\sqrt{\frac{8\beta n^3 \log(n)}{3m}} \|X\|_\infty,$$

where the last inequality holds true with probability at least $1 - 2n^{1-\beta}$ following from Thm. A.3. Consequently,

$$\|X_0 - X\|_F \leq \sqrt{2r} \|X_0 - X\| \leq \sqrt{\frac{64\beta n^3 r \log(n)}{3m}} \|X\|_\infty \leq \sqrt{\frac{64\beta \mu_1^2 n r^2 \log(n)}{3m}} \|X\|,$$

which completes the proof.

4.5 Proof of Lemma 2.6

We first present a lemma which says the matrix returned by Alg. 4 is incoherent and its approximation error can be bounded by the approximation error of the input matrix.

Lemma 4.10. *Let $Z_l = U_l \Sigma_l V_l^*$ be a rank r matrix such that*

$$\|Z_l - X\|_F \leq \frac{\sigma_{\min}(X)}{10\sqrt{2}}.$$

Then the matrix \widehat{Z}_l returned by Alg. 4 satisfies

$$\left\| \mathcal{P}_{\widehat{U}_l}(e_i) \right\| \leq \frac{10}{9} \sqrt{\frac{\mu_0 r}{n}} \quad \text{and} \quad \left\| \mathcal{P}_{\widehat{V}_l}(e_j) \right\| \leq \frac{10}{9} \sqrt{\frac{\mu_0 r}{n}}$$

for $1 \leq i, j \leq n$, where the columns of \widehat{U}_l and \widehat{V}_l are the right and left singular vectors of \widehat{Z}_l respectively. Furthermore,

$$\left\| \widehat{Z}_l - X \right\|_F \leq 8\kappa \|Z_l - X\|_F.$$

Proof. For simplicity, let $d = \|Z_l - X\|_F$. By Lem. 4.1, we have

$$\|U_l U_l^* - U U^*\|_F \leq \frac{\sqrt{2}d}{\sigma_{\min}(X)} \quad \text{and} \quad \|V_l V_l^* - V V^*\|_F \leq \frac{\sqrt{2}d}{\sigma_{\min}(X)}.$$

This together with Lem. 4.5 implies that there exist two unitary matrices $Q_u \in \mathbb{R}^{r \times r}$ and $Q_v \in \mathbb{R}^{r \times r}$ such that

$$\|U_l - U Q_u\|_F \leq \frac{\sqrt{2}d}{\sigma_{\min}(X)} \quad \text{and} \quad \|V_l - V Q_v\|_F \leq \frac{\sqrt{2}d}{\sigma_{\min}(X)}.$$

It follows that

$$\begin{aligned} \|\Sigma_l - Q_u^* \Sigma Q_v\|_F &= \|U_l^* Z_l V_l - (U Q_u)^* X (V Q_v)\|_F \\ &\leq \|U_l^* Z_l V_l - (U Q_u)^* Z_l V_l\|_F + \|(U Q_u)^* Z_l V_l - (U Q_u)^* X V_l\|_F \\ &\quad + \|(U Q_u)^* X V_l - (U Q_u)^* X (V Q_v)\|_F \\ &\leq \|U_l - U Q_u\|_F \|Z_l\| + \|Z_l - X\|_F + \|X\| \|V_l - V Q_v\|_F \\ &\leq \frac{\sqrt{2}d(d + \sigma_{\max}(X))}{\sigma_{\min}(X)} + d + \frac{\sqrt{2}d\sigma_{\max}(X)}{\sigma_{\min}(X)} \\ &\leq 4\kappa d, \end{aligned} \tag{30}$$

where κ is the condition number of X , and we have used the assumption $d \leq \sigma_{\min}(X)/10\sqrt{2} \leq \sigma_{\max}(X)/10\sqrt{2}$ and the fact

$$\|Z_l\| \leq \|X\| + \|Z_l - X\| \leq \sigma_{\max}(X) + d.$$

in the last two inequalities.

Recall that A_l and B_l in Alg. 4 are defined as

$$A_l^{(i)} = \frac{U_l^{(i)}}{\|U_l^{(i)}\|} \min\left(\|U_l^{(i)}\|, \sqrt{\frac{\mu_0 r}{n}}\right), \quad \text{and} \quad B_l^{(i)} = \frac{V_l^{(i)}}{\|V_l^{(i)}\|} \min\left(\|V_l^{(i)}\|, \sqrt{\frac{\mu_0 r}{n}}\right).$$

Because

$$\|(U Q_u)^{(i)}\| \leq \sqrt{\frac{\mu_0 r}{n}} \quad \text{and} \quad \|(V Q_v)^{(i)}\| \leq \sqrt{\frac{\mu_0 r}{n}},$$

we have

$$\|A_l^{(i)} - (U Q_u)^{(i)}\| \leq \|U_l^{(i)} - (U Q_u)^{(i)}\| \quad \text{and} \quad \|B_l^{(i)} - (V Q_v)^{(i)}\| \leq \|V_l^{(i)} - (V Q_v)^{(i)}\|.$$

Therefore

$$\|A_l - U Q_u\|_F \leq \|U_l - U Q_u\|_F \leq \frac{\sqrt{2}d}{\sigma_{\min}(X)}, \tag{31}$$

$$\|B_l - V Q_v\|_F \leq \|V_l - V Q_v\|_F \leq \frac{\sqrt{2}d}{\sigma_{\min}(X)}. \tag{32}$$

Since $\widehat{Z}_l = A_l \Sigma_l B_l^*$ by Alg. 4, we have

$$\begin{aligned}
\left\| \widehat{Z}_l - X \right\|_F &= \|A_l \Sigma_l B_l^* - (UQ_u)(Q_u^* \Sigma Q_v)(VQ_v)^*\|_F \\
&\leq \|A_l \Sigma_l B_l^* - (UQ_u) \Sigma_l B_l^*\|_F + \|(UQ_u) \Sigma_l B_l^* - (UQ_u)(Q_u^* \Sigma Q_v) B_l^*\|_F \\
&\quad + \|(UQ_u)(Q_u^* \Sigma Q_v) B_l^* - (UQ_u)(Q_u^* \Sigma Q_v)(VQ_v)^*\|_F \\
&\leq \|A_l - UQ_u\|_F \|\Sigma_l\| \|B_l\| + \|\Sigma_l - Q_u^* \Sigma Q_v\|_F \|B_l\| + \|\Sigma\| \|B_l - VQ_v\|_F \\
&\leq \frac{\sqrt{2}d}{\sigma_{\min}(X)} (\sigma_{\max}(X) + d) \left(\frac{\sqrt{2}d}{\sigma_{\min}(X)} + 1 \right) + 4\kappa d \left(\frac{\sqrt{2}d}{\sigma_{\min}(X)} + 1 \right) + \sigma_{\max}(X) \frac{\sqrt{2}d}{\sigma_{\min}(X)} \\
&\leq 8\kappa d,
\end{aligned}$$

where in the last two inequalities we use (30), the assumption $d \leq \sigma_{\min}(X)/10\sqrt{2} \leq \sigma_{\max}(X)/10\sqrt{2}$ and the fact

$$\|\Sigma_l\| = \|Z_l\| \leq \|X\| + \|Z_l - X\| \leq \sigma_{\max}(X) + d.$$

It remains to estimate the incoherence of \widehat{Z}_l . Because A_l and B_l are not necessarily orthogonal, we consider their QR factorizations:

$$A_l = \widetilde{U}_l R_u \text{ and } B_l = \widetilde{V}_l R_v.$$

First note that

$$\begin{aligned}
\sigma_{\min}(A_l) &\geq 1 - \|A_l - UQ_u\| \geq 1 - \frac{\sqrt{2}d}{\sigma_{\min}(X)} \geq \frac{9}{10}, \\
\sigma_{\min}(B_l) &\geq 1 - \|B_l - VQ_v\| \geq 1 - \frac{\sqrt{2}d}{\sigma_{\min}(X)} \geq \frac{9}{10}
\end{aligned}$$

following from (31), (32), the assumption $d \leq \sigma_{\min}(X)/10\sqrt{2}$ and the Weyl inequality. Therefore

$$\|R_u^{-1}\| \leq \frac{10}{9} \text{ and } \|R_v^{-1}\| \leq \frac{10}{9}.$$

Consequently,

$$\left\| \widehat{U}_l^{(i)} \right\| = \left\| \widetilde{U}_l^{(i)} \right\| = \left\| A_l^{(i)} R_u^{-1} \right\| \leq \frac{10}{9} \sqrt{\frac{\mu_0 r}{n}} \quad \text{and} \quad \left\| \widehat{V}_l^{(i)} \right\| = \left\| \widetilde{V}_l^{(i)} \right\| = \left\| B_l^{(i)} R_v^{-1} \right\| \leq \frac{10}{9} \sqrt{\frac{\mu_0 r}{n}}$$

following from the construction of A_l and B_l . □

Proof of Lemma 2.6. First assume that

$$\|Z_l - X\|_F \leq \frac{\sigma_{\min}(X)}{256\kappa^2}. \tag{33}$$

It follows immediately from Lem. 4.10 that \widehat{Z}_l is an incoherent matrix with the incoherence parameter $\frac{100}{81} \mu_0$ and

$$\left\| \widehat{Z}_l - X \right\|_F \leq 8\kappa \|Z_l - X\|_F.$$

Analogous to (24), the approximation error at the $(l + 1)$ th iteration can be decomposed as follows

$$\begin{aligned}
\|Z_{l+1} - X\|_F &\leq 2 \left\| \left(\mathcal{P}_{\widehat{T}_l} - \frac{n^2}{\widehat{m}} \mathcal{P}_{\widehat{T}_l} \mathcal{P}_{\Omega_{l+1}} \mathcal{P}_{\widehat{T}_l} \right) (\widehat{Z}_l - X) \right\|_F \\
&\quad + 2 \left\| (\mathcal{I} - \mathcal{P}_{\widehat{T}_l}) (\widehat{Z}_l - X) \right\|_F \\
&\quad + 2 \left\| \frac{n^2}{\widehat{m}} \mathcal{P}_{\widehat{T}_l} \mathcal{P}_{\Omega_{l+1}} (\mathcal{I} - \mathcal{P}_{\widehat{T}_l}) (\widehat{Z}_l - X) \right\|_F \\
&:= I_8 + I_9 + I_{10}.
\end{aligned} \tag{34}$$

Since \widehat{Z}_l (and consequently \widehat{T}_l) is independent of Ω_{l+1} , Thm. A.2 implies

$$\left\| \mathcal{P}_{\widehat{T}_l} - \frac{n^2}{\widehat{m}} \mathcal{P}_{\widehat{T}_l} \mathcal{P}_{\Omega_{l+1}} \mathcal{P}_{\widehat{T}_l} \right\| \leq \sqrt{\frac{3200\beta\mu_0nr \log(n)}{243\widehat{m}}}.$$

with probability at least $1 - 2n^{2-2n}$. So

$$\begin{aligned}
I_8 &\leq 2 \sqrt{\frac{3200\beta\mu_0nr \log(n)}{243\widehat{m}}} \left\| \widehat{Z}_l - X \right\|_F \\
&\leq 16\kappa \sqrt{\frac{3200\beta\mu_0nr \log(n)}{243\widehat{m}}} \|Z_l - X\|_F.
\end{aligned}$$

For I_9 , Lem. 4.1 implies

$$I_9 \leq \frac{2 \left\| \widehat{Z}_l - X \right\|_F^2}{\sigma_{\min}(X)} \leq \frac{128\kappa^2 \|Z_l - X\|_F^2}{\sigma_{\min}(X)} \leq \frac{1}{2} \|Z_l - X\|_F,$$

where the last inequality follows from (33).

To bound I_{10} , note again \widehat{Z}_l (and consequently \widehat{T}_l) is independent of Ω_{l+1} with the incoherence parameter $\frac{100}{81}\mu_0$. So Lem. 4.3 implies

$$\left\| \frac{n^2}{\widehat{m}} \mathcal{P}_{\widehat{T}_l} \mathcal{P}_{\Omega_{l+1}} (\mathcal{P}_U - \mathcal{P}_{U_l}) - \mathcal{P}_{\widehat{T}_l} (\mathcal{P}_U - \mathcal{P}_{U_l}) \right\| \leq \sqrt{\frac{4800\beta\mu_0nr \log(n)}{81\widehat{m}}}.$$

with probability at least $1 - 2n^{2-2\beta}$. Since

$$\begin{aligned}
(\mathcal{I} - \mathcal{P}_{\widehat{T}_l})(\widehat{Z}_l - X) &= -(\mathcal{I} - \mathcal{P}_{\widehat{T}_l})(X) \\
&= -UU^*X + \widehat{U}_l\widehat{U}_l^*X + UU^*X\widehat{V}_l\widehat{V}_l^* - \widehat{U}_l\widehat{U}_l^*X\widehat{V}_l\widehat{V}_l^* \\
&= -(UU^* - \widehat{U}_l\widehat{U}_l^*)X(I - \widehat{V}_l\widehat{V}_l^*) \\
&= (UU^* - \widehat{U}_l\widehat{U}_l^*)(\widehat{Z}_l - X)(I - \widehat{V}_l\widehat{V}_l^*) \\
&= (\mathcal{P}_U - \mathcal{P}_{\widehat{U}_l})(\mathcal{I} - \mathcal{P}_{\widehat{V}_l})(\widehat{Z}_l - X),
\end{aligned}$$

we have

$$I_{10} = 2 \left\| \frac{n^2}{\widehat{m}} \mathcal{P}_{\widehat{T}_l} \mathcal{P}_{\Omega_{l+1}} (\mathcal{P}_U - \mathcal{P}_{\widehat{U}_l})(\mathcal{I} - \mathcal{P}_{\widehat{V}_l})(\widehat{Z}_l - X) \right\|_F$$

$$\begin{aligned}
&= 2 \left\| \frac{n^2}{\widehat{m}} \mathcal{P}_{\widehat{T}_l} \mathcal{P}_{\Omega_{l+1}} (\mathcal{P}_U - \mathcal{P}_{\widehat{U}_l}) (\mathcal{I} - \mathcal{P}_{\widehat{V}_l}) (\widehat{Z}_l - X) - \mathcal{P}_{\widehat{T}_l} (\mathcal{P}_U - \mathcal{P}_{\widehat{U}_l}) (\mathcal{I} - \mathcal{P}_{\widehat{V}_l}) (\widehat{Z}_l - X) \right\|_F \\
&\leq 2 \left\| \frac{n^2}{\widehat{m}} \mathcal{P}_{T_l} \mathcal{P}_{\Omega_{l+1}} (\mathcal{P}_U - \mathcal{P}_{U_l}) - \mathcal{P}_{T_l} (\mathcal{P}_U - \mathcal{P}_{U_l}) \right\| \left\| \widehat{Z}_l - X \right\|_F \\
&\leq 2 \sqrt{\frac{4800\beta\mu_0 nr \log(n)}{81\widehat{m}}} \left\| \widehat{Z}_l - X \right\|_F \\
&\leq 16\kappa \sqrt{\frac{4800\beta\mu_0 nr \log(n)}{81\widehat{m}}} \left\| Z_l - X \right\|_F.
\end{aligned}$$

where the second equality follows from

$$\mathcal{P}_{\widehat{T}_l} (\mathcal{P}_U - \mathcal{P}_{\widehat{U}_l}) (\mathcal{I} - \mathcal{P}_{\widehat{V}_l}) (\widehat{Z}_l - X) = \mathcal{P}_{\widehat{T}_l} (\mathcal{I} - \mathcal{P}_{\widehat{T}_l}) (\widehat{Z}_l - X) = 0.$$

Putting the bounds for I_8 , I_9 and I_{10} together gives

$$\begin{aligned}
\|Z_{l+1} - X\|_F &\leq \left(\frac{1}{2} + 182\kappa \sqrt{\frac{\beta\mu_0 nr \log(n)}{\widehat{m}}} \right) \|Z_l - X\|_F \\
&\leq \frac{5}{6} \|Z_l - X\|_F
\end{aligned}$$

with probability at least $1 - 4n^{2-2\beta}$ provided

$$\widehat{m} \geq C\beta\mu_0\kappa^2 nr \log(n) \quad (35)$$

with C being sufficiently large. Clearly (33) is also valid for $(l+1)$ th iteration.

Since $Z_0 = \mathcal{H}_r \left(\frac{\widehat{m}}{n^2} \mathcal{P}_{\Omega_0}(X) \right)$, (33) is valid for $l=0$ with probability $1 - 2n^{1-\beta}$ provided

$$\widehat{m} \geq C\beta\mu_1^2\kappa^6 nr^2 \log(n) \quad (36)$$

following from Lem. 2.4. Therefore taking a maximum of the right hand sides of (35) and (36) gives

$$\|Z_L - X\|_F \leq \left(\frac{5}{6} \right)^L \frac{\sigma_{\min}(X)}{256\kappa^2}$$

with probability at least $1 - 2n^{1-\beta} - 4Ln^{2-2\beta}$ provided

$$\widehat{m} \geq C\beta \max\{\mu_0, \mu_1^2\} \kappa^6 nr^2 \log(n).$$

□

5 Conclusion and Future Direction

This manuscript presents the sampling complexity for a class of Riemannian gradient descent and conjugate gradient descent algorithms by interpreting them as iterative hard thresholding algorithms with subspace projections. To the best of our knowledge, this is the first work that provides recovery guarantees of the Riemannian optimization algorithms on the embedded low

rank matrix manifold for matrix completion. Our first sampling complexity for the algorithms with one step hard thresholding initialization is $O(n^{1.5}r \log^{1.5}(n))$. We would like to know whether this result can be improved to the nearly optimal one $O(nr \text{polylog}(n))$ [14] since the simulation results in Sec. 3.1 strongly suggests so. The extra r term in the second recovery guarantee result comes from bounding the matrix spectral norm by the Frobenius norm. To further optimize the second sampling complexity, we may need to establish convergence of the algorithms in terms of the matrix spectral norm rather than the Frobenius norm.

It would be of interest to investigate whether we can extend our results to the case whether the unknown matrix is approximately low rank. Since in this case $\mathcal{P}_\Omega(X)$ can be decomposed as $\mathcal{P}_\Omega(X) = \mathcal{P}_\Omega(X_r) + \mathcal{P}_\Omega(X - X_r)$, where X_r is the best rank r approximation of X , it is equivalent to study the robustness of the algorithms under additive noise which is interesting by itself. The empirical observations in Sec. 3.3 suggests that both the Riemannian gradient descent and conjugate gradient descent algorithms are very robust under additive Gaussian white noise. Robustness analysis of the algorithms is scope for future work. It would also be desirable to consider more general and practical sampling models in matrix completion for the Riemannian optimization algorithms (Algs. 1 and 2).

The Riemannian gradient descent and conjugate gradient descent algorithms presented in this manuscript apply equally to other low rank reconstruction problems, such as phase retrieval [20, 10, 11, 48, 16, 41] and blind deconvolution [3, 32]. This line of research will be pursued independently in the future. In phase retrieval and blind deconvolution, the underlying matrix after lifting is rank one. Since the condition number of a rank one matrix is always equal to one, it is worth investigating whether we can obtain recovery guarantees of the form $O(n \text{polylog}(n))$ with a universal constant.

Acknowledgments

KW has been supported by DTRA-NSF grant No.1322393. He would like to thank Xiaodong Li, Shuyang Ling and Thomas Strohmer for stimulating discussions about phase retrieval and blind deconvolution.

References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
- [2] R. Ahlswede and A. Winter. Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48(3):569–579, 2002.
- [3] A. Ahmed, B. Recht, and J. Romberg. Blind deconvolution using convex programming. *IEEE Transactions on Information Theory*, 60(3):1711–1732, 2013.
- [4] Y. Amit, M. Fink, N. Srebro, and S. Ullman. Uncovering shared structures in multiclass classification. In *Proceedings of the 24th International Conference on Machine Learning*, 2007.
- [5] A. Argyriou, T. Evgeniou, and M. Pontil. Multi-task feature learning. In *Advances in Neural Information Processing Systems*, 2007.

- [6] S. Bhojanapalli, A. Kyrillidis, and S. Sanghavi. Dropping convexity for faster semi-definite optimization. arXiv:1509.03917, 2015.
- [7] J. Blanchard, J. Tanner, and K. Wei. CGIHT: Conjugate gradient iterative hard thresholding for compressed sensing and matrix completion. *Information and Inference*, 4(4):289–327, 2015.
- [8] N. Boumal and P.-A. Absil. RTRMC: A Riemannian trust-region method for low-rank matrix completion. In *Advances in Neural Information Processing Systems*, 2011.
- [9] J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.
- [10] E. J. Candès, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM J. on Imaging Sciences*, 6(1):199–225, 2013.
- [11] E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- [12] E. J. Candès and Y. Plan. Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359, 2009.
- [13] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.
- [14] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–1080, 2009.
- [15] Y. Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- [16] Y. Chen and E. Candès. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Communications on Pure and Applied Mathematics*, 2015. To appear.
- [17] Y. Chen and M. J. Wainwright. Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. arXiv:1509.03025, 2015.
- [18] L. Eldén. *Matrix Methods in Data Mining and Pattern Recognition*. SIAM, 2007.
- [19] M. Fazel. Matrix rank minimization with applications. Ph. D. dissertation, Stanford University, 2002.
- [20] R. W. Gerchberg and W. O. Saxton. A practical algorithm for the determination of the phase from image and diffraction plane pictures. *Optik*, 35(237):237–246, 1972.
- [21] D. Goldfarb and S. Ma. Convergence of fixed-point continuation algorithms for matrix rank minimization. *Foundations of Computational Mathematics*, 11(2):183–210, 2011.
- [22] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.

- [23] J. P. Haldar and D. Hernando. Rank-constrained solutions to linear matrix equations using PowerFactorization. *IEEE Signal Processing Letters*, 16(7):584–587, 2009.
- [24] N. J. A. Harvey, D. R. Karger, and S. Yekhanin. The complexity of matrix completion. In *Proceedings of the seventeenth annual ACM-SIAM symposium on discrete algorithms*, 2006.
- [25] P. Jain, R. Meka, and I. Dhillon. Guaranteed rank minimization via singular value projection. In *Proceedings of the Neural Information Processing Systems Conference*, 2010.
- [26] P. Jain and P. Netrapalli. Fast exact matrix completion with finite samples. *JMLR: Workshop and Conference Proceedings*, 40:1–28, 2015.
- [27] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. arXiv:1212.0467, 2012.
- [28] R. H. Keshavan. Efficient algorithms for collaborative filtering. Ph. D. dissertation, Stanford University, 2012.
- [29] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, 2010.
- [30] A. Kyrillidis and V. Cevher. Matrix recipes for hard thresholding methods. *Journal of Mathematical Imaging and Vision*, 48(2):235–265, 2014.
- [31] K. Lee and Y. Bresler. ADMiRA: Atomic decomposition for minimum rank approximation. *IEEE Transactions on Information Theory*, 128(1):4402–4416, 2010.
- [32] S. Ling and T. Strohmer. Blind deconvolution meets blind demixing: Algorithms and performance bounds. arXiv:1512.07730, 2015.
- [33] Z. Liu and L. Vandenberghe. Interior-point method for nuclear norm approximation with application to system identification. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1235–1256, 2009.
- [34] B. Mishra, K. A. Apuroop, and R. Sepulchre. A Riemannian geometry for low-rank matrix completion. arXiv:1306.2672, 2013.
- [35] B. Mishra, G. Meyer, S. Bonnabel, and R. Sepulchre. Fixed-rank matrix factorizations and Riemannian low-rank optimization. *Computational Statistics*, 29(3):591–261, 2014.
- [36] T. Ngo and Y. Saad. Scaled gradients on Grassmann manifolds for matrix completion. In *Advances in Neural Information Processing Systems*, 2012.
- [37] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.
- [38] B. Recht. A simpler approach to matrix completion. *The Journal of Machine Learning Research*, 12:3413–3430, 2011.
- [39] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

- [40] C. De Sa, K. Olukotun, and C. Ré. Global convergence of stochastic gradient descent for some nonconvex matrix problems. In *ICML*, 2015.
- [41] J. Sun, Q. Qu, and J. Wright. A geometric analysis of phase retrieval. arXiv:1602.06664, 2016.
- [42] R. Sun and Z. Luo. Guaranteed matrix completion via non-convex factorization. In *FOCS*, 2015.
- [43] J. Tanner and K. Wei. Normalized iterative hard thresholding for matrix completion. *SIAM Journal on Scientific Computing*, 35(5):S104–S125, 2013.
- [44] J. Tanner and K. Wei. Low rank matrix completion by alternating steepest descent methods. *Applied and Computational Harmonic Analysis*, 40(2):417–429, 2016.
- [45] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht. Low-rank solutions of linear matrix equations via Procrustes flow. arXiv:1507.03566, 2016.
- [46] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- [47] B. Vandereycken. Low rank matrix completion by Riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, 2013.
- [48] K. Wei. Solving systems of phaseless equations via Kaczmarz methods: A proof of concept study. *Inverse Problems*, 31:125008, 2015.
- [49] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung. Guarantees of Riemannian optimization for low rank matrix recovery. arXiv:1511.01562, 2015.
- [50] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a non-linear successive over-relaxation algorithm. *Mathematical Programming Computation*, 4(4):333–361, 2012.
- [51] C. D. White, S. Sanghavi, and R. Ward. The local convexity of solving systems of quadratic equations. arXiv:1506.07868, 2015.
- [52] T. Zhao, Z. Wang, and H. Liu. Nonconvex low rank matrix factorization via inexact first order oracle. In *NIPS*, 2015.
- [53] Q. Zheng and J. Lafferty. A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. In *NIPS*, 2015.

A Supplementary Results

In this section, we will list the tail probability bounds from the literature which have been used in this manuscript. We start with the Noncommutative Bernstein inequality.

Theorem A.1 (Noncommutative Bernstein Inequality, [2, 38]). *Let Z_1, \dots, Z_L be $n \times n$ independent zero-mean random matrices. Suppose $\sigma_k^2 = \max \{ \|\mathbb{E}[Z_k Z_k^*]\|, \|\mathbb{E}[Z_k^* Z_k]\| \}$ and $\|Z_k\| \leq R$ almost surely for all k . Then for any $t > 0$,*

$$\mathbb{P} \left[\left\| \sum_{k=1}^L Z_k \right\| > t \right] \leq 2n \exp \left(\frac{-t^2/2}{\sum_{k=1}^L \sigma_k^2 + Rt/3} \right).$$

Moreover, if $t \leq \frac{1}{R} \sum_{k=1}^L \sigma_k^2$,

$$\mathbb{P} \left[\left\| \sum_{k=1}^L Z_k \right\| > t \right] \leq 2n \exp \left(\frac{-\frac{3}{8}t^2}{\sum_{k=1}^L \sigma_k^2} \right).$$

Two tail probability bounds for \mathcal{P}_Ω can be obtained using the Noncommutative Bernstein inequality under the sampling with replacement model.

Theorem A.2 ([38, 22, 13]). *Suppose Ω with $|\Omega| = m$ is a set of indices sampled independently and uniformly with replacement. Let $X = U\Sigma V^* \in \mathbb{R}^{n \times n}$ be a rank r matrix with the corresponding tangent space T . Assume*

$$\|\mathcal{P}_U(e_i)\| \leq \sqrt{\frac{\mu r}{n}} \quad \text{and} \quad \|\mathcal{P}_V(e_j)\| \leq \sqrt{\frac{\mu r}{n}}$$

for $1 \leq i, j \leq n$. Then for all $\beta > 1$,

$$\left\| \mathcal{P}_T - \frac{n^2}{m} \mathcal{P}_T \mathcal{P}_\Omega \mathcal{P}_T \right\| \leq \sqrt{\frac{32\beta\mu nr \log(n)}{3m}}$$

with probability at least $1 - 2n^{2-2\beta}$ provided $m \geq \frac{32}{3}\beta\mu nr \log(n)$.

Theorem A.3 ([38, 22, 13]). *Suppose Ω with $|\Omega| = m$ is a set of indices sampled independently and uniformly with replacement. Let $Z \in \mathbb{R}^{n \times n}$ be a fixed matrix. Then for all $\beta > 1$,*

$$\left\| \left(\frac{n^2}{m} \mathcal{P}_\Omega - \mathcal{I} \right) (Z) \right\| \leq \sqrt{\frac{8\beta n^3 \log(n)}{3m}} \|Z\|_\infty$$

with probability $1 - 2n^{1-\beta}$ provided $m \geq 6\beta n \log(n)$.

B Proofs of Technical Lemmas

B.1 Proof of Lemma 4.1

The proofs for the first five inequalities in this lemma can be found in [49]. So it only remains to prove the last inequality. For any matrix $Z \in \mathbb{R}^{n \times n}$, we have

$$\begin{aligned} (\mathcal{P}_{T_l} - \mathcal{P}_T)(Z) &= U_l U_l^* Z + Z V_l V_l^* - U_l U_l^* Z V_l V_l^* - U U^* Z - Z V V^* + U U^* Z V V^* \\ &= (U_l U_l^* - U U^*) Z (I - V V^*) + (I - U_l U_l^*) Z (V_l V_l^* - V V^*). \end{aligned}$$

Taking the Frobenius norm on both sides gives

$$\begin{aligned} \|(\mathcal{P}_{T_l} - \mathcal{P}_T)(Z)\|_F &\leq \|U_l U_l^* - U U^*\| \|Z\|_F \|I - V V^*\| + \|I - U_l U_l^*\| \|Z\|_F \|V_l V_l^* - V V^*\| \\ &\leq \frac{2 \|X_l - X\|_F}{\sigma_{\min}(X)} \|Z\|_F. \end{aligned}$$

B.2 Proof of Lemma 4.2

For any $Z \in \mathbb{R}^{n \times n}$, we have

$$\begin{aligned}
\|\mathcal{P}_\Omega \mathcal{P}_T(Z)\|_F^2 &= \langle \mathcal{P}_\Omega \mathcal{P}_T(Z), \mathcal{P}_\Omega \mathcal{P}_T(Z) \rangle \\
&\leq \frac{8}{3} \beta \log(n) \langle \mathcal{P}_T(Z), \mathcal{P}_\Omega \mathcal{P}_T(Z) \rangle \\
&= \frac{8}{3} \beta \log(n) \langle \mathcal{P}_T(Z), \mathcal{P}_T \mathcal{P}_\Omega \mathcal{P}_T(Z) \rangle \\
&\leq \frac{8}{3} \beta \log(n) (1 + \varepsilon_0) p \|\mathcal{P}_T(Z)\|_F^2,
\end{aligned}$$

where the first inequality follows from the first assumption and the second inequality follows from the second assumption. So we have $\|\mathcal{P}_\Omega \mathcal{P}_T\| \leq \sqrt{\frac{8}{3} \beta \log(n) (1 + \varepsilon_0) p}$ and

$$\begin{aligned}
\|\mathcal{P}_\Omega \mathcal{P}_{T_l}\| &\leq \|\mathcal{P}_\Omega (\mathcal{P}_{T_l} - \mathcal{P}_T)\| + \|\mathcal{P}_\Omega \mathcal{P}_T\| \\
&\leq \frac{8}{3} \beta \log(n) \frac{2 \|X_l - X\|_F}{\sigma_{\min}(X)} + \|\mathcal{P}_\Omega \mathcal{P}_T\| \\
&\leq \frac{p^{1/2} \varepsilon_0}{1 + \varepsilon_0} + \sqrt{\frac{8}{3} \beta \log(n) (1 + \varepsilon_0) p} \\
&\leq \frac{8}{3} \beta \log(n) (1 + \varepsilon_0) p^{1/2},
\end{aligned}$$

where the second inequality follows from (16). To prove (18), we use

$$\begin{aligned}
\|\mathcal{P}_{T_l} - p^{-1} \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l}\| &\leq \|\mathcal{P}_{T_l} - \mathcal{P}_T\| + p^{-1} \|\mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_{T_l} - \mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_T\| \\
&\quad + p^{-1} \|\mathcal{P}_{T_l} \mathcal{P}_\Omega \mathcal{P}_T - \mathcal{P}_T \mathcal{P}_\Omega \mathcal{P}_T\| + \|\mathcal{P}_T - p^{-1} \mathcal{P}_T \mathcal{P}_\Omega \mathcal{P}_T\| \\
&\leq \|\mathcal{P}_{T_l} - \mathcal{P}_T\| + p^{-1} \|\mathcal{P}_\Omega \mathcal{P}_{T_l}\| \|\mathcal{P}_{T_l} - \mathcal{P}_T\| \\
&\quad + p^{-1} \|\mathcal{P}_\Omega \mathcal{P}_T\| \|\mathcal{P}_{T_l} - \mathcal{P}_T\| + \|\mathcal{P}_T - p^{-1} \mathcal{P}_T \mathcal{P}_\Omega \mathcal{P}_T\| \\
&\leq \frac{2 \|X_l - X\|_F}{\sigma_{\min}(X)} + p^{-1} \|\mathcal{P}_\Omega \mathcal{P}_{T_l}\| \frac{2 \|X_l - X\|_F}{\sigma_{\min}(X)} \\
&\quad + p^{-1} \|\mathcal{P}_\Omega \mathcal{P}_T\| \frac{2 \|X_l - X\|_F}{\sigma_{\min}(X)} + \|\mathcal{P}_T - p^{-1} \mathcal{P}_T \mathcal{P}_\Omega \mathcal{P}_T\| \\
&\leq 4\varepsilon_0,
\end{aligned}$$

where in the second inequality we utilize the fact $\mathcal{P}_\Omega^* = \mathcal{P}_\Omega$ so that $\|\mathcal{P}_{T_l} \mathcal{P}_\Omega\| = \|\mathcal{P}_\Omega \mathcal{P}_{T_l}\|$; and the last inequality follows from the assumption and the bounds for $\|\mathcal{P}_\Omega \mathcal{P}_{T_l}\|$ and $\|\mathcal{P}_\Omega \mathcal{P}_T\|$.

B.3 Proof of Lemma 4.3

Lemma 4.3 is an asymmetric version of Thm. A.2. We will use the Noncommutative Bernstein inequality (Thm. A.1) to prove it. Let $\Omega = \{(i_k, j_k)\}_{k=1}^m$ be the sampled set of indices. For any $Z \in \mathbb{R}^{n \times n}$, since

$$(\mathcal{P}_U - \mathcal{P}_{U_l})(Z) = \sum_{i,j=1}^n \langle (\mathcal{P}_U - \mathcal{P}_{U_l})(Z), e_i e_j^* \rangle e_i e_j^* = \sum_{i,j=1}^n \langle Z, (\mathcal{P}_U - \mathcal{P}_{U_l})(e_i e_j^*) \rangle e_i e_j^*,$$

we have

$$\mathcal{P}_\Omega (\mathcal{P}_U - \mathcal{P}_{U_l}) (Z) = \sum_{k=1}^m \langle Z, (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*) \rangle e_{i_k} e_{j_k}^*$$

and

$$\mathcal{P}_{T_l} \mathcal{P}_\Omega (\mathcal{P}_U - \mathcal{P}_{U_l}) (Z) = \sum_{k=1}^m \langle Z, (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*) \rangle \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*).$$

Let $\mathcal{T}_{i_k, j_k} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ be a rank one linear operator defined as

$$\mathcal{T}_{i_k, j_k} = \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*) \otimes (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*).$$

Then we have

$$\mathcal{P}_{T_l} \mathcal{P}_\Omega (\mathcal{P}_U - \mathcal{P}_{U_l}) = \sum_{k=1}^m \mathcal{T}_{i_k, j_k}, \quad (37)$$

$$\mathcal{T}_{i_k, j_k}^* = (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*) \otimes \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*), \quad (38)$$

$$\mathbb{E} [\mathcal{T}_{i_k, j_k}] = \frac{1}{n^2} \mathcal{P}_{T_l} (\mathcal{P}_U - \mathcal{P}_{U_l}), \quad (39)$$

$$\mathbb{E} [\mathcal{T}_{i_k, j_k}^*] = \frac{1}{n^2} (\mathcal{P}_U - \mathcal{P}_{U_l}) \mathcal{P}_{T_l}. \quad (40)$$

Moreover,

$$\begin{aligned} \|\mathcal{T}_{i_k, j_k}\| &\leq \|(\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*)\|_F \|\mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*)\|_F \\ &\leq \left(\|\mathcal{P}_U (e_{i_k} e_{j_k}^*)\|_F + \|\mathcal{P}_{U_l} (e_{i_k} e_{j_k}^*)\|_F \right) \|\mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*)\|_F \\ &\leq \frac{4\mu r}{n}, \end{aligned}$$

where the last inequality follows from

$$\|\mathcal{P}_U (e_{i_k} e_{j_k}^*)\|_F = \|\mathcal{P}_U (e_{i_k}) e_{j_k}^*\|_F \leq \sqrt{\frac{\mu r}{n}}, \quad (41)$$

$$\|\mathcal{P}_{U_l} (e_{i_k} e_{j_k}^*)\|_F = \|\mathcal{P}_{U_l} (e_{i_k}) e_{j_k}^*\|_F \leq \sqrt{\frac{\mu r}{n}}, \quad (42)$$

and

$$\begin{aligned} \|\mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*)\|_F^2 &= \langle \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*), e_{i_k} e_{j_k}^* \rangle \\ &= \langle \mathcal{P}_{U_l} (e_{i_k}) e_{j_k}^* + e_{i_k} (\mathcal{P}_{V_l} (e_{j_k}))^* - \mathcal{P}_{U_l} (e_{i_k}) (\mathcal{P}_{V_l} (e_{j_k}))^*, e_{i_k} e_{j_k}^* \rangle \\ &= \|\mathcal{P}_{U_l} (e_{i_k})\|^2 + \|\mathcal{P}_{V_l} (e_{j_k})\|^2 - \|\mathcal{P}_{U_l} (e_{i_k})\|^2 \|\mathcal{P}_{V_l} (e_{j_k})\|^2 \\ &\leq \frac{2\mu r}{n}. \end{aligned} \quad (43)$$

So

$$\|\mathcal{T}_{i_k, j_k} - \mathbb{E} [\mathcal{T}_{i_k, j_k}]\| \leq \|\mathcal{T}_{i_k, j_k}\| + \|\mathbb{E} [\mathcal{T}_{i_k, j_k}]\| \leq \frac{4\mu r}{n} + \frac{2}{n^2} \leq \frac{5\mu r}{n}. \quad (44)$$

To apply the Bernstein inequality, we also need to bound

$$\|\mathbb{E}[(\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k}))^* (\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k}))]\|$$

and

$$\|\mathbb{E}[(\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k})) (\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k}))^*]\|.$$

The first one can be proceeded as follows

$$\begin{aligned} & \|\mathbb{E}[(\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k}))^* (\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k}))]\| \\ &= \|\mathbb{E}[\mathcal{T}_{i_k, j_k}^* \mathcal{T}_{i_k, j_k}] - \mathbb{E}[\mathcal{T}_{i_k, j_k}^*] \mathbb{E}[\mathcal{T}_{i_k, j_k}]\| \\ &\leq \|\mathbb{E}[\mathcal{T}_{i_k, j_k}^* \mathcal{T}_{i_k, j_k}]\| + \left\| \frac{1}{n^4} (\mathcal{P}_U - \mathcal{P}_{U_l}) \mathcal{P}_{T_l} (\mathcal{P}_U - \mathcal{P}_{U_l}) \right\| \\ &\leq \left\| \mathbb{E} \left[\|\mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*)\|_F^2 (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*) \otimes (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*) \right] \right\| + \frac{4}{n^4} \\ &\leq \frac{2\mu r}{n} \|\mathbb{E}[(\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*) \otimes (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*)]\| + \frac{4}{n^4} \\ &\leq \frac{2\mu r}{n} \frac{1}{n^2} \|(\mathcal{P}_U - \mathcal{P}_{U_l})^2\| + \frac{4}{n^4} \\ &\leq \frac{9\mu r}{n^3}, \end{aligned} \tag{45}$$

where the first inequality follows from (39) and (40), the third inequality follows from (43), and the fourth inequality follows from the the fact

$$\mathbb{E}[(\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*) \otimes (\mathcal{P}_U - \mathcal{P}_{U_l}) (e_{i_k} e_{j_k}^*)] = \frac{1}{n^2} (\mathcal{P}_U - \mathcal{P}_{U_l})^2.$$

Similarly, we have

$$\begin{aligned} & \|\mathbb{E}[(\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k})) (\mathcal{T}_{i_k, j_k} - \mathbb{E}(\mathcal{T}_{i_k, j_k}))^*]\| \\ &= \|\mathbb{E}[\mathcal{T}_{i_k, j_k} \mathcal{T}_{i_k, j_k}^*] - \mathbb{E}[\mathcal{T}_{i_k, j_k}] \mathbb{E}[\mathcal{T}_{i_k, j_k}^*]\| \\ &\leq \|\mathbb{E}[\mathcal{T}_{i_k, j_k} \mathcal{T}_{i_k, j_k}^*]\| + \frac{1}{n^4} \|\mathcal{P}_{T_l} (\mathcal{P}_U - \mathcal{P}_{U_l})^2 \mathcal{P}_{T_l}\| \\ &\leq \left\| \mathbb{E} \left[\|\mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*)\|_F^2 \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*) \otimes \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*) \right] \right\| + \frac{4}{n^2} \\ &\leq \frac{4\mu r}{n} \|\mathbb{E}[\mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*) \otimes \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*)]\| + \frac{4}{n^4} \\ &\leq \frac{4\mu r}{n} \frac{1}{n^2} \|\mathcal{P}_{T_l}\| + \frac{4}{n^4} \\ &\leq \frac{5\mu r}{n^3}, \end{aligned} \tag{46}$$

where the first inequality follows from (39) and (40), the third inequality follows from (41) and (42), and the fourth inequality follows from the fact

$$\mathbb{E}[\mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*) \otimes \mathcal{P}_{T_l} (e_{i_k} e_{j_k}^*)] = \frac{1}{n^2} \mathcal{P}_{T_l}.$$

Finally, the theorem follows by applying Thm. A.1 to (37) with the bounds (44), (45) and (46).

B.4 Proof of Lemma 4.4

By the assumption, (19) is valid for $l = 0$. Assume it is valid for $j \leq l$. Then

$$\begin{aligned}
c_{l+1} &\leq \rho_1 c_l + \rho_2 \sum_{j=0}^{l-1} \gamma^{l-j} c_j \\
&\leq \rho_1 \nu^l c_0 + \rho_2 \sum_{j=0}^{l-1} \gamma^{l-j} \nu^j c_0 \\
&= \nu^l c_0 \left(\rho_1 + \rho_2 \sum_{j=0}^{l-1} \left(\frac{\gamma}{\nu} \right)^{l-j} \right) \\
&\leq \nu^l c_0 \left(\rho_1 + \frac{\rho_2 \gamma / \nu}{1 - \gamma / \nu} \right) \\
&= \nu^l c_0 \left(\rho_1 + \frac{\rho_2 \gamma}{\nu - \gamma} \right).
\end{aligned}$$

So it suffices to show that $\rho_1 + \frac{\rho_2 \gamma}{\nu - \gamma} = \nu$. It is equivalent to show that

$$\nu^2 - (\rho_1 + \gamma)\nu - (\rho_2 - \rho_1)\gamma = 0.$$

which is true since

$$\nu = \frac{1}{2} \left(\tau_1 + \sqrt{\tau_1^2 + 4\tau_2} \right)$$

with $\tau_1 = \rho_1 + \gamma$ and $\tau_2 = (\rho_2 - \rho_1)\gamma$.

B.5 Proof of Lemma 4.5

Since

$$\begin{aligned}
\|U_l - UQ\|_F^2 &= 2r - 2\langle U_l, UQ \rangle, \\
\|U_l U_l^* - UU^*\|_F^2 &= 2r - 2\langle U_l U_l^*, UU^* \rangle,
\end{aligned}$$

it suffices to show that there exist a Q such that

$$\langle U_l, UQ \rangle \geq \langle U_l U_l^*, UU^* \rangle.$$

It is equivalent to show that

$$\langle U^* U_l, Q \rangle \geq \langle U^* U_l, U^* U_l \rangle$$

for some unitary matrix $Q \in \mathbb{R}^{r \times r}$. Let $U_l^* U_l = Q_1 \Lambda Q_2^*$ be the singular value decomposition of $U_l^* U_l$. Then we have $\Lambda(i, i) \leq 1$ ($1 \leq i \leq r$), and we can choose $Q = Q_1 Q_2^*$.