

Low-Rank Matrix Recovery with Scaled Subgradient Methods: Fast and Robust Convergence Without the Condition Number

Tian Tong*
Carnegie Mellon University

Cong Ma†
UC Berkeley

Yuejie Chi*
Carnegie Mellon University

October 2020; Revised April 2021

Abstract

Many problems in data science can be treated as estimating a low-rank matrix from highly incomplete, sometimes even corrupted, observations. One popular approach is to resort to matrix factorization, where the low-rank matrix factors are optimized via first-order methods over a smooth loss function, such as the residual sum of squares. While tremendous progress has been made in recent years, the natural smooth formulation suffers from two sources of ill-conditioning, where the iteration complexity of gradient descent scales poorly both with the dimension as well as the condition number of the low-rank matrix. Moreover, the smooth formulation is not robust to corruptions. In this paper, we propose *scaled* subgradient methods to minimize a family of nonsmooth and nonconvex formulations—in particular, the residual sum of absolute errors—which is guaranteed to converge at a fast rate that is almost dimension-free and independent of the condition number, even in the presence of corruptions. We illustrate the effectiveness of our approach when the observation operator satisfies certain mixed-norm restricted isometry properties, and derive state-of-the-art performance guarantees for a variety of problems such as robust low-rank matrix sensing and quadratic sampling.

Keywords: low-rank matrix recovery, nonsmooth and nonconvex optimization, residual sum of absolute errors, scaled subgradient methods

1 Introduction

Many problems in data science can be treated as estimating a low-rank matrix $\mathbf{X}_\star \in \mathbb{R}^{n_1 \times n_2}$ from highly incomplete, sometimes even corrupted, observations $\mathbf{y} = \{y_i\}_{i=1}^m$ given by

$$y_i \approx \mathcal{A}_i(\mathbf{X}_\star), \quad 1 \leq i \leq m. \quad (1)$$

Here, $\mathcal{A}(\cdot) = \{\mathcal{A}_i(\cdot)\}_{i=1}^m : \mathbb{R}^{n_1 \times n_2} \mapsto \mathbb{R}^m$ is the observation operator that models the measurement process. Instead of operating in the full matrix space, i.e. $\mathbb{R}^{n_1 \times n_2}$, a memory-efficient approach is to resort to low-rank matrix factorization, by writing $\mathbf{X}_\star = \mathbf{L}_\star \mathbf{R}_\star^\top$, if the rank r of \mathbf{X}_\star is known *a priori*, where $\mathbf{L}_\star \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{R}_\star \in \mathbb{R}^{n_2 \times r}$ are of a size that is proportional to the degrees of freedom of the low-rank matrix. Furthermore, the low-rank factors can be found by optimizing a smooth loss function, such as the residual sum of squares

$$\underset{\mathbf{L} \in \mathbb{R}^{n_1 \times r}, \mathbf{R} \in \mathbb{R}^{n_2 \times r}}{\text{minimize}} \quad \sum_{i=1}^m (\mathcal{A}_i(\mathbf{L}\mathbf{R}^\top) - y_i)^2, \quad (2)$$

using first-order methods (e.g. gradient descent). While tremendous progress has been made in recent years [CLC19], applying vanilla gradient descent to the above smooth formulation suffers from two sources of ill-conditioning that preclude a desirable computational efficiency from classical optimization principles:

*Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA; Emails: {ttong1,yuejiec}@andrew.cmu.edu.

†Department of Statistics and Department of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, CA 94720, USA; Email: congm@berkeley.edu.

- Due to the heavy-tailed nature of certain measurement operators, such as those encountered in phase retrieval [CLS15] and quadratic sampling [SWW17], the least-squares formulation (2) may suffer from a large smoothness parameter (and hence a large condition number of the loss function) that scales at least linearly with respect to the ambient dimension, leading to a conservative choice of stepsizes and a high iteration complexity when the problem dimension is large.
- Due to the composite nature of the formulation (2), the iteration complexity of vanilla gradient descent is further exacerbated by the condition number of the underlying low-rank matrix \mathbf{X}_* , which could be large in many applications of interest.

While there have been encouraging activities [CCD⁺21, MWCC19, LMCC21, TMC20] that try to alleviate these issues regarding ill-conditioning, none of the existing first-order approaches are able to simultaneously remove both sources of ill-conditioning and achieve fast convergence. Therefore, the goal of the current paper is to develop first-order methods that are guaranteed to converge at a fast rate that is almost dimension-free and independent of the condition number, even in the presence of corruptions.

1.1 Main contributions

In this paper, we propose to minimize the following *nonsmooth and nonconvex* loss function known as the *least absolute deviations*, which measures the residual sum of absolute errors

$$\underset{\mathbf{L} \in \mathbb{R}^{n_1 \times r}, \mathbf{R} \in \mathbb{R}^{n_2 \times r}}{\text{minimize}} \quad f(\mathbf{L}\mathbf{R}^\top) := \sum_{i=1}^m |\mathcal{A}_i(\mathbf{L}\mathbf{R}^\top) - y_i|, \quad (3)$$

via a *scaled subgradient method*:

$$\begin{aligned} \mathbf{L}_{t+1} &:= \mathbf{L}_t - \eta_t \mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1}, \\ \mathbf{R}_{t+1} &:= \mathbf{R}_t - \eta_t \mathbf{S}_t^\top \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1}. \end{aligned} \quad (4)$$

Here, $\mathbf{S}_t \in \partial f(\mathbf{L}_t \mathbf{R}_t^\top)$ is a subgradient of $f(\mathbf{X}) := \sum_{i=1}^m |\mathcal{A}_i(\mathbf{X}) - y_i|$ at $\mathbf{L}_t \mathbf{R}_t^\top$, and $\eta_t > 0$ is a sequence of carefully-chosen stepsizes. Compared with vanilla subgradient methods, our new method (4) scales or preconditions the search directions $\mathbf{S}_t \mathbf{R}_t$ and $\mathbf{S}_t^\top \mathbf{L}_t$ by $(\mathbf{R}_t^\top \mathbf{R}_t)^{-1}$ and $(\mathbf{L}_t^\top \mathbf{L}_t)^{-1}$, respectively.¹ As explained in [TMC20] where a similar preconditioning trick was employed for smooth formulations, the scaled subgradient enables better search directions and therefore larger stepsizes. Our main results can be summarized as follows:

- Under general geometric assumptions on $f(\cdot)$ such as restricted rank- r Lipschitz continuity and sharpness conditions, we demonstrate that the convergence rate of scaled subgradient methods using both Polyak’s and geometrically decaying stepsizes is *independent* of the condition number of \mathbf{X}_* .
- Instantiating our theory under the mixed-norm restricted isometry property (RIP) of the measurement operator, we demonstrate state-of-the-art computational guarantees for low-rank matrix sensing and quadratic sampling even when the observations are noisy and corrupted by outliers. This leads to improvements over the computational complexity of scaled gradient methods in [TMC20] for heavy-tailed measurement ensembles, as well as of vanilla subgradient methods in [CCD⁺21]. Table 1 provides a detailed comparison of the local iteration complexities of the proposed scaled subgradient method in comparison with these prior algorithms, highlighting its robustness to heavy-tailed observations, outliers, as well as a large condition number of the true matrix \mathbf{X}_* .

Our work leverages exciting advances in nonsmooth optimization [CCD⁺21] and scaled first-order methods [TMC20] for low-rank matrix recovery. Our arguments are concise, which avoid the need of sophisticated trajectory-dependent analysis as have been used in [MWCC19, LMCC21] to achieve rapid and robust convergence guarantees.

¹Under appropriate conditions, the inverse matrices always exist; in practice, one can use the pseudo-inverse matrices to avoid numerical instabilities.

Algorithms	matrix sensing		quadratic sampling	
	without corruptions	with corruptions	without corruptions	with corruptions
GD [TBS ⁺ 16, LMCC21]	$\kappa \log \frac{1}{\epsilon}$	N/A	$r^2 \kappa^2 \log \frac{1}{\epsilon}$	N/A
ScaledGD [TMC20]	$\log \frac{1}{\epsilon}$	N/A	$\text{poly}(n) \log \frac{1}{\epsilon}$	N/A
SM [CCD ⁺ 21, LZSV20]	$\kappa \log \frac{1}{\epsilon}$	$\frac{\kappa}{(1-2p_s)^2} \log \frac{1}{\epsilon}$	$r \kappa \log \frac{1}{\epsilon}$	$\frac{r \kappa}{(1-2p_s)^2} \log \frac{1}{\epsilon}$
ScaledSM (this paper)	$\log \frac{1}{\epsilon}$	$\frac{1}{(1-2p_s)^2} \log \frac{1}{\epsilon}$	$r \log \frac{1}{\epsilon}$	$\frac{r}{(1-2p_s)^2} \log \frac{1}{\epsilon}$

Table 1: Local iteration complexities of the proposed scaled subgradient method (**ScaledSM**) in comparison with prior algorithms for matrix sensing and quadratic sampling. **ScaledSM** outperforms the vanilla subgradient method (**SM**) by a factor of κ in both problems, while outperforms scaled gradient descent (**ScaledGD**), and **GD** with additional robustness guarantees. Here, $n = \max\{n_1, n_2\}$, r is the rank, κ is the condition number of \mathbf{X}_* , and $0 \leq p_s < 1/2$ is the fraction of outliers. We say that the output \mathbf{X} of an algorithm reaches ϵ -accuracy, if it satisfies $\|\mathbf{X} - \mathbf{X}_*\|_F \leq \epsilon \sigma_r(\mathbf{X}_*)$, where $\sigma_r(\mathbf{X}_*)$ denotes the r -th largest singular value of \mathbf{X}_* .

1.2 Related work

Low-rank matrix recovery has been a target of intense interest in the last decade; we invite the readers to [DR16, CC18, CLC19] for recent overviews, and limit our discussions to the most relevant literature in the sequel.

Nonsmooth formulations for low-rank matrix recovery. Nonsmooth objective functions, such as the least absolute deviations, have been adopted earlier in both convex and nonconvex formulations of low-rank matrix recovery, including phase retrieval [Han17, DDP17, QZEW17, ZZLC17, DR19], blind deconvolution [Día19], quadratic sampling [LSC17, CL16, CCD⁺21, BL20], low-rank matrix sensing [CCD⁺21, Li13, WGM13, LZSV20], robust synchronization [WS13], to name a few. Our work is most closely related to and generalizes the vanilla subgradient method in [CCD⁺21], by establishing novel performance guarantees of *scaled* subgradient methods for robust low-rank matrix recovery.

Scaled first-order methods for low-rank matrix recovery. Variants of the scaled gradient methods are proposed in [MAS12, TW16, TMC20] for minimizing the least-squares formulation (2), where strong statistical and computational complexities are first established in [TMC20]. To the best of our knowledge, the current paper is the first work that provides rigorous statistical and computational guarantees for scaled subgradient methods for addressing nonsmooth formulations. When it comes to problems with heavy-tail observations such as quadratic sampling, while it is possible to establish faster convergence rates of vanilla gradient descent over the smooth least-squares loss function through a tailored analysis [MWCC19, LMCC21] via leave-one-out arguments, it is unclear if similar analyses are viable for scaled gradient methods (**ScaledGD**) in [TMC20]. Unfortunately, a direct application of the performance guarantee of **ScaledGD** on minimizing the smooth least-squares loss function leads to a much slower rate in terms of the problem dimension (see Table 1) for quadratic sampling. In contrast, our analysis for scaled subgradient methods yields strong guarantees in a more straightforward manner since the nonsmooth loss function has much better geometric properties [CCD⁺21].

Robust low-rank matrix recovery via nonconvex optimization. A pleasant side benefit of nonsmooth formulations is the added robustness to adversarial outliers under a simple algorithm design – the low-rank factors are updated essentially in the same manner regardless of the presence of outliers. In comparison, other nonconvex methods based on smooth formulations often need to introduce some special treatments to mitigate outliers before updating the low-rank factors, e.g. truncation or thresholding [ZCL16, LCZL20, LZSV20], which can be cumbersome to tune properly.

Condition number independent rate of convergence. It is well-known that first-order methods such as gradient descent exhibit poor scaling with respect to the condition number of the low-rank matrix. Possible remedies include alternating least-squares in the factored space [JNS13,HW14], or spectral methods over the matrix space [JMD10]. However, these approaches either require the inversion of a large matrix or a higher memory footprint, compared with the scaled first-order methods adopted herein.

1.3 Paper organization and notation

The rest of this paper is organized as follows. Section 2 describes the proposed scaled subgradient method and its connections to existing methods. Section 3 provides the theoretical guarantees for the scaled subgradient method in terms of both statistical and computational complexities, which are then instantiated to robust low-rank matrix sensing and quadratic sampling. Section 4 illustrates the superior empirical performance of the proposed method. Finally, we conclude in Section 5. The proofs are deferred to the appendix.

Notation. Throughout the paper, we use boldfaced symbols for vectors and matrices. For a vector \mathbf{v} , we use $\|\mathbf{v}\|_p$ to denote its ℓ_p norm. For any matrix \mathbf{A} , we use $\sigma_i(\mathbf{A})$ to denote its i -th largest singular value, and let $\mathbf{A}_{i\cdot}$ and $\mathbf{A}_{\cdot j}$ denote its i -th row and j -th column, respectively. In addition, $\|\mathbf{A}\|_{\text{op}}$ and $\|\mathbf{A}\|_F$ denote the spectral norm and the Frobenius norm of a matrix \mathbf{A} , respectively. For matrices \mathbf{A}, \mathbf{B} of the same size, we use $\langle \mathbf{A}, \mathbf{B} \rangle = \sum_{i,j} \mathbf{A}_{i,j} \mathbf{B}_{i,j} = \text{tr}(\mathbf{A}^\top \mathbf{B})$ to denote their inner product, where $\text{tr}(\cdot)$ denotes the trace. The set of invertible matrices in $\mathbb{R}^{r \times r}$ is denoted by $\text{GL}(r)$.

2 Problem Formulation and Algorithms

In this section, we formulate the low-rank matrix recovery problem, followed by a detailed description of the proposed scaled subgradient method.

2.1 Problem formulation

Let $\mathbf{X}_\star \in \mathbb{R}^{n_1 \times n_2}$ be the ground truth rank- r matrix, whose compact singular value decomposition (SVD) is given by

$$\mathbf{X}_\star = \mathbf{U}_\star \mathbf{\Sigma}_\star \mathbf{V}_\star^\top, \quad (5)$$

where $\mathbf{U}_\star \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{V}_\star \in \mathbb{R}^{n_2 \times r}$ are composed of r left and right singular vectors, respectively, and $\mathbf{\Sigma}_\star \in \mathbb{R}^{r \times r}$ is a diagonal matrix consisting of r singular values of \mathbf{X}_\star organized in a non-increasing order, i.e. $\sigma_1(\mathbf{X}_\star) \geq \dots \geq \sigma_r(\mathbf{X}_\star) > 0$. The condition number of \mathbf{X}_\star is thus defined as

$$\kappa := \sigma_1(\mathbf{X}_\star) / \sigma_r(\mathbf{X}_\star). \quad (6)$$

Without loss of generality, we define the ground truth low-rank factors as

$$\mathbf{L}_\star := \mathbf{U}_\star \mathbf{\Sigma}_\star^{1/2}, \quad \text{and} \quad \mathbf{R}_\star := \mathbf{V}_\star \mathbf{\Sigma}_\star^{1/2}, \quad (7)$$

so that $\mathbf{X}_\star = \mathbf{L}_\star \mathbf{R}_\star^\top$. Moreover, we denote the ground truth stacked factor matrix as

$$\mathbf{F}_\star := [\mathbf{L}_\star^\top, \mathbf{R}_\star^\top]^\top \in \mathbb{R}^{(n_1+n_2) \times r}. \quad (8)$$

Assume that we have access to a number of observations $\mathbf{y} = \{y_i\}_{i=1}^m$ of \mathbf{X}_\star , given as

$$y_i = \mathcal{A}_i(\mathbf{X}_\star) + w_i + s_i, \quad 1 \leq i \leq m, \quad (9)$$

or equivalently,

$$\mathbf{y} = \mathcal{A}(\mathbf{X}_\star) + \mathbf{w} + \mathbf{s}, \quad (10)$$

where $\mathcal{A}(\mathbf{X}_\star) = \{\mathcal{A}_i(\mathbf{X}_\star)\}_{i=1}^m$ is the measurement ensemble, $\mathbf{w} = \{w_i\}_{i=1}^m$ denotes the bounded noise, and $\mathbf{s} = \{s_i\}_{i=1}^m$ models arbitrary corruptions. The goal of low-rank matrix recovery is to reconstruct \mathbf{X}_\star from the noisy and corrupted observations \mathbf{y} in a statistically and computationally efficient manner.

2.2 Scaled subgradient method

Consider the following nonsmooth and nonconvex optimization problem over the factors

$$\underset{\mathbf{L} \in \mathbb{R}^{n_1 \times r}, \mathbf{R} \in \mathbb{R}^{n_2 \times r}}{\text{minimize}} \quad \mathcal{L}(\mathbf{L}, \mathbf{R}) := f(\mathbf{L}\mathbf{R}^\top), \quad (11)$$

where $f(\cdot)$ is a nonsmooth surrogate of the observation residuals. Of particular interest is the residual sum of absolute errors, defined as

$$f(\mathbf{X}) := \|\mathcal{A}(\mathbf{X}) - \mathbf{y}\|_1. \quad (12)$$

Correspondingly, the minimizer is called the least absolute deviations (LAD) solution.

Let us denote the stacked factor matrix in the t -th iterate as $\mathbf{F}_t := [\mathbf{L}_t^\top, \mathbf{R}_t^\top]^\top$. Given an initialization $\mathbf{F}_0 = [\mathbf{L}_0^\top, \mathbf{R}_0^\top]^\top$, the proposed scaled subgradient method (ScaledSM) proceeds as

$$\begin{aligned} \mathbf{L}_{t+1} &:= \mathbf{L}_t - \eta_t \mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1}, \\ \mathbf{R}_{t+1} &:= \mathbf{R}_t - \eta_t \mathbf{S}_t^\top \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1}, \end{aligned} \quad (13)$$

where $\mathbf{S}_t \in \partial f(\mathbf{L}_t \mathbf{R}_t^\top)$ is a subgradient of $f(\cdot)$ at $\mathbf{L}_t \mathbf{R}_t^\top$ (and hence $\mathbf{S}_t \mathbf{R}_t \in \partial_{\mathbf{L}} \mathcal{L}(\mathbf{L}_t, \mathbf{R}_t)$ and $\mathbf{S}_t^\top \mathbf{L}_t \in \partial_{\mathbf{R}} \mathcal{L}(\mathbf{L}_t, \mathbf{R}_t)$), and $\eta_t > 0$ is some properly selected stepsize, which we discuss next.

Stepsize schedules. We consider the following two choices of stepsize schedules:

- If we know the optimal value $f(\mathbf{X}_*)$, we can invoke the following Polyak's stepsize, given by

$$\eta_t^{\text{P}} := \frac{f(\mathbf{L}_t \mathbf{R}_t^\top) - f(\mathbf{X}_*)}{\|\mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1/2}\|_{\mathbb{F}}^2 + \|\mathbf{S}_t^\top \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1/2}\|_{\mathbb{F}}^2}, \quad (14)$$

where the denominator is the squared norm of the subgradient under a scaled metric concerted with the preconditioners. This schedule is implementable, for example, when the observations are noise-free, leading to $f(\mathbf{X}_*) = 0$. However, when the observations are noisy and corrupted, it is not viable to know $f(\mathbf{X}_*)$ beforehand.

- In general, we can apply the geometrically decaying stepsize originally introduced in [Gof77], given by

$$\eta_t^{\text{G}} := \frac{\lambda q^t}{\sqrt{\|\mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1/2}\|_{\mathbb{F}}^2 + \|\mathbf{S}_t^\top \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1/2}\|_{\mathbb{F}}^2}}, \quad (15)$$

where the denominator is similarly scaled as (14), and $\lambda > 0$ and $q \in (0, 1)$ are some parameters to be specified. This choice is broadly applicable when dealing with noisy and corrupted observations.

Compared with the vanilla subgradient method, which proceeds according to

$$\begin{aligned} \mathbf{L}_{t+1} &:= \mathbf{L}_t - \eta_t \mathbf{S}_t \mathbf{R}_t, \\ \mathbf{R}_{t+1} &:= \mathbf{R}_t - \eta_t \mathbf{S}_t^\top \mathbf{L}_t, \end{aligned} \quad (16)$$

the update rule (13) scales the subgradient $\mathbf{S}_t \mathbf{R}_t$ and $\mathbf{S}_t^\top \mathbf{L}_t$ by $(\mathbf{R}_t^\top \mathbf{R}_t)^{-1}$ and $(\mathbf{L}_t^\top \mathbf{L}_t)^{-1}$, respectively; see [TMC20] for its counterpart in smooth problems. An important highlight of the scaled subgradient method is that the update rule is covariant with respect to the ambiguity of low-rank matrix factorization. To see this, imagine that we modify the t -th updates as

$$\tilde{\mathbf{L}}_t := \mathbf{L}_t \mathbf{Q}, \quad \tilde{\mathbf{R}}_t := \mathbf{R}_t \mathbf{Q}^{-\top} \quad (17)$$

for some invertible matrix $\mathbf{Q} \in \text{GL}(r)$. It is easy to check:

- (i) both the Polyak's stepsize (14) and the geometrically decaying stepsize (15) do not change, since

$$\|\mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1/2}\|_{\mathbb{F}}^2 = \langle \mathbf{S}_t, \mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1} \mathbf{R}_t^\top \rangle = \langle \mathbf{S}_t, \mathbf{S}_t \tilde{\mathbf{R}}_t (\tilde{\mathbf{R}}_t^\top \tilde{\mathbf{R}}_t)^{-1} \tilde{\mathbf{R}}_t^\top \rangle = \|\mathbf{S}_t \tilde{\mathbf{R}}_t (\tilde{\mathbf{R}}_t^\top \tilde{\mathbf{R}}_t)^{-1/2}\|_{\mathbb{F}}^2,$$

which holds similarly for $\|\mathbf{S}_t^\top \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1/2}\|_{\mathbb{F}}^2$;

(ii) The next $(t + 1)$ -th iterate can be written as

$$\tilde{\mathbf{L}}_{t+1} = \tilde{\mathbf{L}}_t - \eta_t \mathbf{S}_t \tilde{\mathbf{R}}_t (\tilde{\mathbf{R}}_t^\top \tilde{\mathbf{R}}_t)^{-1} = [\mathbf{L}_t - \eta_t \mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1}] \mathbf{Q} = \mathbf{L}_{t+1} \mathbf{Q},$$

and similarly $\tilde{\mathbf{R}}_{t+1} = \mathbf{R}_{t+1} \mathbf{Q}^{-\top}$. Therefore, all the iterates are covariant with respect to the invertible transform (17).

Remark 1 (Comparison with ScaledGD). Although not our focus, it is instructive to consider the resulting update rule using the nonsmooth ℓ_2 -loss function $f(\mathbf{X}) = \|\mathcal{A}(\mathbf{X}) - \mathbf{y}\|_2$ (which has been studied in [CCD⁺21]), whose subgradient is given by

$$\mathbf{S}_t = \frac{\mathcal{A}^*(\mathbf{r}_t)}{\|\mathbf{r}_t\|_2},$$

where $\mathcal{A}^*(\cdot)$ is the adjoint operator of $\mathcal{A}(\cdot)$, and $\mathbf{r}_t := \mathcal{A}(\mathbf{L}_t \mathbf{R}_t^\top) - \mathbf{y}$ is the residual using the t -th iterate. Consequently, the scaled subgradient method follows the update rule

$$\begin{aligned} \mathbf{L}_{t+1} &= \mathbf{L}_t - \frac{\eta_t}{\|\mathbf{r}_t\|_2} \mathcal{A}^*(\mathbf{r}_t) \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1}, \\ \mathbf{R}_{t+1} &= \mathbf{R}_t - \frac{\eta_t}{\|\mathbf{r}_t\|_2} \mathcal{A}^*(\mathbf{r}_t)^\top \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1}, \end{aligned}$$

for some stepsize η_t . Careful readers might realize that this coincides with the update rule of ScaledGD in [TMC20] when optimizing the smooth squared ℓ_2 -loss function $g(\mathbf{X}) = \frac{1}{2} \|\mathcal{A}(\mathbf{X}) - \mathbf{y}\|_2^2$, except with an adaptive stepsize $\frac{\eta_t}{\|\mathbf{r}_t\|_2}$. Under the same assumption on $\mathcal{A}(\cdot)$ in [TMC20], the convergence behaviors of ScaledSM applied on $f(\mathbf{X})$ match that of ScaledGD on $g(\mathbf{X})$.

Remark 2 (ScaledSM for PSD matrices). When the low-rank matrix of interest is positive semi-definite (PSD), we factorize the matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ as $\mathbf{X} = \mathbf{L} \mathbf{L}^\top$, with $\mathbf{L} \in \mathbb{R}^{n \times r}$. The update rule of ScaledSM simplifies to

$$\mathbf{L}_{t+1} = \mathbf{L}_t - \eta_t \mathbf{S}_t \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1},$$

where $\mathbf{S}_t \in \partial f(\mathbf{L}_t \mathbf{L}_t^\top)$ is a subgradient of $f(\cdot)$ at $\mathbf{L}_t \mathbf{L}_t^\top$. Our theory applies to this PSD case in a straightforward manner.

3 Theoretical Guarantees

In this section, we first provide the theoretical guarantees of the scaled subgradient method under general geometric assumptions on $f(\cdot)$, and then instantiate them to concrete problems including robust low-rank matrix sensing and quadratic sampling.

3.1 Geometric assumptions

We start by introducing the following geometric properties of the loss function $f(\cdot)$, which play a key role in the convergence analysis.

The first condition is similar to the usual Lipschitz property of a function.

Definition 1 (Restricted Lipschitz continuity). A function $f(\cdot) : \mathbb{R}^{n_1 \times n_2} \mapsto \mathbb{R}$ is said to be rank- r restricted L -Lipschitz continuous for some quantity $L > 0$ if

$$|f(\mathbf{X}_1) - f(\mathbf{X}_2)| \leq L \|\mathbf{X}_1 - \mathbf{X}_2\|_F$$

holds for any $\mathbf{X}_1, \mathbf{X}_2 \in \mathbb{R}^{n_1 \times n_2}$ such that $\mathbf{X}_1 - \mathbf{X}_2$ has rank at most $2r$.

The second geometric condition is akin to the (one-point) strong convexity of a function, with the key difference that strong convexity adopts the squared Euclidean norm whereas the following one uses the plain Euclidean norm.

Definition 2 (Restricted sharpness). A function $f(\cdot) : \mathbb{R}^{n_1 \times n_2} \mapsto \mathbb{R}$ is said to be rank- r restricted μ -sharp w.r.t. \mathbf{X}_* for some $\mu > 0$ if

$$f(\mathbf{X}) - f(\mathbf{X}_*) \geq \mu \|\mathbf{X} - \mathbf{X}_*\|_{\text{F}}$$

holds for any $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with rank at most r .

For notational simplicity, if a function f is both restricted L -Lipschitz continuous and μ -sharp, we denote

$$\chi_f := L/\mu. \quad (18)$$

In some cases, e.g. in the presence of noise, the loss function $f(\cdot)$ only satisfies an approximate restricted sharpness property, which is detailed below.

Definition 3 (Approximate restricted sharpness). A function $f(\cdot) : \mathbb{R}^{n_1 \times n_2} \mapsto \mathbb{R}$ is said to be ξ -approximate rank- r restricted μ -sharp for some $\mu, \xi > 0$ if

$$f(\mathbf{X}) - f(\mathbf{X}_*) \geq \mu \|\mathbf{X} - \mathbf{X}_*\|_{\text{F}} - \xi$$

holds for any $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with rank at most r .

As shall be seen in Section 3.3, these conditions can be ensured for proper choices of the loss function as long as the observation operator $\mathcal{A}(\cdot)$ satisfies certain mixed-norm RIP, which holds for a wide number of practical problems.

3.2 Main results

Motivated by [TMC20], we measure the performance of $\mathbf{F} = [\mathbf{L}^\top, \mathbf{R}^\top]^\top$ using the following error metric

$$\text{dist}^2(\mathbf{F}, \mathbf{F}_*) := \inf_{\mathbf{Q} \in \text{GL}(r)} \left\| (\mathbf{L}\mathbf{Q} - \mathbf{L}_*)\boldsymbol{\Sigma}_*^{1/2} \right\|_{\text{F}}^2 + \left\| (\mathbf{R}\mathbf{Q}^{-\top} - \mathbf{R}_*)\boldsymbol{\Sigma}_*^{1/2} \right\|_{\text{F}}^2, \quad (19)$$

which takes into consideration both the representational ambiguity of the factorization up to invertible transforms and the scaling effect of preconditioners. In comparison, the more standard distance metric [MLC21] in the analysis of vanilla gradient methods reads as follows

$$\text{dist}_{\text{u}}^2(\mathbf{F}, \mathbf{F}_*) := \inf_{\mathbf{Q} \in \text{GL}(r)} \|\mathbf{L}\mathbf{Q} - \mathbf{L}_*\|_{\text{F}}^2 + \|\mathbf{R}\mathbf{Q}^{-\top} - \mathbf{R}_*\|_{\text{F}}^2,$$

which is inadequate to delineate the power of preconditioning. See [TMC20] for more discussions.

We start with stating the linear convergence of the scaled subgradient method when $f(\cdot)$ satisfies both the rank- r restricted L -Lipschitz continuity and μ -sharpness. The proof is deferred to Appendix B.

Theorem 1 (Scaled subgradient method with exact convergence). *Suppose that $f(\mathbf{X}) : \mathbb{R}^{n_1 \times n_2} \mapsto \mathbb{R}$ is convex in \mathbf{X} , and satisfies rank- r restricted L -Lipschitz continuity and μ -sharpness (cf. Definitions 1 and 2). In addition, suppose that the initialization \mathbf{F}_0 satisfies*

$$\text{dist}(\mathbf{F}_0, \mathbf{F}_*) \leq 0.02\sigma_r(\mathbf{X}_*)/\chi_f, \quad (20)$$

and the scaled subgradient method in (13) adopts either Polyak's stepsizes in (14) or geometrically decaying stepsizes in (15) with $\lambda = \sqrt{\frac{\sqrt{2}-1}{2}}0.02\sigma_r(\mathbf{X}_*)/\chi_f^2$ and $q = \sqrt{1 - 0.16/\chi_f^2}$. Then for all $t \geq 0$, the iterates satisfy

$$\begin{aligned} \text{dist}(\mathbf{F}_t, \mathbf{F}_*) &\leq (1 - 0.16/\chi_f^2)^{t/2} 0.02\sigma_r(\mathbf{X}_*)/\chi_f, \quad \text{and} \\ \|\mathbf{L}_t \mathbf{R}_t^\top - \mathbf{X}_*\|_{\text{F}} &\leq (1 - 0.16/\chi_f^2)^{t/2} 0.03\sigma_r(\mathbf{X}_*)/\chi_f. \end{aligned}$$

Theorem 1 shows that the iterates of the scaled subgradient method converges at a linear rate; to reach ϵ -accuracy, i.e. $\|\mathbf{L}_t \mathbf{R}_t^\top - \mathbf{X}_*\|_{\text{F}} \leq \epsilon\sigma_r(\mathbf{X}_*)$, it takes at most $O(\chi_f^2 \log \frac{1}{\epsilon})$ iterations, which, importantly, is independent of the condition number κ of \mathbf{X}_* . In addition, it is still possible to ensure approximate reconstruction when only the approximate restricted sharpness property holds, as shown in the next theorem. Again, we postpone the proof to Appendix C.

Theorem 2 (Scaled subgradient method with approximate convergence). *Suppose that $f(\cdot) : \mathbb{R}^{n_1 \times n_2} \mapsto \mathbb{R}$ is convex, and satisfies rank- r restricted L -Lipschitz continuity and ξ -approximate μ -sharpness (cf. Definitions 1 and 3) for some $\xi \leq 10^{-3} \sigma_r(\mathbf{X}_*) \mu / \chi_f$. Suppose that the initialization \mathbf{F}_0 satisfies $\text{dist}(\mathbf{F}_0, \mathbf{F}_*) \leq 0.02 \sigma_r(\mathbf{X}_*) / \chi_f$, and the scaled subgradient method adopts geometrically decaying stepsizes (15) with $\lambda = \sqrt{\frac{\sqrt{2}-1}{2}} 0.02 \sigma_r(\mathbf{X}_*) / \chi_f^2$ and $q = \sqrt{1 - 0.13 / \chi_f^2}$. Then for all $t \geq 0$, the iterates satisfy*

$$\begin{aligned} \text{dist}(\mathbf{F}_t, \mathbf{F}_*) &\leq \max \left\{ (1 - 0.13 / \chi_f^2)^{t/2} 0.02 \sigma_r(\mathbf{X}_*) / \chi_f, 20 \xi / \mu \right\}, \quad \text{and} \\ \|\mathbf{L}_t \mathbf{R}_t^\top - \mathbf{X}_*\|_{\text{F}} &\leq \max \left\{ (1 - 0.13 / \chi_f^2)^{t/2} 0.03 \sigma_r(\mathbf{X}_*) / \chi_f, 30 \xi / \mu \right\}. \end{aligned}$$

Theorem 2 shows that as long as the relaxation parameter ξ is sufficiently small, i.e. $\xi \lesssim \sigma_r(\mathbf{X}_*) \mu / \chi_f$, then the scaled subgradient method with geometrically decaying stepsizes converges at a linear rate until an error floor is hit. In particular, the iterates satisfy $\|\mathbf{L}_t \mathbf{R}_t^\top - \mathbf{X}_*\|_{\text{F}} \leq 30 \xi / \mu$ after at most $O(\chi_f^2)$ iterations up to logarithmic factors.

Remark 3. For simplicity of exposition, we have fixed the values of λ and q for the geometrically decaying stepsizes in the above theorems. It is possible to allow a wider range of λ and q by slightly modifying the arguments without sacrificing the linear convergence. In practice, these parameters should be tuned in order to yield optimal performance.

3.3 A case study: robust low-rank matrix recovery

We now apply the above theorems to robust low-rank matrix recovery, which showcases the superior performance of the scaled subgradient method.

Noise-free case. We start with the observation model (10) with clean measurements, i.e. $\mathbf{w} = \mathbf{0}$ and $\mathbf{s} = \mathbf{0}$. To proceed, we assume that the observation operator $\mathcal{A}(\cdot)$ satisfies the following mixed-norm RIP.

Definition 4 (mixed-norm RIP [RFP10, CCG15, CCD⁺21]). The linear map $\mathcal{A}(\cdot)$ is said to obey the rank- $2r$ mixed-norm RIP with constants δ_1, δ_2 if for all matrices $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2}$ of rank at most $2r$, one has

$$\delta_1 \|\mathbf{M}\|_{\text{F}} \leq \|\mathcal{A}(\mathbf{M})\|_1 \leq \delta_2 \|\mathbf{M}\|_{\text{F}}.$$

The next proposition verifies that the loss function (12) satisfies restricted Lipschitz continuity and sharpness properties under the mixed-norm RIP.

Proposition 1. *If $\mathcal{A}(\cdot)$ satisfies rank- $2r$ mixed-norm RIP with constants (δ_1, δ_2) , then $f(\mathbf{X}) = \|\mathcal{A}(\mathbf{X}) - \mathbf{y}\|_1 = \|\mathcal{A}(\mathbf{X} - \mathbf{X}_*)\|_1$ in (12) satisfies the rank- r restricted L -Lipschitz continuity and μ -sharpness with*

$$L = \delta_2, \quad \text{and} \quad \mu = \delta_1.$$

Proof. See Appendix D. □

With the geometric characterization of $f(\cdot)$ in place, we immediately have the following corollary that captures the performance of the scaled subgradient method when $\mathcal{A}(\cdot)$ satisfies the mixed-norm RIP.

Corollary 1. *If $\mathcal{A}(\cdot)$ satisfies rank- $2r$ mixed-norm RIP with (δ_1, δ_2) , then the scaled subgradient method over the loss function $f(\mathbf{X}) = \|\mathcal{A}(\mathbf{X}) - \mathbf{y}\|_1$ using either Polyak's or geometrically decaying stepsizes achieves $\|\mathbf{L}_t \mathbf{R}_t^\top - \mathbf{X}_*\|_{\text{F}} \leq \epsilon \sigma_r(\mathbf{X}_*)$ in $O\left(\frac{\delta_2^2}{\delta_1^2} \log \frac{1}{\epsilon}\right)$ iterations as long as the initialization satisfies $\text{dist}(\mathbf{F}_0, \mathbf{F}_*) \leq \frac{0.02 \delta_1}{\delta_2} \sigma_r(\mathbf{X}_*)$.*

Noisy and corrupted case. We now consider the observation model (10) where the noise \mathbf{w} is bounded with $\|\mathbf{w}\|_1 \leq \sigma_w$ and $\|\mathbf{s}\|_0 = p_s m$, where $p_s \in [0, 1/2)$ is the fraction of outliers. Following [CCD⁺21], we further introduce another important property of $\mathcal{A}(\cdot)$.

Definition 5 (\mathcal{S} -outlier bound [CCD⁺21]). The linear map $\mathcal{A}(\cdot)$ is said to obey the rank- $2r$ \mathcal{S} -outlier bound w.r.t. a set \mathcal{S} with a constant δ_3 if for all matrices $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2}$ of rank at most $2r$, one has

$$\delta_3 \|\mathbf{M}\|_{\text{F}} \leq \|\mathcal{A}_{\mathcal{S}^c}(\mathbf{M})\|_1 - \|\mathcal{A}_{\mathcal{S}}(\mathbf{M})\|_1,$$

where $\mathcal{A}_{\mathcal{S}}(\mathbf{M}) = \{\mathcal{A}_i(\mathbf{M})\}_{i \in \mathcal{S}}$ and $\mathcal{A}_{\mathcal{S}^c}(\mathbf{M}) = \{\mathcal{A}_i(\mathbf{M})\}_{i \in \mathcal{S}^c}$.

The next proposition verifies that the loss function in (12) satisfies restricted Lipschitz continuity and approximate sharpness properties under the mixed-norm RIP (cf. Definition 4) and the \mathcal{S} -outlier bound (cf. Definition 5).

Proposition 2 (Matrix sensing with outliers). *Denote the support of the outlier \mathbf{s} as \mathcal{S} . Suppose that $\mathcal{A}(\cdot)$ satisfies rank- $2r$ mixed-norm RIP with (δ_1, δ_2) and \mathcal{S} -outlier bound with δ_3 , then $f(\mathbf{X})$ in (12) satisfies rank- r restricted L -Lipschitz continuity and ξ -approximate μ -sharpness with*

$$L = \delta_2, \quad \mu = \delta_3, \quad \text{and} \quad \xi = 2\sigma_w. \quad (21)$$

Proof. See Appendix E. □

Similar to the previous noise-free case, this immediately leads to performance guarantees of the scaled subgradient method when $\mathcal{A}(\cdot)$ satisfies both the mixed-norm RIP and the \mathcal{S} -outlier bound.

Corollary 2. *If $\mathcal{A}(\cdot)$ satisfies rank- $2r$ mixed-norm RIP with (δ_1, δ_2) and \mathcal{S} -outlier bound with δ_3 , and $\|\mathbf{w}\|_1 \leq \sigma_w \leq 10^{-3} \sigma_r(\mathbf{X}_\star) \delta_3^2 / \delta_2$, then the scaled subgradient method over the loss function $f(\mathbf{X}) = \|\mathcal{A}(\mathbf{X}) - \mathbf{y}\|_1$ using the geometrically decaying stepsizes achieves $\|\mathbf{L}_t \mathbf{R}_t^\top - \mathbf{X}_\star\|_{\text{F}} \leq \max\{\epsilon \sigma_r(\mathbf{X}_\star), 60\sigma_w / \delta_3\}$ in $O\left(\frac{\delta_2^2}{\delta_3^2} \log \frac{1}{\epsilon}\right)$ iterations as long as the initialization satisfies $\text{dist}(\mathbf{F}_0, \mathbf{F}_\star) \leq \frac{0.02\delta_3}{\delta_2} \sigma_r(\mathbf{X}_\star)$.*

We now instantiate the above general guarantee to the following two types of observation operators. For simplicity, we assume there is no dense noise, i.e. $\sigma_w = 0$; see Table 1 for a summary of the comparisons.

- *matrix sensing:* the measurement operator $\mathcal{A}_i(\cdot)$ is defined as $\mathcal{A}_i(\mathbf{X}_\star) = \frac{1}{m} \langle \mathbf{A}_i, \mathbf{X}_\star \rangle$, where the matrix \mathbf{A}_i is composed of i.i.d. Gaussian entries $\mathcal{N}(0, 1)$.² It is shown in [CCD⁺21] (see also [LZSV20]) that $\mathcal{A}(\cdot)$ satisfies the mixed-norm RIP and \mathcal{S} -outlier bound with

$$\delta_1 \gtrsim 1, \quad \delta_2 \lesssim 1, \quad \delta_3 \gtrsim 1 - 2p_s,$$

as long as $m \gtrsim \frac{(n_1 + n_2)r}{(1 - 2p_s)^2} \log\left(\frac{1}{1 - 2p_s}\right)$. Hence, the scaled subgradient method converges linearly to ϵ -accuracy in $O\left(\frac{1}{(1 - 2p_s)^2} \log \frac{1}{\epsilon}\right)$ iterations provided that it is initialized properly, making it robust simultaneously to ill-conditioning of the matrix \mathbf{X}_\star and the presence of the outliers.

- *quadratic sampling:* the measurement operator $\mathcal{A}_i(\cdot)$ is defined as $\mathcal{A}_i(\mathbf{X}_\star) = \frac{1}{m} \langle \mathbf{a}_i \mathbf{a}_i^\top, \mathbf{X}_\star \rangle$, where $\mathbf{X}_\star \in \mathbb{R}^{n \times n}$ is PSD and the vector \mathbf{a}_i is composed of i.i.d. Gaussian entries $\mathcal{N}(0, 1)$. It is shown in [CCD⁺21] that $\mathcal{A}(\cdot)$ satisfies the mixed-norm RIP and \mathcal{S} -outlier bound with

$$\delta_1 \gtrsim 1, \quad \delta_2 \lesssim \sqrt{r}, \quad \delta_3 \gtrsim 1 - 2p_s,$$

as long as $m \gtrsim \frac{nr^2}{(1 - 2p_s)^2} \log\left(\frac{\sqrt{r}}{1 - 2p_s}\right)$. Hence, the scaled subgradient method converges linearly to ϵ -accuracy in $O\left(\frac{r}{(1 - 2p_s)^2} \log \frac{1}{\epsilon}\right)$ iterations, as long as it is seeded with a good initialization. In comparison, the iteration complexity of the scaled gradient descent method over the least-squares loss function depends polynomially with respect to n , due to the heavy-tailed nature of the observation operator, let alone its sensitivity to the outliers.

Remark 4 (Initialization). The above discussions are limited to the local iteration complexity, assuming a good initialization satisfying (20) is available. In the absence of outliers, a standard spectral method can be used, as shown in [TMC20]. In the presence of outliers, a truncated spectral method could be used; see e.g. [ZCL16, LCZL20].

²The same guarantee also holds for sub-Gaussian measurements.

4 Numerical Experiments

In this section, we conduct numerical experiments to corroborate our theory.

Comparisons of ScaledSM and VanillaSM. Since the vanilla subgradient method (VanillaSM) has been extensively benchmarked against other methods and established as state-of-the-art in [CCD⁺21], we focus on comparing the proposed scaled subgradient method (ScaledSM) to VanillaSM in the sequel. In general, the geometrically decaying stepsize (15) is a more practical choice than the Polyak’s stepsize (14), especially in the presence of noise and outliers. Nonetheless, using properly tuned geometrically decaying stepsizes essentially matches the performance of using Polyak’s stepsizes, for both VanillaSM [LZSV20] and ScaledSM, the latter of which we shall illustrate in the ensuing experiments. As such, we adopt Polyak’s stepsizes in the comparisons below, to emulate the scenario where both methods are tuned to operate under its largest allowable stepsizes and achieve the fastest convergence. In addition, both algorithms start from the same initialization.

We consider two low-rank matrix estimation tasks discussed in Section 3.3. Recall the observation model in (10) and its entrywise version in (9), which we repeat below for convenience:

$$y_i = \mathcal{A}_i(\mathbf{X}_\star) + w_i + s_i, \quad 1 \leq i \leq m.$$

In both tasks, the noise entry w_i is composed of i.i.d. entries uniformly drawn from $[-\frac{\sigma_w}{m}, \frac{\sigma_w}{m}]$. The outlier $s_i = \bar{s}_i \Omega_i$ is a sparse vector where Ω_i is a Bernoulli random variable with probability $p_s \in [0, 1/2)$, and \bar{s}_i is drawn uniformly at random from $[-10\|\mathcal{A}(\mathbf{X}_\star)\|_\infty, 10\|\mathcal{A}(\mathbf{X}_\star)\|_\infty]$. For ease of presentation, we assume that $\mathbf{X}_\star \in \mathbb{R}^{n \times n}$ is a square matrix with rank as r . We collect $m = 8nr$ measurements using the following respective measurement models. The signal-to-noise ratio is defined as $\text{SNR} := 20 \log_{10} \frac{\|\mathcal{A}(\mathbf{X}_\star)\|_1}{\sigma_w}$ in dB.

- *Matrix sensing.* Here, the measurement operator $\mathcal{A}_i(\cdot)$ is defined as $\mathcal{A}_i(\mathbf{X}_\star) = \frac{1}{m} \langle \mathbf{A}_i, \mathbf{X}_\star \rangle$, where the matrix \mathbf{A}_i is composed of i.i.d. Gaussian entries $\mathcal{N}(0, 1)$. The ground truth matrix \mathbf{X}_\star is generated via its compact SVD $\mathbf{X}_\star = \mathbf{U}_\star \mathbf{\Sigma}_\star \mathbf{V}_\star^\top$, where $\mathbf{U}_\star \in \mathbb{R}^{n \times r}$ is generated as the orthonormal basis vectors of an $n \times r$ matrix with i.i.d. Rademacher entries, $\mathbf{\Sigma}_\star$ is a diagonal matrix with the diagonal entries linearly distributed from 1 to κ , and $\mathbf{V}_\star \in \mathbb{R}^{n \times r}$ is generated in a similar fashion to \mathbf{U}_\star .
- *Quadratic sampling.* Here, the measurement operator $\mathcal{A}_i(\cdot)$ is defined as $\mathcal{A}_i(\mathbf{X}_\star) = \frac{1}{m} \langle \mathbf{a}_i \mathbf{a}_i^\top, \mathbf{X}_\star \rangle$, where \mathbf{a}_i is composed of i.i.d. Gaussian entries $\mathcal{N}(0, 1)$. The ground truth matrix \mathbf{X}_\star is positive semi-definite, and is generated via its compact SVD $\mathbf{X}_\star = \mathbf{U}_\star \mathbf{\Sigma}_\star \mathbf{U}_\star^\top$, where \mathbf{U}_\star and $\mathbf{\Sigma}_\star$ are generated in the same manner described above.

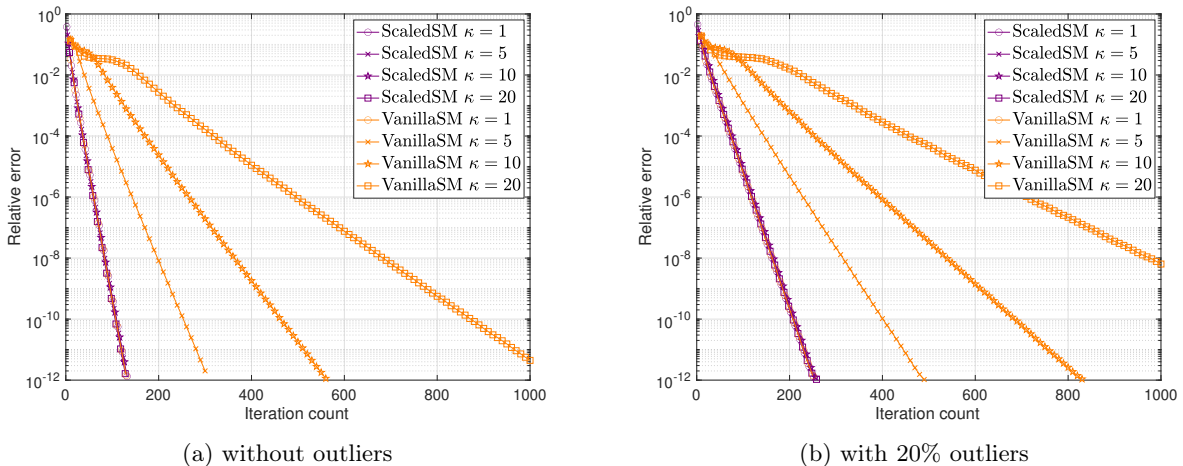


Figure 1: Performance comparisons of ScaledSM and VanillaSM for matrix sensing without or with outliers under different condition numbers $\kappa = 1, 5, 10, 20$, where $n = 100$, $r = 10$, and $m = 8nr$.

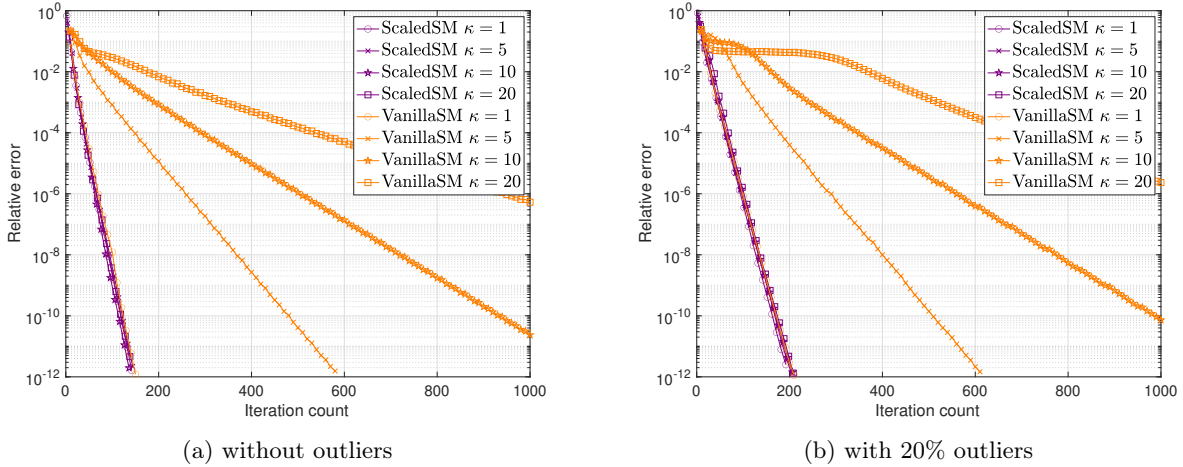


Figure 2: Performance comparisons of **ScaledSM** and **VanillaSM** for quadratic sampling without or with outliers under different condition numbers $\kappa = 1, 5, 10, 20$, where $n = 100$, $r = 5$, and $m = 8nr$.

Denote the index set of the remaining measurements after discarding p_s fraction with largest amplitudes as $\mathcal{I} = \{i : |y_i| \leq |\mathbf{y}|_{([p_s m])}\}$, where $|\mathbf{y}|_{(k)}$ denotes the k th largest amplitude of \mathbf{y} . The truncated spectral method in [ZCL16, LCZL20] is used for initialization, where we apply the standard spectral method only on the subset \mathcal{I} of the measurements. For matrix sensing, it follows the prescription in [LCZL20], and for quadratic sampling, it follows [LMCC21].

Fig. 1 shows the relative reconstruction error $\|\mathbf{X}_t - \mathbf{X}_*\|_F / \|\mathbf{X}_*\|_F$ for matrix sensing without outliers (in (a)) and with 20% outliers (i.e. $p_s = 0.2$ in (b)) under different condition numbers κ , where \mathbf{X}_t is the estimated low-rank matrix at the t -th iteration. Fig. 2 shows the relative reconstruction error for quadratic sampling under the same setting. It can be seen that **ScaledSM** is insensitive to κ and converges as a fast rate that is independent with κ , while the convergence of **VanillaSM** slows down dramatically with the increase of κ . In addition, both algorithms still converge linearly in the presence of outliers, thanks to the robustness of the least absolute deviations.

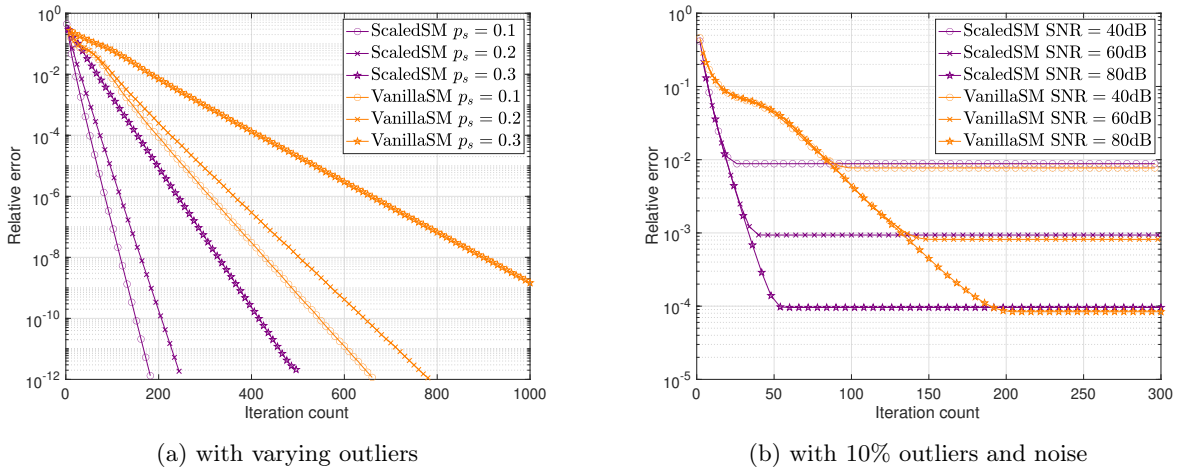


Figure 3: Performance comparisons of **ScaledSM** and **VanillaSM** for matrix sensing under different noise and outlier models, where $n = 100$, $r = 10$, $m = 8nr$, and $\kappa = 10$.

Fig. 3 further examines the impact of the amount of outliers and noise on the convergence speed in matrix sensing with a fixed condition number $\kappa = 10$, where Fig. 3 (a) illustrates the convergence speed at varying

amounts of outliers $p_s = 0.1, 0.2, 0.3$ respectively, and Fig. 3 (b) illustrates the convergence with $p_s = 0.1$ and additional bounded noise with varying SNR = 40, 60, 80dB. Similarly, Fig. 4 shows the same plots for quadratic sampling under the same setting. It can be seen that the convergence rate of **ScaledSM** slows down with the increase of outliers, which is again, consistent with the theory. Furthermore, the reconstruction is robust in the presence of additional bounded noise, where both **ScaledSM** and **VanillaSM** converge to the same accuracy that is proportional to the noise level, with **ScaledSM** converging at a faster speed.

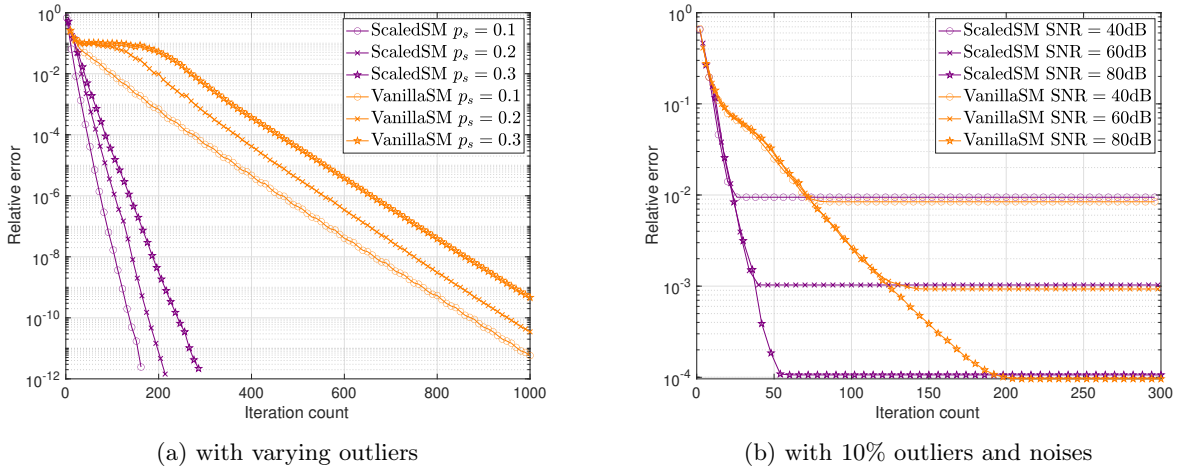


Figure 4: Performance comparisons of **ScaledSM** and **VanillaSM** for quadratic sampling under different noise and outlier models, where $n = 100$, $r = 5$, $m = 8nr$, and $\kappa = 10$.

Comparisons of stepsize schedules. We now compare the geometrically decaying stepsize with the Polyak’s stepsize for **ScaledSM**, which essentially mirrors similar experiments conducted in [LZSV20] for **VanillaSM**. We run **ScaledSM** for at most $T = 1000$ iterations, and stop early if the relative error achieves 10^{-12} . Fig. 5 and Fig. 6 show the performance comparisons of **ScaledSM** under various stepsize schedules for matrix sensing and quadratic sampling, respectively. For both figures, (a) shows the final relative error of **ScaledSM** using geometrically decaying stepsizes under various (λ, q) , where we see that **ScaledSM** converges as long as λ is not too large and q is not too small. We further plot the relative error versus the iteration count for **ScaledSM** using geometrically decaying stepsizes with a fixed q and various λ in (b), and with a fixed λ and various q in (c), where the performance using Polyak’s stepsizes is plotted for comparison. It can be seen that using Polyak’s stepsizes yields the fastest convergence. Indeed, if properly tuned, geometrically decaying stepsizes match Polyak’s stepsizes, as shown in (d). In general, we find that there is a wide range of parameters for geometrically decaying stepsizes where **ScaledSM** converges in a fast speed comparable to that of using Polyak’s stepsizes, as long as λ is not too large and q is not too small.

5 Discussions

This paper proposes scaled subgradient methods to minimize a family of nonsmooth and nonconvex formulations for low-rank matrix recovery—in particular, the residual sum of absolute errors—and guarantees its convergence at a rate that is almost dimension-free and independent of the condition number, even in the presence of corruptions. We illustrate the effectiveness of our approach by providing state-of-the-art performance guarantees for robust low-rank matrix sensing and quadratic sampling. In the future, it is of interest to study the performance of scaled subgradient methods for other signal estimation and statistical inference tasks, such as training student-teacher neural networks [DDKL20], as well as using random initializations [CCFM19].

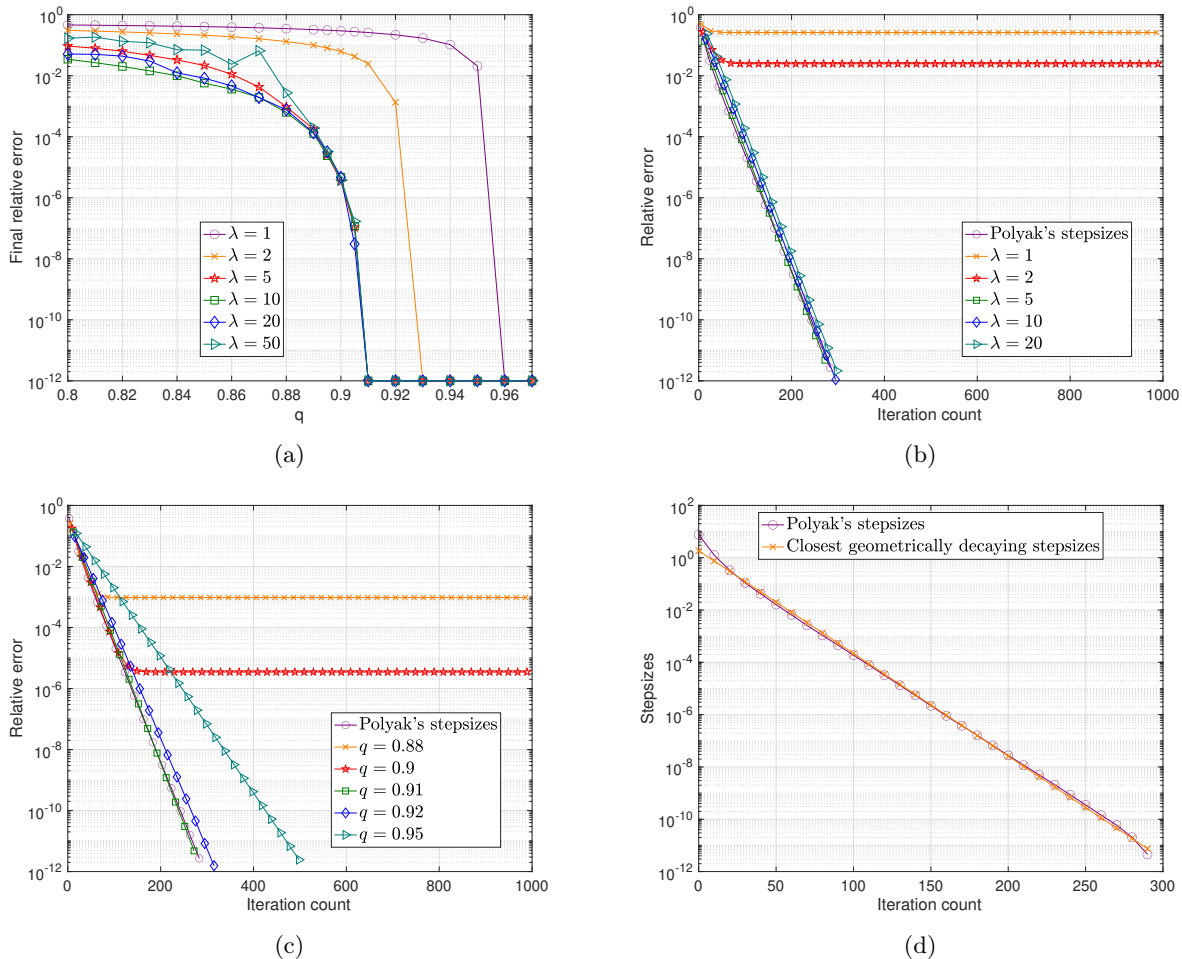


Figure 5: Performance comparisons of **ScaledSM** for matrix sensing using geometrically decaying stepsizes with parameters (λ, q) and Polyak's stepsizes, where we fix $n = 100$, $r = 10$, $m = 8nr$, $\kappa = 10$, and $p_s = 0.2$: (a) the final relative error for various combinations of (λ, q) , (b) the relative error versus iteration count for fixed $q = 0.91$ and varying λ , (c) the relative error versus iteration count for fixed $\lambda = 5$ and varying q , and (d) shows properly tuned geometrically decaying stepsizes with $\lambda = 1.85$ and $q = 0.91$ essentially match Polyak's stepsizes.

Acknowledgements

The work of T. Tong and Y. Chi is supported in part by ONR under the grants N00014-18-1-2142 and N00014-19-1-2404, by ARO under the grant W911NF-18-1-0303, and by NSF under the grants CAREER ECCS-1818571, CCF-1806154 and CCF-1901199.

References

- [BL20] S. Bahmani and K. Lee. Low-rank matrix estimation from rank-one projections by unlifted convex optimization. *arXiv preprint arXiv:2004.02718*, 2020.
- [CC18] Y. Chen and Y. Chi. Harnessing structures in big data via guaranteed low-rank matrix estimation: Recent theory and fast algorithms via convex and nonconvex optimization. *IEEE Signal Processing Magazine*, 35(4):14 – 31, 2018.

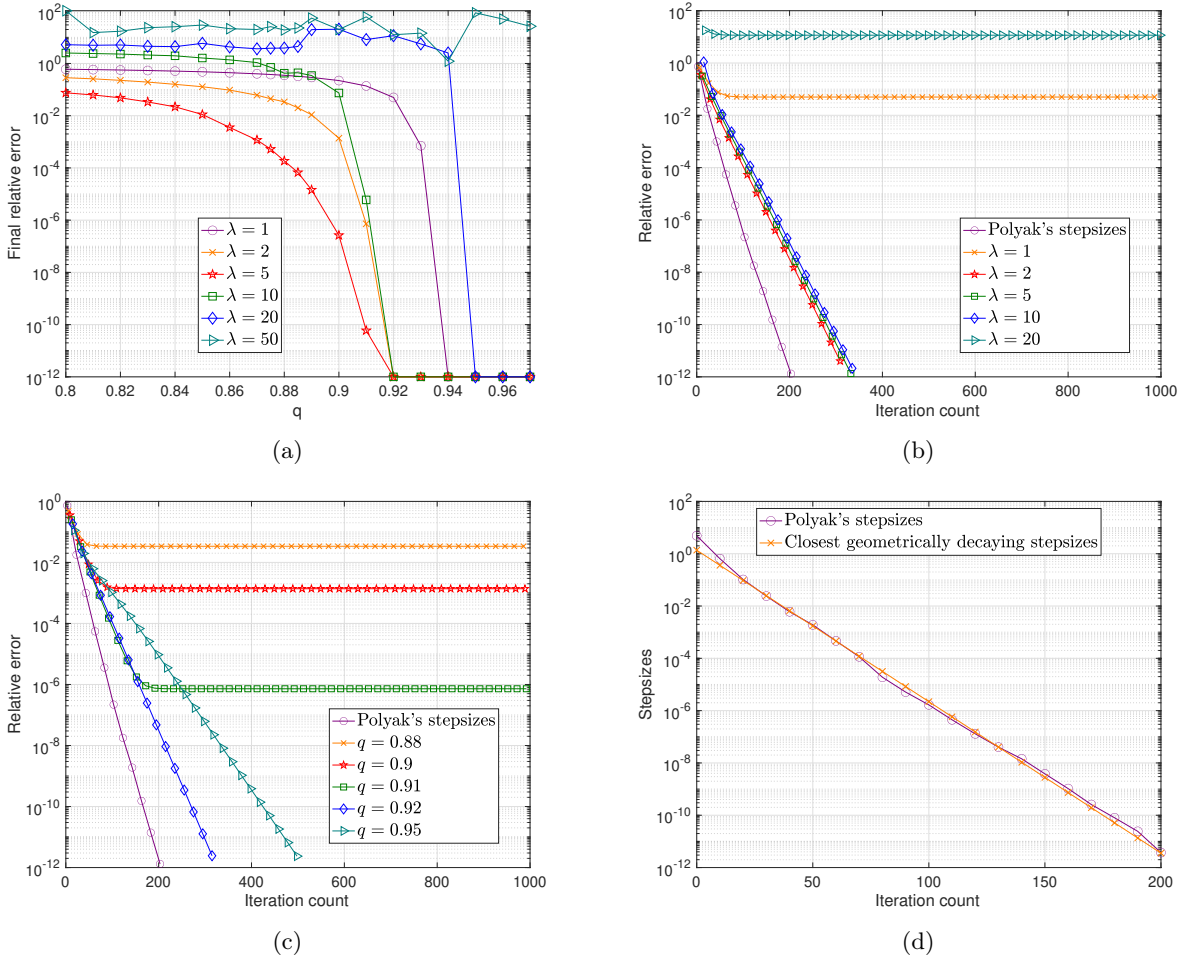


Figure 6: Performance comparisons of **ScaledSM** for quadratic sampling using geometrically decaying stepsizes with parameters (λ, q) and Polyak's stepsizes, where we fix $n = 100$, $r = 5$, $m = 8nr$, $\kappa = 10$, and $p_s = 0.2$: (a) the final relative error for various combinations of (λ, q) , (b) the relative error versus iteration count for fixed $q = 0.92$ and varying λ , (c) the relative error versus iteration count for fixed $\lambda = 2$ and varying q , and (d) shows properly tuned geometrically decaying stepsizes with $\lambda = 1.36$ and $q = 0.88$ essentially match Polyak's stepsizes.

- [CCD⁺21] V. Charisopoulos, Y. Chen, D. Davis, M. Díaz, L. Ding, and D. Drusvyatskiy. Low-rank matrix recovery with composite optimization: good conditioning and rapid convergence. *Foundations of Computational Mathematics*, pages 1–89, 2021.
- [CCFM19] Y. Chen, Y. Chi, J. Fan, and C. Ma. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 176(1-2):5–37, 2019.
- [CCG15] Y. Chen, Y. Chi, and A. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Transactions on Information Theory*, 61(7):4034–4059, July 2015.
- [CL16] Y. Chi and Y. M. Lu. Kaczmarz method for solving quadratic equations. *IEEE Signal Processing Letters*, 23(9):1183–1187, 2016.
- [CLC19] Y. Chi, Y. M. Lu, and Y. Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.

- [CLS15] E. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval via Wirtinger flow: Theory and algorithms. *Information Theory, IEEE Transactions on*, 61(4):1985–2007, 2015.
- [DDKL20] D. Davis, D. Drusvyatskiy, S. Kakade, and J. D. Lee. Stochastic subgradient method converges on tame functions. *Foundations of computational mathematics*, 20(1):119–154, 2020.
- [DDP17] D. Davis, D. Drusvyatskiy, and C. Paquette. The nonsmooth landscape of phase retrieval. *arXiv preprint arXiv:1711.03247*, 2017.
- [Día19] M. Díaz. The nonsmooth landscape of blind deconvolution. *arXiv preprint arXiv:1911.08526*, 2019.
- [DR16] M. A. Davenport and J. Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):608–622, 2016.
- [DR19] J. C. Duchi and F. Ruan. Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. *Information and Inference: A Journal of the IMA*, 8(3):471–529, 2019.
- [Gof77] J.-L. Goffin. On convergence rates of subgradient optimization methods. *Mathematical programming*, 13(1):329–347, 1977.
- [Han17] P. Hand. Phaselift is robust to a constant fraction of arbitrary errors. *Applied and Computational Harmonic Analysis*, 42(3):550–562, 2017.
- [HW14] M. Hardt and M. Wotter. Fast matrix completion without the condition number. In *Proceedings of The 27th Conference on Learning Theory*, pages 638–678, 2014.
- [JMD10] P. Jain, R. Meka, and I. S. Dhillon. Guaranteed rank minimization via singular value projection. In *Advances in Neural Information Processing Systems*, pages 937–945, 2010.
- [JNS13] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 665–674. ACM, 2013.
- [LCZL20] Y. Li, Y. Chi, H. Zhang, and Y. Liang. Non-convex low-rank matrix recovery with arbitrary outliers via median-truncated gradient descent. *Information and Inference: A Journal of the IMA*, 9(2):289–325, 2020.
- [Li13] X. Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Constructive Approximation*, 37(1):73–99, 2013.
- [LMCC21] Y. Li, C. Ma, Y. Chen, and Y. Chi. Nonconvex matrix factorization from rank-one measurements. *IEEE Transactions on Information Theory*, 67(3):1928–1950, 2021.
- [LSC17] Y. Li, Y. Sun, and Y. Chi. Low-rank positive semidefinite matrix recovery from corrupted rank-one measurements. *IEEE Transactions on Signal Processing*, 65(2):397–408, 2017.
- [LZSV20] X. Li, Z. Zhu, A. M.-C. So, and R. Vidal. Nonconvex robust low-rank matrix recovery. *SIAM Journal on Optimization*, 30(1):660–686, 2020.
- [MAS12] B. Mishra, K. A. Apuroop, and R. Sepulchre. A Riemannian geometry for low-rank matrix completion. *arXiv preprint arXiv:1211.1550*, 2012.
- [MLC21] C. Ma, Y. Li, and Y. Chi. Beyond Procrustes: Balancing-free gradient descent for asymmetric low-rank matrix sensing. *IEEE Transactions on Signal Processing*, 69:867–877, 2021.
- [MWCC19] C. Ma, K. Wang, Y. Chi, and Y. Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion, and blind deconvolution. *Foundations of Computational Mathematics*, pages 1–182, 2019.

- [QZEW17] Q. Qing, Y. Zhang, Y. Eldar, and J. Wright. Convolutional phase retrieval via gradient descent. *Neural Information Processing Systems*, 2017.
- [RFP10] 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.
- [SWW17] S. Sanghavi, R. Ward, and C. D. White. The local convexity of solving systems of quadratic equations. *Results in Mathematics*, 71(3-4):569–608, 2017.
- [TBS⁺16] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht. Low-rank solutions of linear matrix equations via Procrustes flow. In *International Conference Machine Learning*, pages 964–973, 2016.
- [TMC20] T. Tong, C. Ma, and Y. Chi. Accelerating ill-conditioned low-rank matrix estimation via scaled gradient descent. *arXiv preprint arXiv:2005.08898*, 2020.
- [TW16] J. Tanner and K. Wei. Low rank matrix completion by alternating steepest descent methods. *Applied and Computational Harmonic Analysis*, 40(2):417–429, 2016.
- [WGMM13] J. Wright, A. Ganesh, K. Min, and Y. Ma. Compressive principal component pursuit. *Information and Inference: A Journal of the IMA*, 2(1):32–68, 2013.
- [WS13] L. Wang and A. Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2(2):145–193, 2013.
- [ZCL16] H. Zhang, Y. Chi, and Y. Liang. Provable non-convex phase retrieval with outliers: Median truncated Wirtinger flow. In *International conference on machine learning*, pages 1022–1031, 2016.
- [ZZLC17] H. Zhang, Y. Zhou, Y. Liang, and Y. Chi. A nonconvex approach for phase retrieval: Reshaped Wirtinger flow and incremental algorithms. *Journal of Machine Learning Research*, 18(141):1–35, 2017.

A Technical Lemmas

This section gathers several useful lemmas that will be used in the proof. Throughout the appendix, we use \mathbf{X}_\star to denote the ground truth, with its compact SVD as $\mathbf{X}_\star = \mathbf{U}_\star \Sigma_\star \mathbf{V}_\star^\top$, and $\mathbf{F}_\star = \begin{bmatrix} \mathbf{L}_\star \\ \mathbf{R}_\star \end{bmatrix} = \begin{bmatrix} \mathbf{U}_\star \Sigma_\star^{1/2} \\ \mathbf{V}_\star \Sigma_\star^{1/2} \end{bmatrix}$. For any factor matrix $\mathbf{F} := \begin{bmatrix} \mathbf{L} \\ \mathbf{R} \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times r}$, we define the optimal alignment matrix \mathbf{Q} between \mathbf{F} and \mathbf{F}_\star as

$$\mathbf{Q} := \operatorname{argmin}_{\mathbf{Q} \in \text{GL}(r)} \left\| (\mathbf{L}\mathbf{Q} - \mathbf{L}_\star) \Sigma_\star^{1/2} \right\|_{\mathbf{F}}^2 + \left\| (\mathbf{R}\mathbf{Q}^{-\top} - \mathbf{R}_\star) \Sigma_\star^{1/2} \right\|_{\mathbf{F}}^2, \quad (22)$$

whenever the minimum is achieved.³

Lemma 1 ([TMC20]). *Fix any factor matrix $\mathbf{F} := \begin{bmatrix} \mathbf{L} \\ \mathbf{R} \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times r}$. Suppose that*

$$\operatorname{dist}(\mathbf{F}, \mathbf{F}_\star) = \sqrt{\inf_{\mathbf{Q} \in \text{GL}(r)} \left\| (\mathbf{L}\mathbf{Q} - \mathbf{L}_\star) \Sigma_\star^{1/2} \right\|_{\mathbf{F}}^2 + \left\| (\mathbf{R}\mathbf{Q}^{-\top} - \mathbf{R}_\star) \Sigma_\star^{1/2} \right\|_{\mathbf{F}}^2} < \sigma_r(\mathbf{X}_\star),$$

then the minimizer of the above minimization problem is attained at some $\mathbf{Q} \in \text{GL}(r)$, i.e. the optimal alignment matrix \mathbf{Q} between \mathbf{F} and \mathbf{F}_\star exists.

³If there exist multiple minimizers, we arbitrarily choose one as \mathbf{Q} .

Lemma 2 ([TMC20]). For any factor matrix $\mathbf{F} := \begin{bmatrix} \mathbf{L} \\ \mathbf{R} \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times r}$, the following relation holds

$$\text{dist}(\mathbf{F}, \mathbf{F}_\star) \leq \sqrt{\sqrt{2} + 1} \|\mathbf{L}\mathbf{R}^\top - \mathbf{X}_\star\|_{\mathbb{F}}.$$

Lemma 3 ([TMC20]). For any $\mathbf{L} \in \mathbb{R}^{n_1 \times r}$, $\mathbf{R} \in \mathbb{R}^{n_2 \times r}$, denote $\Delta_L := \mathbf{L} - \mathbf{L}_\star$ and $\Delta_R := \mathbf{R} - \mathbf{R}_\star$. Suppose that $\max\{\|\Delta_L \Sigma_\star^{-1/2}\|_{\text{op}}, \|\Delta_R \Sigma_\star^{-1/2}\|_{\text{op}}\} < 1$, then one has

$$\left\| \mathbf{L}(\mathbf{L}^\top \mathbf{L})^{-1} \Sigma_\star^{1/2} \right\|_{\text{op}} \leq \frac{1}{1 - \|\Delta_L \Sigma_\star^{-1/2}\|_{\text{op}}}; \quad (23a)$$

$$\left\| \mathbf{R}(\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_\star^{1/2} \right\|_{\text{op}} \leq \frac{1}{1 - \|\Delta_R \Sigma_\star^{-1/2}\|_{\text{op}}}; \quad (23b)$$

$$\left\| \mathbf{L}(\mathbf{L}^\top \mathbf{L})^{-1} \Sigma_\star^{1/2} - \mathbf{U}_\star \right\|_{\text{op}} \leq \frac{\sqrt{2} \|\Delta_L \Sigma_\star^{-1/2}\|_{\text{op}}}{1 - \|\Delta_L \Sigma_\star^{-1/2}\|_{\text{op}}}; \quad (23c)$$

$$\left\| \mathbf{R}(\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_\star^{1/2} - \mathbf{V}_\star \right\|_{\text{op}} \leq \frac{\sqrt{2} \|\Delta_R \Sigma_\star^{-1/2}\|_{\text{op}}}{1 - \|\Delta_R \Sigma_\star^{-1/2}\|_{\text{op}}}. \quad (23d)$$

Lemma 4 ([TMC20]). Recall the partial Frobenius norm

$$\|\mathbf{X}\|_{\mathbb{F},r} := \sqrt{\sum_{i=1}^r \sigma_i^2(\mathbf{X})} = \|\mathcal{P}_r(\mathbf{X})\|_{\mathbb{F}}, \quad (24)$$

where $\mathcal{P}_r(\mathbf{X})$ is the best rank- r approximation of \mathbf{X} . For any $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ and $\mathbf{R} \in \mathbb{R}^{n_2 \times r}$, one has

$$\|\mathbf{X}\mathbf{R}\|_{\mathbb{F}} \leq \|\mathbf{X}\|_{\mathbb{F},r} \|\mathbf{R}\|_{\text{op}}. \quad (25)$$

In addition, for any $\mathbf{X}, \bar{\mathbf{X}} \in \mathbb{R}^{n_1 \times n_2}$ with $\text{rank}(\bar{\mathbf{X}}) \leq r$, one has

$$|\langle \mathbf{X}, \bar{\mathbf{X}} \rangle| \leq \|\mathbf{X}\|_{\mathbb{F},r} \|\bar{\mathbf{X}}\|_{\mathbb{F}}. \quad (26)$$

Lemma 5. Suppose that $f(\cdot) : \mathbb{R}^{n_1 \times n_2} \mapsto \mathbb{R}$ is convex and rank- r restricted L -Lipschitz continuous (cf. Definition 1). Then for any subgradient $\mathbf{S} \in \partial f(\mathbf{X})$, one has $\|\mathbf{S}\|_{\mathbb{F},r} \leq L$.

Proof. Fix any subgradient $\mathbf{S} \in \partial f(\mathbf{X})$. By the definition of a subgradient, for any $\tilde{\mathbf{X}} \in \mathbb{R}^{n_1 \times n_2}$, one has

$$f(\tilde{\mathbf{X}}) \geq f(\mathbf{X}) + \langle \mathbf{S}, \tilde{\mathbf{X}} - \mathbf{X} \rangle.$$

In particular, taking $\tilde{\mathbf{X}} = \mathbf{X} + \mathcal{P}_r(\mathbf{S})$ arrives at

$$f(\mathbf{X} + \mathcal{P}_r(\mathbf{S})) \geq f(\mathbf{X}) + \langle \mathbf{S}, \mathcal{P}_r(\mathbf{S}) \rangle = f(\mathbf{X}) + \|\mathbf{S}\|_{\mathbb{F},r}^2, \quad (27)$$

where the last equality follows from the definition (24). Note that $\mathcal{P}_r(\mathbf{S})$ has rank at most r . By the rank- r restricted L -Lipschitz continuity of $f(\cdot)$, we have

$$f(\mathbf{X} + \mathcal{P}_r(\mathbf{S})) - f(\mathbf{X}) \leq L \|\mathcal{P}_r(\mathbf{S})\|_{\mathbb{F}} = L \|\mathbf{S}\|_{\mathbb{F},r}.$$

Combining the above inequality with (27), we conclude $\|\mathbf{S}\|_{\mathbb{F},r} \leq L$. \square

B Proof of Theorem 1

Suppose that the t -th iterate \mathbf{F}_t obeys the condition

$$\text{dist}(\mathbf{F}_t, \mathbf{F}_\star) \leq 0.02\sigma_r(\mathbf{X}_\star)/\chi_f. \quad (28)$$

Lemma 1 ensures that \mathbf{Q}_t , the optimal alignment matrix between \mathbf{F}_t and \mathbf{F}_* exists. For notational convenience, we denote $\mathbf{L} := \mathbf{L}_t \mathbf{Q}_t$, $\mathbf{R} := \mathbf{R}_t \mathbf{Q}_t^{-\top}$, $\Delta_L := \mathbf{L} - \mathbf{L}_*$, $\Delta_R := \mathbf{R} - \mathbf{R}_*$, $\mathbf{S} := \mathbf{S}_t$, and $\epsilon := 0.02/\chi_f$. By the definition

$$\text{dist}(\mathbf{F}_t, \mathbf{F}_*) = \sqrt{\|\Delta_L \Sigma_*^{1/2}\|_{\mathbb{F}}^2 + \|\Delta_R \Sigma_*^{1/2}\|_{\mathbb{F}}^2} \quad (29)$$

and the relation $\|\mathbf{A}\mathbf{B}\|_{\mathbb{F}} \geq \|\mathbf{A}\|_{\mathbb{F}} \sigma_r(\mathbf{B}) \geq \|\mathbf{A}\|_{\sigma_r}(\mathbf{B})$, we have

$$\max\{\|\Delta_L \Sigma_*^{-1/2}\|_{\text{op}}, \|\Delta_R \Sigma_*^{-1/2}\|_{\text{op}}\} \leq \epsilon. \quad (30)$$

We start by relating $\|\mathbf{L}\mathbf{R}^\top - \mathbf{X}_*\|_{\mathbb{F}}$ to $\text{dist}(\mathbf{F}_t, \mathbf{F}_*)$ given (30). Applying the triangle inequality to the basic relation $\mathbf{L}\mathbf{R}^\top - \mathbf{X}_* = \mathbf{L}_t \mathbf{R}_t^\top - \mathbf{X}_* = \Delta_L \mathbf{R}_*^\top + \mathbf{L}_* \Delta_R^\top + \Delta_L \Delta_R^\top$, we have

$$\begin{aligned} \|\mathbf{L}\mathbf{R}^\top - \mathbf{X}_*\|_{\mathbb{F}} &\leq \|\Delta_L \mathbf{R}_*^\top\|_{\mathbb{F}} + \|\mathbf{L}_* \Delta_R^\top\|_{\mathbb{F}} + \|\Delta_L \Delta_R^\top\|_{\mathbb{F}} \\ &\leq \|\Delta_L \mathbf{R}_*^\top\|_{\mathbb{F}} + \|\mathbf{L}_* \Delta_R^\top\|_{\mathbb{F}} + \frac{1}{2} \|\Delta_L \Sigma_*^{-1/2}\|_{\text{op}} \|\Delta_R \Sigma_*^{1/2}\|_{\mathbb{F}} + \frac{1}{2} \|\Delta_L \Sigma_*^{1/2}\|_{\mathbb{F}} \|\Delta_R \Sigma_*^{-1/2}\|_{\text{op}} \\ &\leq \left(1 + \frac{1}{2} \max\{\|\Delta_L \Sigma_*^{-1/2}\|_{\text{op}}, \|\Delta_R \Sigma_*^{-1/2}\|_{\text{op}}\}\right) \left(\|\Delta_L \Sigma_*^{1/2}\|_{\mathbb{F}} + \|\Delta_R \Sigma_*^{1/2}\|_{\mathbb{F}}\right) \\ &\leq \left(1 + \frac{\epsilon}{2}\right) \sqrt{2} \sqrt{\|\Delta_L \Sigma_*^{1/2}\|_{\mathbb{F}}^2 + \|\Delta_R \Sigma_*^{1/2}\|_{\mathbb{F}}^2} \leq 1.5 \text{dist}(\mathbf{F}_t, \mathbf{F}_*), \end{aligned} \quad (31)$$

where the last line uses the basic inequality $\|\Delta_L \Sigma_*^{1/2}\|_{\mathbb{F}} + \|\Delta_R \Sigma_*^{1/2}\|_{\mathbb{F}} \leq \sqrt{2} \text{dist}(\mathbf{F}_t, \mathbf{F}_*)$ and (30).

From now on, we focus on proving the distance contraction. By the definition of $\text{dist}(\mathbf{F}_{t+1}, \mathbf{F}_*)$, one has

$$\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_*) \leq \left\| (\mathbf{L}_{t+1} \mathbf{Q}_t - \mathbf{L}_*) \Sigma_*^{1/2} \right\|_{\mathbb{F}}^2 + \left\| (\mathbf{R}_{t+1} \mathbf{Q}_t^{-\top} - \mathbf{R}_*) \Sigma_*^{1/2} \right\|_{\mathbb{F}}^2. \quad (32)$$

We expand the first square in (32) as

$$\begin{aligned} \left\| (\mathbf{L}_{t+1} \mathbf{Q}_t - \mathbf{L}_*) \Sigma_*^{1/2} \right\|_{\mathbb{F}}^2 &= \left\| (\mathbf{L} - \eta_t \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} - \mathbf{L}_*) \Sigma_*^{1/2} \right\|_{\mathbb{F}}^2 \\ &= \|\Delta_L \Sigma_*^{1/2}\|_{\mathbb{F}}^2 - 2\eta_t \langle \mathbf{S}, \Delta_L \Sigma_* (\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top \rangle + \eta_t^2 \left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_*^{1/2} \right\|_{\mathbb{F}}^2 \\ &= \|\Delta_L \Sigma_*^{1/2}\|_{\mathbb{F}}^2 - 2\eta_t \left\langle \mathbf{S}, \Delta_L \mathbf{R}_*^\top + \frac{1}{2} \Delta_L \Delta_R^\top \right\rangle + \eta_t^2 \underbrace{\left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_*^{1/2} \right\|_{\mathbb{F}}^2}_{\mathfrak{S}_1} \\ &\quad - 2\eta_t \underbrace{\left\langle \mathbf{S}, \Delta_L \Sigma_* (\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top - \Delta_L \mathbf{R}_*^\top - \frac{1}{2} \Delta_L \Delta_R^\top \right\rangle}_{\mathfrak{S}_2}, \end{aligned} \quad (33)$$

where in the first line, we used the fact that the update rule (13) is covariant with respect to \mathbf{Q}_t , implying that

$$\mathbf{L}_{t+1} \mathbf{Q}_t = \mathbf{L} - \eta_t \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1}.$$

We proceed to bound \mathfrak{S}_1 and \mathfrak{S}_2 . The term \mathfrak{S}_1 can be bounded by

$$\begin{aligned} \mathfrak{S}_1 &\leq \left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\mathbb{F}}^2 \left\| (\mathbf{R}^\top \mathbf{R})^{-1/2} \Sigma_*^{1/2} \right\|_{\text{op}}^2 \\ &\leq \left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\mathbb{F}}^2 \frac{1}{(1 - \epsilon)^2}, \end{aligned}$$

where the second line follows from the condition (30) and Lemma 3 (cf. (23b)):

$$\left\| (\mathbf{R}^\top \mathbf{R})^{-1/2} \Sigma_*^{1/2} \right\|_{\text{op}} = \left\| \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_*^{1/2} \right\|_{\text{op}} \leq \frac{1}{1 - \epsilon}.$$

For the term \mathfrak{S}_2 , note that

$$\Delta_L \Sigma_\star (\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top - \Delta_L \mathbf{R}_\star^\top - \frac{1}{2} \Delta_L \Delta_R^\top = \Delta_L \Sigma_\star^{1/2} \left(\mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_\star^{1/2} - \mathbf{V}_\star - \frac{1}{2} \Delta_R \Sigma_\star^{-1/2} \right)^\top$$

has rank at most r . Hence we can invoke Lemma 4 (cf. (26)) to obtain

$$\begin{aligned} |\mathfrak{S}_2| &\leq \|\mathbf{S}\|_{\mathbb{F},r} \left\| \Delta_L \Sigma_\star (\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top - \Delta_L \mathbf{R}_\star^\top - \frac{1}{2} \Delta_L \Delta_R^\top \right\|_{\mathbb{F}} \\ &\leq \|\mathbf{S}\|_{\mathbb{F},r} \|\Delta_L \Sigma_\star^{1/2}\|_{\mathbb{F}} \left(\left\| \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_\star^{1/2} - \mathbf{V}_\star \right\|_{\text{op}} + \frac{1}{2} \|\Delta_R \Sigma_\star^{-1/2}\|_{\text{op}} \right) \\ &\leq L \left(\frac{\sqrt{2}\epsilon}{1-\epsilon} + \frac{\epsilon}{2} \right) \|\Delta_L \Sigma_\star^{1/2}\|_{\mathbb{F}}, \end{aligned}$$

where the second line follows from the triangle inequality, and the third line follows from $\|\mathbf{S}\|_{\mathbb{F},r} \leq L$ (cf. Lemma 5), (30), and

$$\left\| \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \Sigma_\star^{1/2} - \mathbf{V}_\star \right\|_{\text{op}} \leq \frac{\sqrt{2}\epsilon}{1-\epsilon}$$

from Lemma 3 (cf. (23d)).

Plugging collectively the bounds for \mathfrak{S}_1 and \mathfrak{S}_2 into (33) yields

$$\begin{aligned} \left\| (\mathbf{L}_{t+1} \mathbf{Q}_t - \mathbf{L}_\star) \Sigma_\star^{1/2} \right\|_{\mathbb{F}}^2 &\leq \|\Delta_L \Sigma_\star^{1/2}\|_{\mathbb{F}}^2 - 2\eta_t \left\langle \mathbf{S}, \Delta_L \mathbf{R}_\star^\top + \frac{1}{2} \Delta_L \Delta_R^\top \right\rangle + \frac{\eta_t^2}{(1-\epsilon)^2} \left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\mathbb{F}}^2 \\ &\quad + \eta_t L \left(\frac{2\sqrt{2}\epsilon}{1-\epsilon} + \epsilon \right) \|\Delta_L \Sigma_\star^{1/2}\|_{\mathbb{F}}. \end{aligned}$$

Similarly, we can obtain the control of $\|(\mathbf{R}_{t+1} \mathbf{Q}_t^{-\top} - \mathbf{R}_\star) \Sigma_\star^{1/2}\|_{\mathbb{F}}^2$. Combine them together to reach

$$\begin{aligned} \text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) &\leq \|\Delta_L \Sigma_\star^{1/2}\|_{\mathbb{F}}^2 + \|\Delta_R \Sigma_\star^{1/2}\|_{\mathbb{F}}^2 - 2\eta_t \left\langle \mathbf{S}, \Delta_L \mathbf{R}_\star^\top + \mathbf{L}_\star \Delta_R^\top + \Delta_L \Delta_R^\top \right\rangle \\ &\quad + \frac{\eta_t^2}{(1-\epsilon)^2} \left(\left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\mathbb{F}}^2 + \left\| \mathbf{S}^\top \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1/2} \right\|_{\mathbb{F}}^2 \right) + \eta_t L \left(\frac{2\sqrt{2}\epsilon}{1-\epsilon} + \epsilon \right) \left(\|\Delta_L \Sigma_\star^{1/2}\|_{\mathbb{F}} + \|\Delta_R \Sigma_\star^{1/2}\|_{\mathbb{F}} \right). \end{aligned}$$

Using the subgradient optimality of \mathbf{S} , we obtain

$$\left\langle \mathbf{S}, \Delta_L \mathbf{R}_\star^\top + \mathbf{L}_\star \Delta_R^\top + \Delta_L \Delta_R^\top \right\rangle = \left\langle \mathbf{S}, \mathbf{L} \mathbf{R}^\top - \mathbf{X}_\star \right\rangle \geq f(\mathbf{L} \mathbf{R}^\top) - f(\mathbf{X}_\star),$$

together with (29), which further implies that

$$\begin{aligned} \text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) &\leq \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star) - 2\eta_t (f(\mathbf{L} \mathbf{R}^\top) - f(\mathbf{X}_\star)) \\ &\quad + \frac{\eta_t^2}{(1-\epsilon)^2} \left(\left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\mathbb{F}}^2 + \left\| \mathbf{S}^\top \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1/2} \right\|_{\mathbb{F}}^2 \right) + \eta_t L \left(\frac{4\epsilon}{1-\epsilon} + \sqrt{2}\epsilon \right) \text{dist}(\mathbf{F}_t, \mathbf{F}_\star), \end{aligned} \quad (34)$$

where the last term uses the basic inequality $\|\Delta_L \Sigma_\star^{1/2}\|_{\mathbb{F}} + \|\Delta_R \Sigma_\star^{1/2}\|_{\mathbb{F}} \leq \sqrt{2} \text{dist}(\mathbf{F}_t, \mathbf{F}_\star)$.

Before proceeding to different cases of stepsize schedules, we record two useful properties. First, by the restricted μ -sharpness of $f(\cdot)$ together with Lemma 2, we have

$$f(\mathbf{L} \mathbf{R}^\top) - f(\mathbf{X}_\star) \geq \mu \|\mathbf{L} \mathbf{R}^\top - \mathbf{X}_\star\|_{\mathbb{F}} \geq \mu \sqrt{\sqrt{2} - 1} \text{dist}(\mathbf{F}_t, \mathbf{F}_\star). \quad (35)$$

On the other end, by Lemma 4 (cf. (25)), we have

$$\begin{aligned} \left\| \mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\mathbb{F}}^2 + \left\| \mathbf{S}^\top \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1/2} \right\|_{\mathbb{F}}^2 &\leq \|\mathbf{S}\|_{\mathbb{F},r}^2 \left(\left\| \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\text{op}}^2 + \left\| \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1/2} \right\|_{\text{op}}^2 \right) \\ &\leq 2L^2, \end{aligned} \quad (36)$$

where the second line follows from $\|\mathbf{S}\|_{\mathbb{F},r} \leq L$ (cf. Lemma 5) and

$$\left\| \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2} \right\|_{\text{op}}^2 = \left\| \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top \right\|_{\text{op}} = 1, \quad \left\| \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1/2} \right\|_{\text{op}}^2 = \left\| \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1} \mathbf{L}^\top \right\|_{\text{op}} = 1.$$

B.1 Convergence with Polyak's stepsizes

Let $\eta_t = \eta_t^P$ be the Polyak's stepsize in (14), which is

$$\begin{aligned}\eta_t &= \frac{f(\mathbf{L}_t \mathbf{R}_t^\top) - f(\mathbf{X}_\star)}{\|\mathbf{S}_t \mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1/2}\|_{\mathbb{F}}^2 + \|\mathbf{S}_t^\top \mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1/2}\|_{\mathbb{F}}^2} \\ &= \frac{f(\mathbf{L} \mathbf{R}^\top) - f(\mathbf{X}_\star)}{\|\mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2}\|_{\mathbb{F}}^2 + \|\mathbf{S}^\top \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1/2}\|_{\mathbb{F}}^2},\end{aligned}\quad (37)$$

where the second line follows since $\mathbf{L}_t \mathbf{R}_t^\top = \mathbf{L} \mathbf{R}^\top$, $\mathbf{L}_t (\mathbf{L}_t^\top \mathbf{L}_t)^{-1} \mathbf{L}_t^\top = \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1} \mathbf{L}^\top$ and $\mathbf{R}_t (\mathbf{R}_t^\top \mathbf{R}_t)^{-1} \mathbf{R}_t^\top = \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top$. Plugging (37) into (34), we have

$$\begin{aligned}\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) &\leq \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star) - \eta_t \left(2 - \frac{1}{(1-\epsilon)^2}\right) (f(\mathbf{L} \mathbf{R}^\top) - f(\mathbf{X}_\star)) + \eta_t L \left(\frac{4\epsilon}{1-\epsilon} + \sqrt{2}\epsilon\right) \text{dist}(\mathbf{F}_t, \mathbf{F}_\star) \\ &\leq \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star) - \eta_t \mu \left(\sqrt{\sqrt{2}-1} \left(2 - \frac{1}{(1-\epsilon)^2}\right) - \chi_f \left(\frac{4\epsilon}{1-\epsilon} + \sqrt{2}\epsilon\right)\right) \text{dist}(\mathbf{F}_t, \mathbf{F}_\star),\end{aligned}\quad (38)$$

where the second line follows from (35) and $\chi_f = L/\mu$.

To continue, combining (35) and (36), we can lower bound the Polyak's stepsize (37) as

$$\eta_t \geq \frac{\sqrt{\sqrt{2}-1} \mu \text{dist}(\mathbf{F}_t, \mathbf{F}_\star)}{2L^2}.$$

This, combined with (38), leads to

$$\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq \rho(\epsilon, \chi_f) \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star),$$

where the contraction rate $\rho(\epsilon, \chi_f)$ is

$$\rho(\epsilon, \chi_f) := 1 - \frac{\sqrt{\sqrt{2}-1}}{2\chi_f^2} \left(\sqrt{\sqrt{2}-1} \left(2 - \frac{1}{(1-\epsilon)^2}\right) - \chi_f \left(\frac{4\epsilon}{1-\epsilon} + \sqrt{2}\epsilon\right)\right). \quad (39)$$

Under the condition $\epsilon = 0.02/\chi_f$, we calculate $(1 - \rho(\epsilon, \chi_f))\chi_f^2$ as

$$\begin{aligned}&\frac{\sqrt{\sqrt{2}-1}}{2} \left(\sqrt{\sqrt{2}-1} \left(2 - \frac{1}{(1-\epsilon)^2}\right) - \chi_f \left(\frac{4\epsilon}{1-\epsilon} + \sqrt{2}\epsilon\right)\right) \\ &\geq 0.32 \left(0.64 \times \left(2 - \frac{1}{0.98^2}\right) - 0.02 \left(\frac{4}{0.98} + \sqrt{2}\right)\right) \geq 0.16,\end{aligned}$$

thus $\rho(\epsilon, \chi_f) \leq 1 - 0.16/\chi_f^2$. We conclude that

$$\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq (1 - 0.16/\chi_f^2) \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star),$$

which is the desired claim.

B.2 Convergence with geometrically decaying stepsizes

Let $\eta_t = \eta_t^G$ be the geometrically decaying stepsize in (15), which is

$$\eta_t = \frac{\lambda q^t}{\sqrt{\|\mathbf{S} \mathbf{R} (\mathbf{R}^\top \mathbf{R})^{-1/2}\|_{\mathbb{F}}^2 + \|\mathbf{S}^\top \mathbf{L} (\mathbf{L}^\top \mathbf{L})^{-1/2}\|_{\mathbb{F}}^2}}.$$

Plugging the above into (34), we have

$$\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star) - \eta_t \mu \left(2\sqrt{\sqrt{2}-1} - \chi_f \left(\frac{4\epsilon}{1-\epsilon} + \sqrt{2}\epsilon\right)\right) \text{dist}(\mathbf{F}_t, \mathbf{F}_\star) + \frac{\lambda^2 q^{2t}}{(1-\epsilon)^2}$$

$$\leq \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star) - \frac{\lambda q^t}{\sqrt{2}\chi_f} \left(2\sqrt{\sqrt{2}-1} - \chi_f \left(\frac{4\epsilon}{1-\epsilon} + \sqrt{2}\epsilon \right) \right) \text{dist}(\mathbf{F}_t, \mathbf{F}_\star) + \frac{\lambda^2 q^{2t}}{(1-\epsilon)^2},$$

where the first line follows from (35) and $\chi_f = L/\mu$, and the second line follows from $\eta_t \geq \frac{\lambda q^t}{\sqrt{2}L}$ due to (36). We now aim to show that

$$\text{dist}(\mathbf{F}_t, \mathbf{F}_\star) \leq (1 - 0.16/\chi_f^2)^{t/2} 0.02\sigma_r(\mathbf{X}_\star)/\chi_f$$

in an inductive manner. Assume the above induction hypothesis holds at the t -iteration. By the setting of parameters, i.e.

$$\lambda q^t = \sqrt{\frac{\sqrt{2}-1}{2}} (1 - 0.16/\chi_f^2)^{t/2} 0.02\sigma_r(\mathbf{X}_\star)/\chi_f^2,$$

we have

$$\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq \rho(\epsilon, \chi_f) (1 - 0.16/\chi_f^2)^t (0.02\sigma_r(\mathbf{X}_\star)/\chi_f)^2,$$

where the contraction rate $\rho(\epsilon, \chi_f)$ matches exactly (39). Therefore, under the condition $\epsilon = 0.02/\chi_f$, we have $\rho(\epsilon, \chi_f) \leq 1 - 0.16/\chi_f^2$, thus we conclude that

$$\text{dist}(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq (1 - 0.16/\chi_f^2)^{\frac{t+1}{2}} 0.02\sigma_r(\mathbf{X}_\star)/\chi_f,$$

which is the desired claim.

C Proof of Theorem 2

We start by introducing the short-hand notation $d_t := (1 - 0.13/\chi_f^2)^{t/2} 0.02\sigma_r(\mathbf{X}_\star)/\chi_f$. The parameters are set as

$$\lambda q^t = \sqrt{\frac{\sqrt{2}-1}{2}} (1 - 0.13/\chi_f^2)^{t/2} 0.02\sigma_r(\mathbf{X}_\star)/\chi_f^2 = \sqrt{\frac{\sqrt{2}-1}{2}} \frac{d_t}{\chi_f}.$$

Therefore, the geometric stepsize

$$\eta_t = \frac{\lambda q^t}{\sqrt{\|\mathbf{S}\mathbf{R}(\mathbf{R}^\top \mathbf{R})^{-1/2}\|_F^2 + \|\mathbf{S}^\top \mathbf{L}(\mathbf{L}^\top \mathbf{L})^{-1/2}\|_F^2}},$$

in view of (36), satisfies

$$\eta_t \geq \frac{\lambda q^t}{\sqrt{2}L} = \frac{\sqrt{\sqrt{2}-1}}{2} \frac{d_t}{\chi_f^2 \mu}. \quad (40)$$

Follow the same derivations as the proof of Theorem 1 until (34). Plugging the stepsize (40) into (34), together with the approximate restricted sharpness property

$$f(\mathbf{L}\mathbf{R}^\top) - f(\mathbf{X}_\star) \geq \mu \|\mathbf{L}\mathbf{R}^\top - \mathbf{X}_\star\|_F - \xi \geq \sqrt{\sqrt{2}-1} \mu \text{dist}(\mathbf{F}_t, \mathbf{F}_\star) - \xi,$$

we have

$$\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star) - \eta_t \mu \left(2\sqrt{\sqrt{2}-1} - \left(\frac{4}{1-\epsilon} + \sqrt{2} \right) \epsilon \chi_f \right) \text{dist}(\mathbf{F}_t, \mathbf{F}_\star) + \frac{\lambda^2 q^{2t}}{(1-\epsilon)^2} + 2\eta_t \xi.$$

Under the conditions $\chi_f \geq 1$ and $\epsilon = 0.02/\chi_f \leq 0.02$, the above relation can be simplified to

$$\text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq \text{dist}^2(\mathbf{F}_t, \mathbf{F}_\star) - 1.177\eta_t \mu \text{dist}(\mathbf{F}_t, \mathbf{F}_\star) + \frac{0.216}{\chi_f^2} d_t^2 + 2\eta_t \xi. \quad (41)$$

We next prove the theorem by induction, where the base case is established trivially by the initial condition. By the induction hypothesis, the distance at the t -th iterate is bounded by

$$\text{dist}(\mathbf{F}_t, \mathbf{F}_\star) \leq \max\{d_t, 20\xi/\mu\}.$$

To obtain the control of $\text{dist}(\mathbf{F}_{t+1}, \mathbf{F}_\star)$, we split the discussion in two cases.

1. If $d_t \geq 20\xi/\mu$, or equivalently, $\xi \leq 0.05\mu d_t$, in view of (41), we have

$$\begin{aligned} \text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) &\stackrel{(i)}{\leq} d_t^2 - 1.177\eta_t\mu d_t + \frac{0.216}{\chi_f^2}d_t^2 + 0.1\eta_t\mu d_t \\ &= d_t^2 - 1.077\eta_t\mu d_t + \frac{0.216}{\chi_f^2}d_t^2 \\ &\stackrel{(ii)}{\leq} d_t^2 - \frac{0.346}{\chi_f^2}d_t^2 + \frac{0.216}{\chi_f^2}d_t^2 \\ &= (1 - 0.13/\chi_f^2)d_t^2, \end{aligned}$$

where (i) uses $\xi \leq 0.05\mu d_t$, and (ii) uses the condition (40). We conclude that $\text{dist}(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq (1 - 0.13/\chi_f^2)^{1/2}d_t$.

2. If $0 \leq d_t < 20\xi/\mu$, we have

$$\begin{aligned} \text{dist}^2(\mathbf{F}_{t+1}, \mathbf{F}_\star) &\leq \left(\frac{20\xi}{\mu}\right)^2 - 1.177\eta_t\mu \frac{20\xi}{\mu} + \frac{0.216}{\chi_f^2}d_t^2 + 2\eta_t\xi \\ &= \left(\frac{20\xi}{\mu}\right)^2 - 1.077\eta_t\mu \frac{20\xi}{\mu} + \frac{0.216}{\chi_f^2}d_t^2 \\ &\leq \left(\frac{20\xi}{\mu}\right)^2 - \frac{1.077\sqrt{\sqrt{2}-1}}{2\chi_f^2}d_t \frac{20\xi}{\mu} + \frac{0.216}{\chi_f^2}d_t^2 \\ &\leq \left(\frac{20\xi}{\mu}\right)^2 - 0.13d_t \frac{20\xi}{\mu} \\ &\leq \left(\frac{20\xi}{\mu}\right)^2, \end{aligned}$$

where the third line uses the condition (40), and the last line holds since $d_t \geq 0$.

In sum, we conclude

$$\text{dist}(\mathbf{F}_{t+1}, \mathbf{F}_\star) \leq \max\left\{(1 - 0.13/\chi_f^2)^{\frac{t+1}{2}} 0.02\sigma_r(\mathbf{X}_\star)/\chi_f, 20\xi/\mu\right\},$$

which is the desired claim.

D Proof of Proposition 1

For \mathbf{X}_1 and \mathbf{X}_2 where $\mathbf{X}_1 - \mathbf{X}_2$ has rank at most $2r$, we have

$$\begin{aligned} |f(\mathbf{X}_1) - f(\mathbf{X}_2)| &= \left| \|\mathcal{A}(\mathbf{X}_1 - \mathbf{X}_\star)\|_1 - \|\mathcal{A}(\mathbf{X}_2 - \mathbf{X}_\star)\|_1 \right| \\ &\leq \|\mathcal{A}(\mathbf{X}_1 - \mathbf{X}_2)\|_1 \leq \delta_2 \|\mathbf{X}_1 - \mathbf{X}_2\|_F, \end{aligned}$$

where the second line follows from the inverse triangle inequality and the assumed rank- $2r$ mixed-norm RIP (cf. Definition 4) of $\mathcal{A}(\cdot)$. As a result, we have $L = \delta_2$. On the other end, we note

$$f(\mathbf{X}) - f(\mathbf{X}_\star) = \|\mathcal{A}(\mathbf{X} - \mathbf{X}_\star)\|_1 \geq \delta_1 \|\mathbf{X} - \mathbf{X}_\star\|_F,$$

where the first equality uses $f(\mathbf{X}_\star) = 0$ and the second inequality follows from the rank- $2r$ mixed-norm RIP; thus $\mu = \delta_1$.

E Proof of Proposition 2

For \mathbf{X}_1 and \mathbf{X}_2 with $\text{rank}(\mathbf{X}_1 - \mathbf{X}_2) \leq 2r$, we have

$$\begin{aligned} |f(\mathbf{X}_1) - f(\mathbf{X}_2)| &= \left| \|\mathcal{A}(\mathbf{X}_1 - \mathbf{X}_*) - \mathbf{w} - \mathbf{s}\|_1 - \|\mathcal{A}(\mathbf{X}_2 - \mathbf{X}_*) - \mathbf{w} - \mathbf{s}\|_1 \right| \\ &\leq \|\mathcal{A}(\mathbf{X}_1 - \mathbf{X}_2)\|_1 \leq \delta_2 \|\mathbf{X}_1 - \mathbf{X}_2\|_F, \end{aligned}$$

where the second line follows from the inverse triangle inequality and the rank- $2r$ mixed-norm RIP; hence $L = \delta_2$. For approximate restricted sharpness, note that

$$\begin{aligned} f(\mathbf{X}) - f(\mathbf{X}_*) &= \|\mathcal{A}(\mathbf{X} - \mathbf{X}_*) - \mathbf{w} - \mathbf{s}\|_1 - \|\mathbf{w} + \mathbf{s}\|_1 \\ &\geq \|\mathcal{A}(\mathbf{X} - \mathbf{X}_*) - \mathbf{s}\|_1 - \|\mathbf{w}\|_1 - \|\mathbf{s}\|_1 - \|\mathbf{w}\|_1 \\ &= \|\mathcal{A}_{\mathcal{S}^c}(\mathbf{X} - \mathbf{X}_*)\|_1 + \|\mathcal{A}_{\mathcal{S}}(\mathbf{X} - \mathbf{X}_*) - \mathbf{s}\|_1 - \|\mathbf{s}\|_1 - 2\|\mathbf{w}\|_1 \\ &\geq \|\mathcal{A}_{\mathcal{S}^c}(\mathbf{X} - \mathbf{X}_*)\|_1 - \|\mathcal{A}_{\mathcal{S}}(\mathbf{X} - \mathbf{X}_*)\|_1 - 2\|\mathbf{w}\|_1 \\ &\geq \delta_3 \|\mathbf{X} - \mathbf{X}_*\|_F - 2\|\mathbf{w}\|_1 \\ &\geq \delta_3 \|\mathbf{X} - \mathbf{X}_*\|_F - 2\sigma_w, \end{aligned}$$

where the second and the fourth lines follow from the triangle inequality, the third line follows from the definition of \mathcal{S} , and the last line follows from the definition of the \mathcal{S} -outlier bound and the noise upper bound $\|\mathbf{w}\|_1 \leq \sigma_w$. Therefore, we have $\mu = \delta_3$ and $\xi = 2\sigma_w$.