

---

# Faster Projection-free Convex Optimization over the Spectrahedron

---

Dan Garber

Toyota Technological Institute at Chicago  
dgarber@ttic.edu

## Abstract

Minimizing a convex function over the spectrahedron, i.e., the set of all  $d \times d$  positive semidefinite matrices with unit trace, is an important optimization task with many applications in optimization, machine learning, and signal processing. It is also notoriously difficult to solve in large-scale since standard techniques require to compute expensive matrix decompositions. An alternative is the conditional gradient method (aka Frank-Wolfe algorithm) that regained much interest in recent years, mostly due to its application to this specific setting. The key benefit of the CG method is that it avoids expensive matrix decompositions all together, and simply requires a single eigenvector computation per iteration, which is much more efficient. On the downside, the CG method, in general, converges with an inferior rate. The error for minimizing a  $\beta$ -smooth function after  $t$  iterations scales like  $\beta/t$ . This rate does not improve even if the function is also strongly convex. In this work we present a modification of the CG method tailored for the spectrahedron. The per-iteration complexity of the method is essentially identical to that of the standard CG method: only a single eigenvector computation is required. For minimizing an  $\alpha$ -strongly convex and  $\beta$ -smooth function, the *expected* error of the method after  $t$  iterations is:

$$O\left(\min\left\{\frac{\beta}{t}, \left(\frac{\beta\sqrt{\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4}t}\right)^{4/3}, \left(\frac{\beta}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)t}\right)^2\right\}\right),$$

where  $\text{rank}(\mathbf{X}^*)$ ,  $\lambda_{\min}(\mathbf{X}^*)$  are the rank of the optimal solution and smallest non-zero eigenvalue, respectively. Beyond the significant improvement in convergence rate, it also follows that when the optimum is low-rank, our method provides better accuracy-rank tradeoff than the standard CG method. To the best of our knowledge, this is the first result that attains provably faster convergence rates for a CG variant for optimization over the spectrahedron. We also present encouraging preliminary empirical results.

## 1 Introduction

Minimizing a convex function over the set of positive semidefinite matrices with unit trace, aka the spectrahedron, is an important optimization task which lies at the heart of many optimization, machine learning, and signal processing tasks such as matrix completion [1, 13], metric learning [21, 22], kernel matrix learning [16, 9], multiclass classification [2, 23], and more.

Since modern applications are mostly of very large scale, first-order methods are the obvious choice to deal with this optimization problem. However, even these are notoriously difficult to apply, since most of the popular gradient schemes require the computation of an orthogonal projection on each iteration to enforce feasibility, which for the spectrahedron, amounts to computing a full eigen-decomposition

of a real symmetric matrix. Such a decomposition requires  $O(d^3)$  arithmetic operations for a  $d \times d$  matrix and thus is prohibitive for high-dimensional problems. An alternative is to use first-order methods that do not require expensive decompositions, but rely only on computationally-cheap leading eigenvector computations. These methods are mostly based on the conditional gradient method, also known as the Frank-Wolfe algorithm [3, 12], which is a generic method for constrained convex optimization given an oracle for minimizing linear functions over the feasible domain. Indeed, linear minimization over the spectrahedron amounts to a single leading eigenvector computation. While the CG method has been discovered already in the 1950's [3], it has regained much interest in recent years in the machine learning and optimization communities, in particular due to its applications to semidefinite optimization and convex optimization with a nuclear norm constraint / regularization<sup>1</sup>, e.g., [10, 13, 17, 19, 22, 2, 11]. This regained interest is not surprising: while a full eigen-decomposition for  $d \times d$  matrix requires  $O(d^3)$  arithmetic operations, leading eigenvector computations can be carried out, roughly speaking, in worst-case time that is only linear in the number of non-zeros in the input matrix multiplied by either  $\epsilon^{-1}$  for the popular Power Method or by  $\epsilon^{-1/2}$  for the more efficient Lanczos method, where  $\epsilon$  is the target accuracy. These running times improve exponentially to only depend on  $\log(1/\epsilon)$  when the eigenvalues of the input matrix are well distributed [14]. Indeed, in several important machine learning applications, such as matrix completion, the CG method requires eigenvector computations of very sparse matrices [13]. Also, very recently, new eigenvector algorithms with significantly improved performance guarantees were introduced which are applicable for matrices with certain popular structure [5, 8, 20].

The main drawback of the CG method is that its convergence rate is, in general, inferior compared to projection-based gradient methods. The convergence rate for minimizing a smooth function, roughly speaking, scales only like  $1/t$ . In particular, this rate does not improve even when the function is also strongly convex. On the other hand, the convergence rate of optimal projection-based methods, such as Nesterov's accelerated gradient method, scales like  $1/t^2$  for smooth functions, and can be improved exponentially to  $\exp(-\Theta(t))$  when the objective is also strongly convex.

Very recently, several successful attempts were made to devise natural modifications of the CG method that retain the overall low per-iteration complexity, while enjoying provably faster convergence rates, usually under a strong-convexity assumption, or a slightly weaker one. These results exhibit provably-faster rates for optimization over polyhedral sets [7, 15] and strongly-convex sets [6], but do not apply to the spectrahedron. For the specific setting considered in this work, several heuristic improvements of the CG method were suggested which show promising empirical evidence, however, non of them provably improve over the rate of the standard CG method [19, 17, 4].

In this work we present a new non-trivial variant of the CG method, which, to the best of our knowledge, is the first to exhibit provably faster convergence rates for optimization over the spectrahedron, under standard smoothness and strong convexity assumptions. The per-iteration complexity of the method is essentially identical to that of the standard CG method, i.e., only a single leading eigenvector computation per iteration is required. Our method is tailored for optimization over the spectrahedron, and can be seen as a certain hybridization of the standard CG method and the projected gradient method. From a high-level view, we take advantage of the fact that solving a  $\ell_2$ -regularized linear problem over the set of extreme points of the spectrahedron is equivalent to linear optimization over this set, i.e., amounts to a single eigenvector computation. We then show via a novel and non-trivial analysis, that includes new decomposition concepts for positive semidefinite matrices, that such an algorithmically-cheap regularization is sufficient, in presence of strong convexity, to derive faster convergence rates.

## 2 Preliminaries and Notation

For vectors we let  $\|\cdot\|$  denote the standard Euclidean norm, while for matrices we let  $\|\cdot\|$  denote the spectral norm,  $\|\cdot\|_F$  denote the Frobenius norm, and  $\|\cdot\|_*$  denote the nuclear norm. We denote by  $\mathbb{S}_d$  the space of  $d \times d$  real symmetric matrices, and by  $\mathcal{S}_d$  the *spectrahedron* in  $\mathbb{S}_d$ , i.e.,  $\mathcal{S}_d := \{\mathbf{X} \in \mathbb{S}_d \mid \mathbf{X} \succeq 0, \text{Tr}(\mathbf{X}) = 1\}$ . We let  $\text{Tr}(\cdot)$  and  $\text{rank}(\cdot)$  denote the trace and rank of a given matrix in  $\mathbb{S}_d$ , respectively. We let  $\bullet$  denote the standard inner-product for matrices. Given a matrix  $\mathbf{X} \in \mathcal{S}_d$ , we let  $\lambda_{\min}(\mathbf{X})$  denote the smallest non-zero eigenvalue of  $\mathbf{X}$ .

---

<sup>1</sup>minimizing a convex function subject to a nuclear norm constraint is efficiently reducible to the minimization of the function over the spectrahedron, as fully detailed in [13].

Given a matrix  $\mathbf{A} \in \mathbb{S}_d$ , we denote by  $\mathbf{EV}(\mathbf{A})$  an eigenvector of  $\mathbf{A}$  that corresponds to the largest (signed) eigenvalue of  $\mathbf{A}$ , i.e.,  $\mathbf{EV}(\mathbf{A}) \in \arg \max_{\mathbf{v}: \|\mathbf{v}\|=1} \mathbf{v}^\top \mathbf{A} \mathbf{v}$ . Given a scalar  $\xi > 0$ , we also denote by  $\mathbf{EV}_\xi(\mathbf{A})$  an  $\xi$ -approximation to the largest (in terms of eigenvalue) eigenvector of  $\mathbf{A}$ , i.e.,  $\mathbf{EV}_\xi(\mathbf{A})$  returns a unit vector  $\mathbf{v}$  such that  $\mathbf{v}^\top \mathbf{A} \mathbf{v} \geq \lambda_{\max}(\mathbf{A}) - \xi$ .

**Definition 1.** We say that a function  $f(\mathbf{X}) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$  is  $\alpha$ -strongly convex w.r.t. a norm  $\|\cdot\|$ , if for all  $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{m \times n}$  it holds that

$$f(\mathbf{Y}) \geq f(\mathbf{X}) + (\mathbf{Y} - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \frac{\alpha}{2} \|\mathbf{X} - \mathbf{Y}\|^2.$$

**Definition 2.** We say that a function  $f(\mathbf{X}) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$  is  $\beta$ -smooth w.r.t. a norm  $\|\cdot\|$ , if for all  $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{m \times n}$  it holds that

$$f(\mathbf{Y}) \leq f(\mathbf{X}) + (\mathbf{Y} - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \frac{\beta}{2} \|\mathbf{X} - \mathbf{Y}\|^2.$$

The first-order optimality condition implies that for a  $\alpha$ -strongly convex  $f$ , if  $\mathbf{X}^*$  is the unique minimizer of  $f$  over a convex set  $\mathcal{K} \subset \mathbb{R}^{m \times n}$ , then for all  $\mathbf{X} \in \mathcal{K}$  it holds that

$$f(\mathbf{X}) - f(\mathbf{X}^*) \geq \frac{\alpha}{2} \|\mathbf{X} - \mathbf{X}^*\|^2. \quad (1)$$

## 2.1 Problem setting

The main focus of this work is the following optimization problem:

$$\min_{\mathbf{X} \in \mathcal{S}_d} f(\mathbf{X}), \quad (2)$$

where we assume that  $f(\mathbf{X})$  is both  $\alpha$ -strongly convex and  $\beta$ -smooth w.r.t.  $\|\cdot\|_F$ . We denote the (unique) minimizer of  $f$  over  $\mathcal{S}_d$  by  $\mathbf{X}^*$ .

## 3 Our Approach

We begin by briefly describing the conditional gradient and projected-gradient methods, pointing out their advantages and short-comings for solving Problem (2) in Subsection 3.1. We then present our new method which is a certain combination of ideas from both methods in Subsection 3.2.

### 3.1 Conditional gradient and projected gradient descent

The standard conditional gradient algorithm is detailed below in Algorithm 1.

---

#### Algorithm 1 Conditional Gradient

---

- 1: input: sequence of step-sizes  $\{\eta_t\}_{t \geq 1} \subset [0, 1]$
  - 2: let  $\mathbf{X}_1$  be an arbitrary matrix in  $\mathcal{S}_d$
  - 3: **for**  $t = 1 \dots$  **do**
  - 4:    $\mathbf{v}_t \leftarrow \mathbf{EV}(-\nabla f(\mathbf{X}_t))$
  - 5:    $\mathbf{X}_{t+1} \leftarrow \mathbf{X}_t + \eta_t(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{X}_t)$
  - 6: **end for**
- 

Let us denote the approximation error of Algorithm 1 after  $t$  iterations by  $h_t := f(\mathbf{X}_t) - f(\mathbf{X}^*)$ .

The convergence result of Algorithm 1 is based on the following simple observations:

$$h_{t+1} = f(\mathbf{X}_t + \eta_t(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{X}_t)) - f(\mathbf{X}^*) \quad (3)$$

$$\leq h_t + \eta_t(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{X}_t) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t^2 \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{X}_t\|_F^2$$

$$\leq h_t + \eta_t(\mathbf{X}^* - \mathbf{X}_t) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t^2 \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{X}_t\|_F^2 \leq (1 - \eta_t)h_t + \frac{\eta_t^2 \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{X}_t\|_F^2,$$

where the first inequality follows from the  $\beta$ -smoothness of  $f(\mathbf{X})$ , the second one follows for the optimal choice of  $\mathbf{v}_t$ , and the third one follows from convexity of  $f(\mathbf{X})$ . Unfortunately, while we

expect the error  $h_t$  to rapidly converge to zero, the term  $\|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{X}_t\|_F^2$  in Eq. (3), in principal, might remain as large as the diameter of  $\mathcal{S}_d$ , which, given a proper choice of step-size  $\eta_t$ , results in the well-known convergence rate of  $O(\beta/t)$  [12, 10]. This consequence holds also in case  $f(\mathbf{X})$  is not only smooth, but also strongly-convex.

However, in case  $f$  is strongly convex, a non-trivial modification of Algorithm 1 can lead to a much faster convergence rate. In this case, it follows from Eq. (1), that on any iteration  $t$ ,  $\|\mathbf{X}_t - \mathbf{X}^*\|_F^2 \leq \frac{2}{\alpha} h_t$ . Thus, if we consider replacing the choice of  $\mathbf{X}_{t+1}$  in Algorithm 1 with the following update rule:

$$\mathbf{V}_t \leftarrow \arg \min_{\mathbf{V} \in \mathcal{S}_d} \mathbf{V} \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{V} - \mathbf{X}_t\|_F^2, \quad \mathbf{X}_{t+1} \leftarrow \mathbf{X}_t + \eta_t (\mathbf{V}_t - \mathbf{X}_t), \quad (4)$$

then, following basically the same steps as in Eq. (3), we will have that

$$h_{t+1} \leq h_t + \eta_t (\mathbf{X}^* - \mathbf{X}_t) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t^2 \beta}{2} \|\mathbf{X}^* - \mathbf{X}_t\|_F^2 \leq \left(1 - \eta_t + \frac{\eta_t^2 \beta}{\alpha}\right) h_t, \quad (5)$$

and thus by a proper choice of  $\eta_t$ , a linear convergence rate will be attained. Of course the issue now, is that computing  $\mathbf{V}_t$  is no longer a computationally-cheap leading eigenvalue problem (in particular  $\mathbf{V}_t$  is not rank-one), but requires a full eigen-decomposition of  $\mathbf{X}_t$ , which is much more expensive. In fact, the update rule in Eq. (4) is nothing more than the projected gradient decent method.

### 3.2 A new hybrid approach: rank one-regularized conditional gradient algorithm

At the heart of our new method is the combination of ideas from both of the above approaches: on one hand, solving a certain regularized linear problem in order to avoid the shortcomings of the CG method, i.e., slow convergence rate, and on the other hand, maintaining the simple structure of a leading eigenvalue computation that avoids the shortcoming of the computationally-expensive projected-gradient method.

Towards this end, suppose that we have an explicit decomposition of the current iterate  $\mathbf{X}_t = \sum_{i=1}^k a_i \mathbf{x}_i \mathbf{x}_i^\top$ , where  $(a_1, a_2, \dots, a_k)$  is a probability distribution over  $[k]$ , and each  $\mathbf{x}_i$  is a unit vector. Note in particular that the standard CG method (Algorithm 1) naturally produces such an explicit decomposition of  $\mathbf{X}_t$  (provided  $\mathbf{X}_1$  is chosen to be rank-one). Consider now the update rule in Eq. (4), but with the additional restriction that  $\mathbf{V}_t$  is rank one, i.e.  $\mathbf{V}_t \leftarrow \arg \min_{\mathbf{V} \in \mathcal{S}_d, \text{rank}(\mathbf{V})=1} \mathbf{V} \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{V} - \mathbf{X}_t\|_F^2$ . Note that in this case it follows that  $\mathbf{V}_t$  is a unit trace rank-one matrix which corresponds to the leading eigenvector of the matrix  $-\nabla f(\mathbf{X}_t) + \eta_t \beta \mathbf{X}_t$ . However, when  $\mathbf{V}_t$  is rank-one, the regularization  $\|\mathbf{V}_t - \mathbf{X}_t\|_F^2$  makes little sense in general, since unless  $\mathbf{X}^*$  is rank-one, we do not expect  $\mathbf{X}_t$  to be such (note however, that if  $\mathbf{X}^*$  is rank one, this modification will already result in a linear convergence rate). However, we can think of solving a set of decoupled component-wise regularized problems:

$$\begin{aligned} \forall i \in [k] : \mathbf{v}_t^{(i)} &\leftarrow \arg \min_{\|\mathbf{v}\|=1} \mathbf{v}^\top \nabla f(\mathbf{X}_t) \mathbf{v} + \frac{\eta_t \beta}{2} \|\mathbf{v} \mathbf{v}^\top - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \equiv \mathbf{E} \mathbf{V} (-\nabla f(\mathbf{X}_t) + \eta_t \beta \mathbf{x}_i \mathbf{x}_i^\top) \\ \mathbf{X}_{t+1} &\leftarrow \sum_{i=1}^k a_i \left( (1 - \eta_t) \mathbf{x}_i \mathbf{x}_i^\top + \eta_t \mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} \right), \end{aligned} \quad (6)$$

where the equivalence in the first line follows since  $\|\mathbf{v} \mathbf{v}^\top\|_F = 1$ , and thus the minimizer of the LHS is w.l.o.g. a leading eigenvector of the matrix on the RHS. Following the lines of Eq. (3), we will now have that

$$\begin{aligned} h_{t+1} &\leq h_t + \eta_t \sum_{i=1}^k a_i (\mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t^2 \beta}{2} \left\| \sum_{i=1}^k a_i (\mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} - \mathbf{x}_i \mathbf{x}_i^\top) \right\|_F^2 \\ &\leq h_t + \eta_t \sum_{i=1}^k a_i (\mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t^2 \beta}{2} \sum_{i=1}^k a_i \|\mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \\ &= h_t + \eta_t \mathbb{E}_{i \sim (a_1, \dots, a_k)} \left[ (\mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right], \end{aligned} \quad (7)$$

where the second inequality follows from convexity of the squared Frobenius norm, and the last equality follows since  $(a_1, \dots, a_k)$  is a probability distribution over  $[k]$ .

While the approach in Eq. (6) relies only on leading eigenvector computations, the benefit in terms of potential convergence rates is not trivial, since it is not immediate that we can get non-trivial bounds for the individual distances  $\|\mathbf{v}_t^{(i)} \mathbf{v}_t^{(i)\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F$ . Indeed, the main novelty in our analysis is dedicated precisely to this issue. A motivation, if any, is that there might exist a decomposition of  $\mathbf{X}^*$  as  $\mathbf{X}^* = \sum_{i=1}^k b_i \mathbf{x}^{*(i)} \mathbf{x}^{*(i)\top}$ , which is close in some sense to the decomposition of  $\mathbf{X}_t$ . We can then think of the regularized problem in Eq. (6), as an attempt to push each individual component  $\mathbf{x}^{(i)}$  towards its corresponding component in the decomposition of  $\mathbf{X}^*$ , and as an overall result, bring the following iterate  $\mathbf{X}_{t+1}$  closer to  $\mathbf{X}^*$ .

Note that Eq. (7) implicitly describes a randomized algorithm in which, instead of solving a regularized EV problem for each rank-one matrix in the decomposition of  $\mathbf{X}_t$ , which is expensive as this decomposition grows large with the number of iterations, we pick a single rank-one component according to its weight in the decomposition, and only update it. This directly brings us to our proposed algorithm, Algorithm 2, which is given below.

---

**Algorithm 2** Randomized Rank one-regularized Conditional Gradient

---

- 1: input: sequence of step-sizes  $\{\eta_t\}_{t \geq 1}$ , sequence of error tolerances  $\{\xi_t\}_{t \geq 0}$
- 2: let  $\mathbf{x}_0$  be an arbitrary unit vector
- 3:  $\mathbf{X}_1 \leftarrow \mathbf{x}_1 \mathbf{x}_1^\top$  such that  $\mathbf{x}_1 \leftarrow \mathbf{E} \mathbf{V}_{\xi_0}(-\nabla f(\mathbf{x}_0 \mathbf{x}_0^\top))$
- 4: **for**  $t = 1 \dots$  **do**
- 5:   suppose  $\mathbf{X}_t$  is given by  $\mathbf{X}_t = \sum_{i=1}^k a_i \mathbf{x}_i \mathbf{x}_i^\top$ , where each  $\mathbf{x}_i$  is a unit vector, and  $(a_1, a_2, \dots, a_k)$  is a probability distribution over  $[k]$ , for some integer  $k$
- 6:   pick  $i_t \in [k]$  according to the probability distribution  $(a_1, a_2, \dots, a_k)$
- 7:   set a new step-size  $\tilde{\eta}_t$  as follows:

$$\tilde{\eta}_t \leftarrow \begin{cases} \eta_t/2 & \text{if } a_{i_t} \geq \eta_t \\ a_{i_t} & \text{else} \end{cases}$$

- 8:    $\mathbf{v}_t \leftarrow \mathbf{E} \mathbf{V}_{\xi_t}(-\nabla f(\mathbf{X}_t) + \eta_t \beta \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top)$
  - 9:    $\mathbf{X}_{t+1} \leftarrow \mathbf{X}_t + \tilde{\eta}_t (\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top)$
  - 10: **end for**
- 

We have the following guarantee for Algorithm 2 which is the main result of this paper.

**Theorem 1.** [Main Theorem] Consider the sequence of step-sizes  $\{\eta_t\}_{t \geq 1}$  defined by  $\eta_t = 18/(t + 8)$ , and suppose that  $\xi_0 = \beta$  and for any iteration  $t \geq 1$  it holds that  $\xi_t = O\left(\min\left\{\frac{\beta}{t}, \left(\frac{\beta \sqrt{\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4} t}\right)^{4/3}, \left(\frac{\beta}{\sqrt{\alpha} \lambda_{\min}(\mathbf{X}^*) t}\right)^2\right\}\right)$ . Then, all iterates are feasible, and

$$\forall t \geq 1: \quad \mathbb{E}[f(\mathbf{X}_t) - f(\mathbf{X}^*)] = O\left(\min\left\{\frac{\beta}{t}, \left(\frac{\beta \sqrt{\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4} t}\right)^{4/3}, \left(\frac{\beta}{\sqrt{\alpha} \lambda_{\min}(\mathbf{X}^*) t}\right)^2\right\}\right).$$

It is important to note that the step-size choice in Theorem 1 does not require any knowledge on the parameters  $\alpha, \beta, \text{rank}(\mathbf{X}^*)$ , and  $\lambda_{\min}(\mathbf{X}^*)$ . The knowledge of  $\beta$  is required however for the EV computations. While it follows from Theorem 1 that the knowledge of  $\alpha, \text{rank}(\mathbf{X}^*), \lambda_{\min}(\mathbf{X}^*)$  is needed to set the accuracy parameters -  $\xi_t$ , in practice, iterative eigenvector methods are very efficient and are much less sensitive to exact knowledge of parameters than the choice of step-size for instance.

While the eigenvalue problem in Algorithm 2 is different from the one in Algorithm 1, due to the additional term in  $\mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top$ , the efficiency of solving both problems is essentially the same since efficient EV procedures are based on iteratively multiplying the input matrix with a vector. In particular, multiplying a vector with a rank-one matrix takes  $O(d)$  time. Thus, as long as  $\text{nnz}(\nabla f(\mathbf{X}_t)) = \Omega(d)$ , which is highly reasonable, both EV computations run in essentially the same time.

Finally, note also that aside from the computation of the gradient direction and the leading eigenvector computation, all other operations on any iteration  $t$ , can be carried out in  $O(d^2 + t)$  additional time.

## 4 Analysis

The complete proof of Theorem 1 and all supporting lemmas are given in full detail in the appendix. Here we only detail the two main ingredients in the analysis of Algorithm 2.

Throughout this section, given a matrix  $\mathbf{Y} \in \mathcal{S}_d$ , we let  $\mathbf{P}_{\mathbf{Y},\tau} \in \mathbb{S}_d$  denote the projection matrix onto all eigenvectors of  $\mathbf{Y}$  that correspond to eigenvalues of magnitude at least  $\tau$ . Similarly, we let  $\mathbf{P}_{\mathbf{Y},\tau}^\perp$  denote the projection matrix onto the eigenvectors of  $\mathbf{Y}$  that correspond to eigenvalues of magnitude smaller than  $\tau$  (including eigenvectors that correspond to zero-valued eigenvalues).

### 4.1 A new decomposition for positive semidefinite matrices with locality properties

The analysis of Algorithm 2 relies heavily on a new decomposition idea of matrices in  $\mathcal{S}_d$  that suggests that given a matrix  $\mathbf{X}$  in the form of a convex combination of rank-one matrices:  $\mathbf{X} = \sum_{i=1}^k \alpha_i \mathbf{x}_i \mathbf{x}_i^\top$ , and another matrix  $\mathbf{Y} \in \mathcal{S}_d$ , roughly speaking, we can decompose  $\mathbf{Y}$  as the sum of rank-one matrices, such that the components in the decomposition of  $\mathbf{Y}$  are close to those in the decomposition of  $\mathbf{X}$  in terms of the overall distance  $\|\mathbf{X} - \mathbf{Y}\|_F$ . This decomposition and corresponding property justifies the idea of solving rank-one regularized problems, as suggested in Eq. (6), and applied in Algorithm 2.

**Lemma 1.** *Let  $\mathbf{X}, \mathbf{Y} \in \mathcal{S}_d$  such that  $\mathbf{X}$  is given as  $\mathbf{X} = \sum_{i=1}^k a_i \mathbf{x}_i \mathbf{x}_i^\top$ , where each  $\mathbf{x}_i$  is a unit vector, and  $(a_1, \dots, a_k)$  is a distribution over  $[k]$ , and let  $\tau, \gamma \in [0, 1]$  be scalars that satisfy  $\frac{\gamma\tau}{1-\gamma} \geq \|\mathbf{X} - \mathbf{Y}\|_F$ . Then,  $\mathbf{Y}$  can be written as  $\mathbf{Y} = \sum_{i=1}^k b_i \mathbf{y}_i \mathbf{y}_i^\top + \sum_{j=1}^k (a_j - b_j) \mathbf{W}_j$ , such that*

1. each  $\mathbf{y}_i$  is a unit vector and  $\mathbf{W} \in \mathcal{S}_d$
2.  $\forall i \in [k] : 0 \leq b_i \leq a_i$  and  $\sum_{j=1}^k (a_j - b_j) \leq \sqrt{\text{rank}(\mathbf{Y})} (\|\mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}^\perp\|_F + \|\mathbf{X} - \mathbf{Y}\|_F) + \gamma$
3.  $\sum_{i=1}^k b_i \|\mathbf{x}_i \mathbf{x}_i^\top - \mathbf{y}_i \mathbf{y}_i^\top\|_F^2 \leq 2\sqrt{\text{rank}(\mathbf{Y})} (\|\mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}^\perp\|_F + \|\mathbf{X} - \mathbf{Y}\|_F)$

### 4.2 Bounding the per-iteration improvement

The main step in the proof of Theorem 1, is understanding the per-iteration improvement, as captured in Eq. (7), achievable by applying the update rule in Eq. (6), which updates on each iteration all of the rank-one components in the decomposition of the current iterate.

**Lemma 2.** *[full deterministic update] Fix a scalar  $\eta > 0$ . Let  $\mathbf{X} \in \mathcal{S}_d$  such that  $\mathbf{X} = \sum_{i=1}^k a_i \mathbf{x}_i \mathbf{x}_i^\top$ , where each  $\mathbf{x}_i$  is a unit vector, and  $(a_1, \dots, a_k)$  is a probability distribution over  $[k]$ . For any  $i \in [k]$ , let  $\mathbf{v}_i := \mathbf{E} \mathbf{V} (-\nabla f(\mathbf{X}) + \eta\beta \mathbf{x}_i \mathbf{x}_i^\top)$ . Then, it holds that*

$$\begin{aligned} \sum_{i=1}^k a_i \left[ (\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] &\leq -(f(\mathbf{X}) - f(\mathbf{X}^*)) \\ + \eta\beta \cdot \min\{1, 5\sqrt{\frac{2}{\alpha} \text{rank}(\mathbf{X}^*)} \sqrt{f(\mathbf{X}) - f(\mathbf{X}^*)}, \frac{3\sqrt{2}}{\sqrt{\alpha} \lambda_{\min}(\mathbf{X}^*)} \sqrt{f(\mathbf{X}) - f(\mathbf{X}^*)}\}. \end{aligned}$$

*proof sketch.* The proof is divided to three parts, each corresponding to a different term in the min expression in the bound in the Lemma. The first bound, at a high-level, follows from the standard conditional gradient analysis (see Eq. (3)). We continue to derive the second and third bounds.

From Lemma 1 we know we can write  $\mathbf{X}^*$  in the following way:

$$\mathbf{X}^* = \sum_{i=1}^k b_i^* \mathbf{y}_i^* \mathbf{y}_i^{*\top} + \sum_{j=1}^k (a_j - b_j^*) \mathbf{W}_j^*, \quad (8)$$

where for all  $i \in [k]$ ,  $b_i^* \in [0, a_i]$  and  $\mathbf{y}_i^*$  is a unit vector, and  $\mathbf{W}^* \in \mathcal{S}_d$ .

Using nothing more than Eq. (8), the optimality of  $\mathbf{v}_i$  for each  $i \in [k]$ , and the bounds in Lemma 1, it can be shown that

$$\begin{aligned} & \sum_{i=1}^k a_i \left[ (\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \leq \\ & (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \sum_{i=1}^k b_i^* \|\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 + \eta\beta \sum_{i=1}^k (a_i - b_i^*) \leq \\ & (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \eta\beta \left( 2\sqrt{\text{rank}(\mathbf{X}^*)} (\|\mathbf{X}^* \mathbf{P}_{\mathbf{X}^*, \tau}^\perp\|_F + \|\mathbf{X} - \mathbf{X}^*\|_F) + \gamma \right). \end{aligned} \quad (9)$$

Now we can optimize the above bound in terms of  $\tau, \gamma$ . One option is to upper bound  $\|\mathbf{X}^* \mathbf{P}_{\mathbf{X}^*, \tau}^\perp\|_F \leq \sqrt{\text{rank}(\mathbf{X}^*)} \tau$ , which together with the choice  $\tau_1 = \sqrt{\frac{\|\mathbf{X} - \mathbf{X}^*\|_F}{2\text{rank}(\mathbf{X}^*)}}$ ,  $\gamma_1 = \sqrt{2\text{rank}(\mathbf{X}^*)} \|\mathbf{X} - \mathbf{X}^*\|_F$ , give us:

$$\text{RHS of (9)} \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + 5\eta\beta \sqrt{\text{rank}(\mathbf{X}^*)} \|\mathbf{X} - \mathbf{X}^*\|_F. \quad (10)$$

Another option, is to choose  $\tau_2 = \lambda_{\min}(\mathbf{X}^*)$ ,  $\gamma_2 = \frac{\|\mathbf{X} - \mathbf{X}^*\|_F}{\lambda_{\min}(\mathbf{X}^*)}$  which gives us  $\|\mathbf{X}^* \mathbf{P}_{\mathbf{X}^*, \tau}^\perp\|_F = 0$ . This results in the bound:

$$\text{RHS of (9)} \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \frac{3\eta\beta \|\mathbf{X} - \mathbf{X}^*\|_F}{\lambda_{\min}(\mathbf{X}^*)}. \quad (11)$$

Now, using the convexity of  $f$  to upper bound  $(\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) \leq -(f(\mathbf{X}) - f(\mathbf{X}^*))$  and Eq. (1) in both Eq. (10) and (11), gives the second the third parts of the bound in the lemma.  $\square$

## 5 Preliminary Empirical Evaluation

We evaluate our method, along with other conditional gradient variants, on the task of matrix completion [13].

**Setting** The underlying optimization problem for the matrix completion task is the following:

$$\min_{\mathbf{Z} \in \mathcal{NB}_{d_1, d_2}(\theta)} \{f(\mathbf{Z}) := \frac{1}{2} \sum_{l=1}^n (\mathbf{Z} \bullet \mathbf{E}_{i_l, j_l} - r_l)^2\}, \quad (12)$$

where  $\mathbf{E}_{i, j}$  is the indicator matrix for the entry  $(i, j)$  in  $\mathbb{R}^{d_1 \times d_2}$ ,  $\{(i_l, j_l, r_l)\}_{l=1}^n \subset [d_1] \times [d_2] \times \mathbb{R}$ , and  $\mathcal{NB}_{d_1, d_2}(\theta)$  denotes the nuclear-norm ball of radius  $\theta$  in  $\mathbb{R}^{d_1 \times d_2}$ , i.e.,

$$\mathcal{NB}_{d_1, d_2}(\theta) := \{\mathbf{Z} \in \mathbb{R}^{d_1 \times d_2} \mid \|\mathbf{Z}\|_* := \sum_{i=1}^{\min\{d_1, d_2\}} \sigma_i(\mathbf{Z}) \leq \theta\},$$

where we let  $\sigma(\mathbf{Z})$  denote the vector of singular values of  $\mathbf{Z}$ . That is, our goal is to find a matrix with bounded nuclear norm (which serves as a convex surrogate for bounded rank) which matches best the partial observations given by  $\{(i_l, j_l, r_l)\}_{l=1}^n$ .

In order to transform Problem (12) to optimization over the spectrahedron, we use the reduction specified in full detail in [13], and also described in Section A in the appendix.

The objective function in Eq. (12) is known to have a smoothness parameter  $\beta$  with respect to  $\|\cdot\|_F$ , which satisfies  $\beta = O(1)$ , see for instance [13]. While the objective function in Eq. (12) is not strongly convex, it is known that under certain conditions, the matrix completion problem exhibit properties very similar to strong convexity, in the sense of Eq. (1) (which is indeed the only consequence of strong convexity that we use in our analysis) [18].

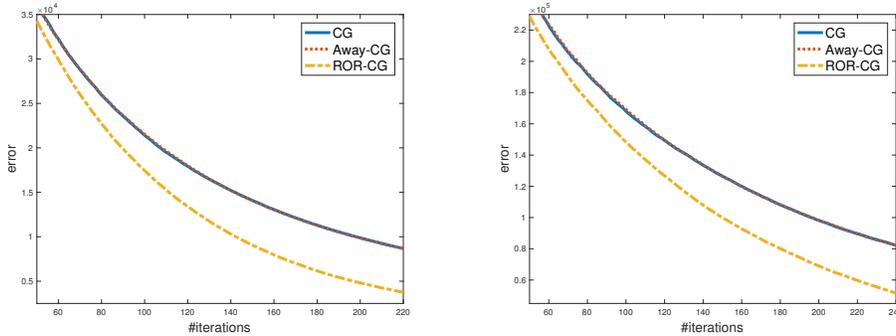


Figure 1: Comparison between conditional gradient variants for solving the matrix completion problem on the MOVIELENS100K (left) and MOVIELENS1M (right) datasets.

**Two modifications of Algorithm 2** We implemented our rank one-regularized conditional gradient variant, Algorithm 2 (denoted ROR-CG in our figures) with two modifications. First, on each iteration  $t$ , instead of picking an index  $i_t$  of a rank-one matrix in the decomposition of the current iterate at random according to the distribution  $(a_1, a_2, \dots, a_k)$ , we choose it in a greedy way, i.e., we choose the rank-one component that has the largest product with the current gradient direction. While this approach is computationally more expensive, it could be easily parallelized since all dot-product computations are independent of each other. Second, after computing the eigenvector  $\mathbf{v}_t$  using the step-size  $\eta_t = 1/t$  (which is very close to that prescribed in Theorem 1), we apply a line-search, as detailed in [13], in order to determine the optimal step-size given the direction  $\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top$ .

**Baselines** As baselines for comparison we used the standard conditional gradient method with exact line-search for setting the step-size (denoted CG in our figures)[13], and the conditional gradient with away-steps variant, recently studied in [15] (denoted Away-CG in our figures). While the away-steps variant was studied in the context of optimization over polyhedral sets, and its formal improved guarantees apply only in that setting, the concept of away-steps still makes sense for any convex feasible set. This variant also allows the incorporation of an exact line-search procedure to choose the optimal step-size.

**Datasets** We have experimented with two well known datasets for the matrix completion task: the MOVIELENS100K dataset for which  $d_1 = 943, d_2 = 1682, n = 10^5$ , and the MOVIELENS1M dataset for which  $d_1 = 6040, d_2 = 3952, n \approx 10^6$ . The MOVIELENS1M dataset was further sub-sampled to contain roughly half of the observations. We have set the parameter  $\theta$  in Problem (12) to  $\theta = 10000$  for the ML100K dataset, and  $\theta = 35000$  for the ML1M dataset.

Figure 1 presents the objective (12) vs. the number of iterations executed. Each graph is the average over 5 independent experiments<sup>2</sup>. It can be seen that our approach indeed improves significantly over the baselines in terms of convergence rate, for the setting under consideration.

## References

- [1] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- [2] Miroslav Dudík, Zaïd Harchaoui, and Jérôme Malick. Lifted coordinate descent for learning with trace-norm regularization. *Journal of Machine Learning Research - Proceedings Track*, 22:327–336, 2012.
- [3] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3:149–154, 1956.
- [4] Robert M Freund, Paul Grigas, and Rahul Mazumder. An extended frank-wolfe method with "in-face" directions, and its application to low-rank matrix completion. *arXiv preprint arXiv:1511.02204*, 2015.

<sup>2</sup>We ran several experiments since the leading eigenvector computation in each one of the CG variants is randomized.

- [5] Dan Garber and Elad Hazan. Fast and simple pca via convex optimization. *arXiv preprint arXiv:1509.05647*, 2015.
- [6] Dan Garber and Elad Hazan. Faster rates for the frank-wolfe method over strongly-convex sets. In *Proceedings of the 32nd International Conference on Machine Learning, ICML*, pages 541–549, 2015.
- [7] Dan Garber and Elad Hazan. A linearly convergent variant of the conditional gradient algorithm under strong convexity, with applications to online and stochastic optimization. *SIAM Journal on Optimization*, 26(3):1493–1528, 2016.
- [8] Dan Garber, Elad Hazan, Chi Jin, Sham M. Kakade, Cameron Musco, Praneeth Netrapalli, and Aaron Sidford. Faster eigenvector computation via shift-and-invert preconditioning. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA*, pages 2626–2634, 2016.
- [9] Mehmet Gönen and Ethem Alpaydın. Multiple kernel learning algorithms. *The Journal of Machine Learning Research*, 12:2211–2268, 2011.
- [10] Elad Hazan. Sparse approximate solutions to semidefinite programs. In *8th Latin American Theoretical Informatics Symposium, LATIN*, 2008.
- [11] Elad Hazan and Satyen Kale. Projection-free online learning. In *Proceedings of the 29th International Conference on Machine Learning, ICML*, 2012.
- [12] Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning, ICML*, 2013.
- [13] Martin Jaggi and Marek Sulovský. A simple algorithm for nuclear norm regularized problems. In *Proceedings of the 27th International Conference on Machine Learning, ICML*, 2010.
- [14] J. Kuczynski and H. Woźniakowski. Estimating the largest eigenvalues by the power and lanczos algorithms with a random start. *SIAM J. Matrix Anal. Appl.*, 13:1094–1122, October 1992.
- [15] Simon Lacoste-Julien and Martin Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. In *Advances in Neural Information Processing Systems*, pages 496–504, 2015.
- [16] Gert RG Lanckriet, Nello Cristianini, Peter Bartlett, Laurent El Ghaoui, and Michael I Jordan. Learning the kernel matrix with semidefinite programming. *The Journal of Machine Learning Research*, 5:27–72, 2004.
- [17] Sören Laue. A hybrid algorithm for convex semidefinite optimization. In *Proceedings of the 29th International Conference on Machine Learning, ICML*, 2012.
- [18] Sahand Negahban, Bin Yu, Martin J Wainwright, and Pradeep K. Ravikumar. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 1348–1356. 2009.
- [19] Shai Shalev-Shwartz, Alon Gonen, and Ohad Shamir. Large-scale convex minimization with a low-rank constraint. In *Proceedings of the 28th International Conference on Machine Learning, ICML*, 2011.
- [20] Ohad Shamir. A stochastic PCA and SVD algorithm with an exponential convergence rate. In *Proceedings of the 32nd International Conference on Machine Learning, ICML*, 2015.
- [21] Eric P Xing, Andrew Y Ng, Michael I Jordan, and Stuart Russell. Distance metric learning with application to clustering with side-information. *Advances in neural information processing systems*, 15:505–512, 2003.
- [22] Yiming Ying and Peng Li. Distance metric learning with eigenvalue optimization. *J. Mach. Learn. Res.*, 13(1):1–26, January 2012.
- [23] Xinhua Zhang, Dale Schuurmans, and Yao-liang Yu. Accelerated training for matrix-norm regularization: A boosting approach. In *Advances in Neural Information Processing Systems*, pages 2906–2914, 2012.

## A Convex Optimization with a Nuclear Norm Constraint

An important optimization problem highly-related to Problem (2), is the problem of minimizing a convex function over the set of  $d_1 \times d_2$  real-valued matrices with bounded nuclear norm, i.e.,

$$\min_{\mathbf{Z} \in \mathcal{NB}_{d_1, d_2}(\theta)} f(\mathbf{Z}). \quad (13)$$

Here we let  $\mathcal{NB}_{d_1, d_2}(\theta)$  denote the nuclear-norm ball of radius  $\theta$  in  $\mathbb{R}^{d_1 \times d_2}$ , i.e.,

$$\mathcal{NB}_{d_1, d_2}(\theta) := \{\mathbf{Z} \in \mathbb{R}^{d_1 \times d_2} \mid \|\mathbf{Z}\|_* := \sum_{i=1}^{\min\{d_1, d_2\}} \sigma_i(\mathbf{Z}) \leq \theta\},$$

where we let  $\sigma(\mathbf{Z})$  denote the vector of singular values of  $\mathbf{Z}$ .

Problem (13) could be directly formulated as convex optimization over the spectrahedron. Towards this end, consider now the following convex optimization problem:

$$\begin{aligned} & \min_{\mathbf{X} \in \mathcal{S}_{d_1+d_2}} \hat{f}(\mathbf{X}), \\ \hat{f}(\mathbf{X}) & := f(2\theta \cdot \mathbf{M}_1 \mathbf{X} \mathbf{M}_2), \quad \mathbf{M}_1 := \begin{pmatrix} \mathbf{I}_{d_1} & \mathbf{0}_{d_1 \times d_2} \end{pmatrix}, \quad \mathbf{M}_2 := \begin{pmatrix} \mathbf{0}_{d_1 \times d_2} \\ \mathbf{I}_{d_2} \end{pmatrix}. \end{aligned}$$

The following Lemma, whose proof can be found in [13], shows the equivalence between the two problems.

**Lemma 3.** *Let  $\mathbf{X} \in \mathcal{S}_{d_1+d_2}$  such that  $\hat{f}(\mathbf{X}) - \hat{f}(\mathbf{X}^*) = \epsilon$ , for some  $\epsilon > 0$ , where  $\mathbf{X}^*$  is the minimizer of  $\hat{f}$  over  $\mathcal{S}_{d_1+d_2}$ . Consider the following factorization of  $\mathbf{X}$ :*

$$\mathbf{X} = \begin{pmatrix} \mathbf{X}_1 & \mathbf{X}_2 \\ \mathbf{X}_2^\top & \mathbf{X}_3 \end{pmatrix},$$

where  $\mathbf{X}_1$  is  $d_1 \times d_1$ ,  $\mathbf{X}_2$  is  $d_1 \times d_2$ , and  $\mathbf{X}_3$  is  $d_2 \times d_2$ . Define  $\mathbf{Z} := 2\theta \cdot \mathbf{X}_2$ . Then it follows that  $\mathbf{Z} \in \mathcal{NB}_{d_1, d_2}(\theta)$ , and  $f(\mathbf{Z}) - f(\mathbf{Z}^*) = \epsilon$ , where  $\mathbf{Z}^*$  is the minimizer of  $f$  over  $\mathcal{NB}_{d_1, d_2}(\theta)$ .

## B Omitted Lemmas and Proofs from Subsection 4.1

Here we give a complete proof of Lemma 1. Before we can prove the lemma, we first need two more technical lemmas.

**Lemma 4.** *Let  $\mathbf{X}, \mathbf{Y} \in \mathcal{S}_d$ . Let  $\tau, \gamma \in [0, 1]$  be scalars that satisfy  $\frac{\gamma\tau}{1-\gamma} \geq \|\mathbf{X} - \mathbf{Y}\|_F$ . Then it holds that  $\mathbf{Y} \succeq (1-\gamma)\mathbf{P}_{\mathbf{Y}, \tau} \mathbf{X} \mathbf{P}_{\mathbf{Y}, \tau}$ .*

*Proof.* Given a vector  $\mathbf{w} \in \mathbb{R}^d$  let us write it as  $\mathbf{w} = \mathbf{w}^+ + \mathbf{w}^-$  where  $\mathbf{w}^+ = \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{w}$  and  $\mathbf{w}^- = \mathbf{P}_{\mathbf{Y}, \tau}^\perp \mathbf{w} = \mathbf{w} - \mathbf{w}^+$ .

It holds that

$$\begin{aligned} \mathbf{w}^\top \mathbf{Y} \mathbf{w} &= \mathbf{w}^{+\top} \mathbf{Y} \mathbf{w}^+ + \mathbf{w}^{-\top} \mathbf{Y} \mathbf{w}^- + 2\mathbf{w}^{-\top} \mathbf{Y} \mathbf{w}^+ \\ &= \mathbf{w}^{+\top} \mathbf{Y} \mathbf{w}^+ + \mathbf{w}^{-\top} \mathbf{Y} \mathbf{w}^- + 2\mathbf{w}^\top \mathbf{P}_{\mathbf{Y}, \tau}^\perp \mathbf{Y} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{w} \\ &\geq \mathbf{w}^{+\top} \mathbf{Y} \mathbf{w}^+, \end{aligned} \quad (14)$$

where the inequality follows since  $\mathbf{P}_{\mathbf{Y}, \tau}^\perp \mathbf{Y} \mathbf{P}_{\mathbf{Y}, \tau} = 0$  and  $\mathbf{Y}$  is positive semidefinite.

Similarly, since  $\mathbf{P}_{\mathbf{Y}, \tau} \mathbf{w}^- = 0$ , we have that

$$\mathbf{w}^{-\top} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{X} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{w}^- = \mathbf{w}^{-\top} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{X} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{w}^+ = \mathbf{w}^{+\top} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{X} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{w}^- = 0. \quad (15)$$

Note also that

$$\mathbf{w}^{+\top} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{X} \mathbf{P}_{\mathbf{Y}, \tau} \mathbf{w}^+ = \mathbf{w}^{+\top} \mathbf{X} \mathbf{w}^+. \quad (16)$$

Thus, we have that

$$\begin{aligned}
\mathbf{w}^\top [(1-\gamma)\mathbf{P}_{\mathbf{Y},\tau}\mathbf{X}\mathbf{P}_{\mathbf{Y},\tau}] \mathbf{w} &= (1-\gamma)\mathbf{w}^{+\top}\mathbf{P}_{\mathbf{Y},\tau}\mathbf{X}\mathbf{P}_{\mathbf{Y},\tau}\mathbf{w}^+ \\
&= (1-\gamma)\mathbf{w}^{+\top}\mathbf{X}\mathbf{w}^+ \\
&\leq (1-\gamma)(\mathbf{w}^{+\top}\mathbf{Y}\mathbf{w}^+ + \|\mathbf{X}-\mathbf{Y}\|_F \cdot \|\mathbf{w}^+\|^2) \\
&\leq \mathbf{w}^\top\mathbf{Y}\mathbf{w} + (1-\gamma)\|\mathbf{X}-\mathbf{Y}\|_F \cdot \|\mathbf{w}^+\|^2 - \gamma\mathbf{w}^{+\top}\mathbf{Y}\mathbf{w}^+ \\
&\leq \mathbf{w}^\top\mathbf{Y}\mathbf{w} + (1-\gamma)\|\mathbf{X}-\mathbf{Y}\|_F \cdot \|\mathbf{w}^+\|^2 - \gamma\tau\|\mathbf{w}^+\|^2,
\end{aligned}$$

where the first equality follows from Eq. (15), the second equality follows from Eq. (16), the first inequality follows from the Cauchy-Schwarz ineq., the second inequality follows from Eq. (14), and the last inequality follows from the definitions of  $\mathbf{w}^+$  and  $\tau$ .

Thus, we can see that if  $\frac{\gamma\tau}{1-\gamma} \geq \|\mathbf{X}-\mathbf{Y}\|_F$ , the lemma follows.  $\square$

**Lemma 5.** Let  $\mathbf{X}, \mathbf{Y} \in \mathcal{S}_d$  and suppose  $\mathbf{X}$  is given in the form  $\mathbf{X} = \sum_{i=1}^k a_i \mathbf{x}_i \mathbf{x}_i^\top$  where each  $\mathbf{x}_i$  is a unit vector, and the weights  $(a_1, \dots, a_k)$  are a distribution over  $[k]$ . Let  $\mathbf{P} \in \mathcal{S}_d$  be a projection matrix onto a subset of the eigenvectors of  $\mathbf{Y}$ , and define for any  $i \in [k]$ ,  $\tilde{\mathbf{x}}_i := \mathbf{P}\mathbf{x}_i$ . Then, it holds that

$$\sum_{i=1}^k a_i (1 - \|\tilde{\mathbf{x}}_i\|^2) \leq \sqrt{\text{rank}(\mathbf{Y})} \|\mathbf{Y} - \mathbf{P}\mathbf{X}\mathbf{P}\|_F.$$

*Proof.* Let us write the eigen-decomposition of  $\mathbf{Y}$  as  $\mathbf{Y} = \sum_{j=1}^{\text{rank}(\mathbf{Y})} \lambda_j \mathbf{v}_j \mathbf{v}_j^\top$ . Using simple algebraic manipulations we have that

$$\begin{aligned}
\|\mathbf{Y} - \mathbf{P}\mathbf{X}\mathbf{P}\|_F^2 &\geq \sum_{j=1}^{\text{rank}(\mathbf{Y})} ((\mathbf{Y} - \mathbf{P}\mathbf{X}\mathbf{P}) \bullet \mathbf{v}_j \mathbf{v}_j^\top)^2 = \sum_{j=1}^{\text{rank}(\mathbf{Y})} \left( \lambda_j - \sum_{i=1}^k a_i \mathbf{v}_j^\top \mathbf{P} \mathbf{x}_i \mathbf{x}_i^\top \mathbf{P} \mathbf{v}_j \right)^2 \\
&= \sum_{j=1}^{\text{rank}(\mathbf{Y})} \left( \lambda_j - \sum_{i=1}^k a_i (\mathbf{v}_j^\top \mathbf{P} \mathbf{x}_i)^2 \right)^2 \\
&\geq \frac{1}{\text{rank}(\mathbf{Y})} \left( \sum_{j=1}^{\text{rank}(\mathbf{Y})} \left( \lambda_j - \sum_{i=1}^k a_i (\mathbf{v}_j^\top \mathbf{P} \mathbf{x}_i)^2 \right) \right)^2 \\
&= \frac{1}{\text{rank}(\mathbf{Y})} \left( 1 - \sum_{j=1}^{\text{rank}(\mathbf{Y})} \sum_{i=1}^k a_i (\mathbf{v}_j^\top \mathbf{P} \mathbf{x}_i)^2 \right)^2 \\
&= \frac{1}{\text{rank}(\mathbf{Y})} \left( \sum_{i=1}^k a_i \left( 1 - \sum_{j=1}^{\text{rank}(\mathbf{Y})} (\mathbf{v}_j^\top \mathbf{P} \mathbf{x}_i)^2 \right) \right)^2 \\
&= \frac{1}{\text{rank}(\mathbf{Y})} \left( \sum_{i=1}^k a_i (1 - \|\tilde{\mathbf{x}}_i\|^2) \right)^2,
\end{aligned}$$

where the first inequality follows since  $\{\mathbf{v}_j \mathbf{v}_j^\top\}_{j=1}^{\text{rank}(\mathbf{Y})}$  is a subset of an orthonormal basis for  $\mathbb{R}^{d \times d}$ , the second inequality follows since for any finite set of reals  $\{r_i\}_{i=1}^k$  it holds that  $\sum_{i=1}^k r_i^2 \geq (\sum_{i=1}^k r_i)^2 / k$ , and the third and fourth equalities hold since  $\sum_{j=1}^{\text{rank}(\mathbf{Y})} \lambda_j = \sum_{i=1}^k a_i = 1$ .

Thus we have that

$$\sum_{i=1}^k a_i (1 - \|\tilde{\mathbf{x}}_i\|^2) \leq \sqrt{\text{rank}(\mathbf{Y})} \|\mathbf{Y} - \mathbf{P}\mathbf{X}\mathbf{P}\|_F,$$

which gives the bound in the lemma.  $\square$

We can now prove Lemma 1.

*Proof.* For each  $i \in [k]$  let  $\tilde{\mathbf{x}}_i = \mathbf{P}_{\mathbf{Y},\tau} \mathbf{x}_i$ . It follows from Lemma 4 that as long as  $\frac{\gamma\tau}{1-\gamma} \geq \|\mathbf{X} - \mathbf{Y}\|_F$ , it holds that

$$\mathbf{Y} \succeq \sum_{i=1}^k a_i (1 - \gamma) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top.$$

Since  $\mathbf{Y} \in \mathcal{S}_d$  and  $\text{Tr} \left( \sum_{i=1}^k a_i (1 - \gamma) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \right) = \sum_{i=1}^k a_i (1 - \gamma) \|\tilde{\mathbf{x}}_i\|^2$ , it follows that  $\mathbf{Y}$  can be written as:

$$\mathbf{Y} = \sum_{i=1}^k a_i (1 - \gamma) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top + \left( \sum_{j=1}^k a_j (1 - (1 - \gamma) \|\tilde{\mathbf{x}}_j\|^2) \right) \mathbf{W},$$

where  $\mathbf{W} \in \mathcal{S}_d$ .

Let us now define  $\mathbf{y}_i := \frac{\tilde{\mathbf{x}}_i}{\|\tilde{\mathbf{x}}_i\|}$  and  $b_i := a_i (1 - \gamma) \|\tilde{\mathbf{x}}_i\|^2$ . Then indeed it follows that

$$\mathbf{Y} = \sum_{i=1}^k b_i \mathbf{y}_i \mathbf{y}_i^\top + \sum_{j=1}^k (a_j - b_j) \mathbf{W}.$$

We are going to apply Lemma 5 to derive the bounds listed in the lemma. As a first step, we need to bound the distance  $\|\mathbf{Y} - \mathbf{P}_{\mathbf{Y},\tau} \mathbf{X} \mathbf{P}_{\mathbf{Y},\tau}\|_F$ .

$$\begin{aligned} \|\mathbf{Y} - \mathbf{P}_{\mathbf{Y},\tau} \mathbf{X} \mathbf{P}_{\mathbf{Y},\tau}\|_F &\leq \|\mathbf{Y} - \mathbf{P}_{\mathbf{Y},\tau} \mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}\|_F + \|\mathbf{P}_{\mathbf{Y},\tau} \mathbf{X} \mathbf{P}_{\mathbf{Y},\tau} - \mathbf{P}_{\mathbf{Y},\tau} \mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}\|_F \\ &\leq \|\mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}^\perp\|_F + \|\mathbf{X} - \mathbf{Y}\|_F, \end{aligned} \quad (17)$$

where the bound on  $\|\mathbf{Y} - \mathbf{P}_{\mathbf{Y},\tau} \mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}\|_F$  follows from the definition of  $\mathbf{P}_{\mathbf{Y},\tau}^\perp$ , and the bound on  $\|\mathbf{P}_{\mathbf{Y},\tau} \mathbf{X} \mathbf{P}_{\mathbf{Y},\tau} - \mathbf{P}_{\mathbf{Y},\tau} \mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}\|_F$  follows from the inequality  $\|\mathbf{A} \mathbf{B}\|_F \leq \|\mathbf{A}\| \cdot \|\mathbf{B}\|_F$  (applied twice, each time with  $\mathbf{A} = \mathbf{P}_{\mathbf{Y},\tau}$ , and recalling that  $\|\mathbf{P}_{\mathbf{Y},\tau}\| = 1$ ).

By definition of  $\{b_i\}_{i \in [k]}$  it holds that

$$\begin{aligned} \sum_{i=1}^k (a_i - b_i) &= \sum_{i=1}^k a_i (1 - (1 - \gamma) \|\tilde{\mathbf{x}}_i\|^2) \leq \sum_{i=1}^k a_i (1 - \|\tilde{\mathbf{x}}_i\|^2) + \gamma \\ &\leq \sqrt{\text{rank}(\mathbf{Y})} (\|\mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}^\perp\|_F + \|\mathbf{X} - \mathbf{Y}\|_F) + \gamma, \end{aligned}$$

where the last inequality follows from Lemma 5 and the bound in Eq. (17).

We continue to upper-bound  $\sum_{i=1}^k b_i \|\mathbf{x}_i \mathbf{x}_i^\top - \mathbf{y}_i \mathbf{y}_i^\top\|_F^2$ :

$$\begin{aligned} \sum_{i=1}^k b_i \|\mathbf{x}_i \mathbf{x}_i^\top - \mathbf{y}_i \mathbf{y}_i^\top\|_F^2 &\leq \sum_{i=1}^k a_i \|\mathbf{x}_i \mathbf{x}_i^\top - \mathbf{y}_i \mathbf{y}_i^\top\|_F^2 \\ &= 2 \sum_{i=1}^k a_i (1 - (\mathbf{x}_i^\top \mathbf{y}_i)^2) \\ &= 2 \sum_{i=1}^k a_i \left( 1 - \left( \frac{\mathbf{x}_i^\top \tilde{\mathbf{x}}_i}{\|\tilde{\mathbf{x}}_i\|} \right)^2 \right) \\ &= 2 \sum_{i=1}^k a_i (1 - \|\tilde{\mathbf{x}}_i\|^2) \\ &\leq 2 \sqrt{\text{rank}(\mathbf{Y})} (\|\mathbf{Y} \mathbf{P}_{\mathbf{Y},\tau}^\perp\|_F + \|\mathbf{X} - \mathbf{Y}\|_F), \end{aligned}$$

where the last inequality follows again from the application of Lemma 5 and the bound in Eq. (17).  $\square$

## C Omitted Lemmas and Proofs from Subsection 4.2

We now turn to analyze the per-iteration improvement of Algorithm 2. We start by first analyzing a deterministic, and much less efficient, version that updates all of the rank-one components on each iteration  $t$ , as suggested in Eq. (6). This is done in Lemma 2, whose complete proof is given here. Then in Lemma 7, we apply Lemma 2, to analyze the randomized step of Algorithm 2. However, first we need a simple observation regarding Algorithm 2, that shows that it can always take sufficiently large step-sizes, i.e., step-size of magnitude at least  $\eta_t/2$  on iteration  $t$ .

**Observation 1.** *In case the input sequence of step-sizes in Algorithm 2 -  $\{\eta_t\}_{t \geq 1}$ , is monotonically non-increasing and  $\eta_t \in [0, 2]$  for all  $t \geq 1$ , it follows that on each iteration  $t$  of the algorithm, the iterate  $\mathbf{X}_t$  admits an explicitly-given factorization into a convex sum of rank-one matrices, as described in the algorithm, such that for every rank-one coefficient  $a_i$ , it holds that  $a_i \geq \eta_t/2$ .*

*Proof.* The proof is by a simple induction. Since  $\mathbf{X}_1$  is just a rank-one matrix, it follows that the corresponding coefficient in the convex sum is  $a_1 = 1$ . Thus, for any  $\eta_1 \in [0, 2]$  it indeed follows that  $a_1 \geq \eta_1/2$ . Assume now that the induction holds for time  $t \geq 1$ . On time  $t$  we choose a coefficient  $a_{i_t}$  and move a mass of  $\tilde{\eta}_t$  from it to a new rank-one matrix  $\mathbf{y}_t \mathbf{y}_t^\top$ , and all other coefficients remain unchanged. Since we assume that the step-size sequence is monotonically non-increasing, it directly follows that the induction step holds for all unchanged coefficients. Regarding the affected coefficients  $a_{i_t}$  and the coefficient of the new rank-one matrix, we consider two cases. First, if  $a_{i_t} \geq \eta_t$  then by the definition of  $\tilde{\eta}_t$  we have that the mass of the new coefficient is going to be exactly  $\eta_t/2$  and the mass of the old coefficient is going to be  $a_{i_t} - \eta_t/2 \geq \eta_t/2$ , and thus the induction holds. In the other case, we have that  $\tilde{\eta}_t = a_{i_t} < \eta_t$ . By the induction hypothesis we know that  $a_{i_t} \geq \eta_t/2 \geq \eta_{t+1}/2$ . Since we are moving now all the mass from  $a_{i_t}$  to the new rank-one matrix, it follows that its weight is also going to be at least  $\eta_{t+1}/2$ , and thus the induction follows.  $\square$

We now restate Lemma 2 and give its complete proof. In this version of the lemma we also take into consideration the approximation errors in the eigenvector computations.

**Lemma 6.** *[full deterministic update] Fix a scalar  $\eta > 0$ . Let  $\mathbf{X} \in \mathcal{S}_d$  such that  $\mathbf{X} = \sum_{i=1}^k a_i \mathbf{x}_i \mathbf{x}_i^\top$ , where each  $\mathbf{x}_i$  is a unit vector, and  $(a_1, \dots, a_k)$  is a probability distribution over  $[k]$ . For any  $i \in [k]$ , let*

$$\mathbf{v}_i := \mathbf{E} \mathbf{V}_\xi (-\nabla f(\mathbf{X}) + \eta \beta \mathbf{x}_i \mathbf{x}_i^\top), \quad (18)$$

for some parameter  $\xi > 0$ . Then, it holds that

$$\begin{aligned} & \sum_{i=1}^k a_i \left[ (\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta \beta}{2} \|\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \leq -(f(\mathbf{X}) - f(\mathbf{X}^*)) \\ & + \eta \beta \cdot \min\left\{1, 5 \sqrt{\frac{2}{\alpha} \text{rank}(\mathbf{X}^*)} \sqrt{f(\mathbf{X}) - f(\mathbf{X}^*)}, \frac{3\sqrt{2}}{\sqrt{\alpha} \lambda_{\min}(\mathbf{X}^*)} \sqrt{f(\mathbf{X}) - f(\mathbf{X}^*)}\right\} + \xi. \end{aligned}$$

*Proof.* For the sake of clarity, throughout the proof we treat each  $\mathbf{v}_i$  as the result of an exact eigenvector computation, i.e., we assume  $\xi = 0$ , and at the end we discuss the effect of the approximation error in the computation of  $\mathbf{v}_i$ .

Let  $\mathbf{w}^* \in \arg \min_{\mathbf{w}: \|\mathbf{w}\|=1} \mathbf{w}^\top \nabla f(\mathbf{X}) \mathbf{w}$ . Using the optimality of  $\mathbf{v}_i$ , and the fact that for every  $i \in [k]$ , both  $\mathbf{v}_i$  and  $\mathbf{x}_i$  are unit vectors, we have that

$$\begin{aligned}
& \sum_{i=1}^k a_i \left[ (\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \leq \\
& \sum_{i=1}^k a_i \left[ (\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \leq \\
& \sum_{i=1}^k a_i \left[ (\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \eta\beta \right] = \\
& (\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \eta\beta \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \eta\beta \leq -(f(\mathbf{X}) - f(\mathbf{X}^*)) + \eta\beta,
\end{aligned} \tag{19}$$

where the third inequality follows from the optimality of  $\mathbf{w}^*$  (note that the minimizer of a linear function  $\mathbf{A} \bullet \mathbf{X}$  when  $\mathbf{X} \in \mathcal{S}_d$  and  $\mathbf{A} \in \mathbb{S}_d$ , is w.l.o.g. a rank-one matrix corresponding to a leading eigenvector of  $\mathbf{A}$ ), and the last inequality follows from the convexity of  $f$ . Thus, Eq. (19) gives us the first part of the bound stated in the lemma. We now move on to prove the second part.

From Lemma 1 we know we can write  $\mathbf{X}^*$  in the following way:

$$\mathbf{X}^* = \sum_{i=1}^k b_i^* \mathbf{y}_i^* \mathbf{y}_i^{*\top} + \sum_{j=1}^k (a_j - b_j^*) \mathbf{W}^*, \tag{20}$$

where for all  $i \in [k]$ ,  $b_i^* \in [0, a_i]$  and  $\mathbf{y}_i^*$  is a unit vector, and  $\mathbf{W}^* \in \mathcal{S}_d$ .

Using again the optimality of  $\mathbf{v}_i$  for each  $i \in [k]$ , we have that

$$\begin{aligned}
& \sum_{i=1}^k a_i \left[ (\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{v}_i \mathbf{v}_i^\top - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \leq \\
& \sum_{i=1}^k a_i \cdot \min \left\{ (\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2, \right. \\
& \left. (\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right\} \leq \\
& \sum_{i=1}^k b_i^* \left[ (\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \\
& + \sum_{i=1}^k (a_i - b_i^*) \left[ (\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \\
& \leq \sum_{i=1}^k b_i^* \left[ (\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right] \\
& + \sum_{i=1}^k (a_i - b_i^*) \left[ (\mathbf{W}^* - \mathbf{x}_i \mathbf{x}_i^\top) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \|\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \right],
\end{aligned} \tag{21}$$

where the second inequality follows since  $\min\{a, b\} \leq \lambda a + (1 - \lambda)b$  for any  $a, b \in \mathbb{R}, \lambda \in [0, 1]$ , and the third inequality follows from the optimality of  $\mathbf{w}^*$ . Using Eq. (20) we have that

$$\begin{aligned}
\text{RHS of (21)} &\leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \sum_{i=1}^k b_i^* \frac{\eta\beta}{2} \|\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \\
&\quad + \sum_{i=1}^k (a_i - b_i^*) \frac{\eta\beta}{2} \|\mathbf{w}^* \mathbf{w}^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 \\
&\leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \frac{\eta\beta}{2} \sum_{i=1}^k b_i^* \|\mathbf{y}_i^* \mathbf{y}_i^{*\top} - \mathbf{x}_i \mathbf{x}_i^\top\|_F^2 + \eta\beta \sum_{i=1}^k (a_i - b_i^*) \\
&\leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \eta\beta \sqrt{\text{rank}(\mathbf{X}^*)} \left( \sqrt{\text{rank}(\mathbf{X}^*)} \tau + \|\mathbf{X} - \mathbf{X}^*\|_F \right) \\
&\quad + \eta\beta \left( \sqrt{\text{rank}(\mathbf{X}^*)} \left( \sqrt{\text{rank}(\mathbf{X}^*)} \tau + \|\mathbf{X} - \mathbf{X}^*\|_F \right) + \gamma \right) \\
&= (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \eta\beta \left( 2\sqrt{\text{rank}(\mathbf{X}^*)} (\|\mathbf{X}^* \mathbf{P}_{\mathbf{X}^*, \tau}^\perp\|_F + \|\mathbf{X} - \mathbf{X}^*\|_F) + \gamma \right), \tag{22}
\end{aligned}$$

where the last inequality follows from plugging the bounds in Lemma 1 and holds for any  $\tau, \gamma \in [0, 1]$  such that  $\frac{\tau\gamma}{1-\gamma} \geq \|\mathbf{X} - \mathbf{X}^*\|_F$ .

Now we can optimize the above bound in terms  $\tau, \gamma$  under the constraint that  $\frac{\tau\gamma}{1-\gamma} \geq \|\mathbf{X} - \mathbf{X}^*\|_F$ . One option is to upper bound  $\|\mathbf{X}^* \mathbf{P}_{\mathbf{X}^*, \tau}^\perp\|_F \leq \sqrt{\text{rank}(\mathbf{X}^*)} \tau$ , which gives us

$$\text{RHS of (21)} \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \eta\beta \left( 2\sqrt{\text{rank}(\mathbf{X}^*)} \left( \sqrt{\text{rank}(\mathbf{X}^*)} \tau + \|\mathbf{X} - \mathbf{X}^*\|_F \right) + \gamma \right).$$

We can then set:

$$\tau_1 = \sqrt{\frac{\|\mathbf{X} - \mathbf{X}^*\|_F}{2\text{rank}(\mathbf{X}^*)}}, \quad \gamma_1 = \sqrt{2\text{rank}(\mathbf{X}^*) \|\mathbf{X} - \mathbf{X}^*\|_F},$$

as long as  $\|\mathbf{X} - \mathbf{X}^*\|_F \leq \frac{1}{2\text{rank}(\mathbf{X}^*)}$ , which gives us:

$$\text{RHS of (21)} \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + 2\eta\beta \sqrt{\text{rank}(\mathbf{X}^*)} \left( \sqrt{2\|\mathbf{X} - \mathbf{X}^*\|_F} + \|\mathbf{X} - \mathbf{X}^*\|_F \right).$$

Note that in order for the above bound to improve over that in Eq. (19), it indeed must in particular hold that  $\|\mathbf{X} - \mathbf{X}^*\|_F < \frac{1}{2\text{rank}(\mathbf{X}^*)}$ . In that case it follows that

$$\text{RHS of (21)} \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + 5\eta\beta \sqrt{\text{rank}(\mathbf{X}^*) \|\mathbf{X} - \mathbf{X}^*\|_F}. \tag{23}$$

Another option, is to choose

$$\tau_2 = \lambda_{\min}(\mathbf{X}^*), \quad \gamma_2 = \frac{\|\mathbf{X} - \mathbf{X}^*\|_F}{\lambda_{\min}(\mathbf{X}^*)},$$

as long as  $\|\mathbf{X} - \mathbf{X}^*\|_F < \lambda_{\min}(\mathbf{X}^*)$ . In this case, it holds that  $\|\mathbf{X}^* \mathbf{P}_{\mathbf{X}^*, \tau}^\perp\|_F = 0$ . Plugging into Eq. (22) we have that

$$\text{RHS of (21)} \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \eta\beta \|\mathbf{X} - \mathbf{X}^*\|_F \left( 2\sqrt{\text{rank}(\mathbf{X}^*)} + \frac{1}{\lambda_{\min}(\mathbf{X}^*)} \right).$$

Note that since  $\mathbf{X}^* \in \mathcal{S}_d$  it holds that  $\lambda_{\min}(\mathbf{X}^*)^{-1} \geq \text{rank}(\mathbf{X}^*)$  and thus we have that

$$\text{RHS of (21)} \leq (\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) + \frac{3\eta\beta \|\mathbf{X} - \mathbf{X}^*\|_F}{\lambda_{\min}(\mathbf{X}^*)}. \tag{24}$$

Note that here also, the above bound improves over the one in Eq. (19) only when indeed  $\|\mathbf{X} - \mathbf{X}^*\|_F < \lambda_{\min}(\mathbf{X}^*)$ .

Now, by using the convexity of  $f$  to upper bound  $(\mathbf{X}^* - \mathbf{X}) \bullet \nabla f(\mathbf{X}) \leq -(f(\mathbf{X}) - f(\mathbf{X}^*))$  and Eq. (1) to upper bound  $\|\mathbf{X} - \mathbf{X}^*\|_F \leq \sqrt{\frac{2}{\alpha}}(f(\mathbf{X}) - f(\mathbf{X}^*))$  in both Eq. (23) and (24), gives the rest of the bound in the lemma.

By going through the analysis above again (basically Eq. (19) and Eq. (21)), it's clear that an  $\xi$  additive error in the computation of each eigenvector  $\mathbf{v}_i$  results in a single additive term  $\xi$  in all of the above bounds, and hence the lemma follows.  $\square$

**Lemma 7.** *[randomized update] Consider an iteration  $t$  of Algorithm 2. Fix a step-size  $\eta_t$  and assume that the iterate of the algorithm on this iteration is feasible and given in the following explicit form:  $\mathbf{X}_t = \sum_{i=1}^k a_i \mathbf{x}_i \mathbf{x}_i^\top$ , where each  $\mathbf{x}_i$  is a unit vector, and  $(a_1, \dots, a_k)$  is a distribution over  $[k]$ . Further, suppose that each  $a_i$  satisfies that  $a_i \geq \eta_t/2$ . Then,*

$$\mathbb{E}[h_{t+1}] \leq \left(1 - \frac{\eta_t}{2}\right) \mathbb{E}[h_t] + \frac{\eta_t^2 \beta}{2} \min\left\{1, \frac{5\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4}} \mathbb{E}[h_t]^{1/4}, \frac{3\sqrt{2}}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} \mathbb{E}[h_t]^{1/2}\right\} + \eta_t \xi_t,$$

where  $\forall t \geq 1$   $h_t := f(\mathbf{X}_t) - f(\mathbf{X}^*)$ .

*Proof.* Using the update step of Algorithm 2 we have that

$$\begin{aligned} h_{t+1} &= f(\mathbf{X}_{t+1}) - f(\mathbf{X}^*) = f(\mathbf{X}_t + \tilde{\eta}_t(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top)) - f(\mathbf{X}^*) \\ &\leq f(\mathbf{X}_t) - f(\mathbf{X}^*) + \tilde{\eta}_t(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\tilde{\eta}_t^2 \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top\|_F^2 \\ &\leq h_t + \tilde{\eta}_t \left[ (\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top\|_F^2 \right], \end{aligned}$$

where the first inequality follows from the smoothness of  $f$  and the second one follows since by definition,  $\eta_t \geq \tilde{\eta}_t$ .

By the choice of  $\mathbf{v}_t$  we have that

$$(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top\|_F^2 \leq \xi_t. \quad (25)$$

To see why this is true, observe that since  $\|\mathbf{v}_t\| = \|\mathbf{x}_{i_t}\| = 1$ , we have that

$$\begin{aligned} &(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top\|_F^2 = \\ &(\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) \bullet \nabla f(\mathbf{X}_t) - \eta_t \beta \mathbf{v}_t \mathbf{v}_t^\top \bullet \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top + \eta_t \beta = \\ &\mathbf{v}_t \mathbf{v}_t^\top \bullet (\nabla f(\mathbf{X}_t) - \eta_t \beta \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top \bullet \nabla f(\mathbf{X}_t) + \eta_t \beta \leq \\ &(\mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top \bullet (\nabla f(\mathbf{X}_t) - \eta_t \beta \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) + \xi_t) - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top \bullet \nabla f(\mathbf{X}_t) + \eta_t \beta = \xi_t, \end{aligned}$$

where the inequality follows from the approximated optimality of  $\mathbf{v}_t$ .

Thus, since by our assumption on  $\{a_i\}_{i \in [k]}$  it also holds that  $\tilde{\eta}_t \geq \eta_t/2$ , we have using Eq. (25) that

$$\begin{aligned} h_{t+1} &\leq h_t + \frac{\eta_t}{2} \left[ (\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top\|_F^2 \right] + \left(\tilde{\eta}_t - \frac{\eta_t}{2}\right) \xi_t \\ &\leq h_t + \frac{\eta_t}{2} \left[ (\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top) \bullet \nabla f(\mathbf{X}_t) + \frac{\eta_t \beta}{2} \|\mathbf{v}_t \mathbf{v}_t^\top - \mathbf{x}_{i_t} \mathbf{x}_{i_t}^\top\|_F^2 \right] + \frac{\eta_t}{2} \xi_t, \quad (26) \end{aligned}$$

where the last inequality follows again by using  $\eta_t \geq \tilde{\eta}_t$ .

Taking expectation over the random choice of  $i_t$  in Eq. (26), and plugging Lemma 6, we have that

$$\mathbb{E}_{i_t}[h_{t+1} | \mathbf{X}_t] \leq h_t - \frac{\eta_t}{2} h_t + \frac{\eta_t^2 \beta}{2} \min\left\{1, \frac{5\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4}} h_t^{1/4}, \frac{3\sqrt{2}}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} h_t^{1/2}\right\} + \frac{\eta_t}{2} \xi_t + \frac{\eta_t}{2} \xi_t.$$

Taking expectation over the randomness introduced on iterations  $1, \dots, t-1$  we have that

$$\begin{aligned}\mathbb{E}[h_{t+1}] &\leq \left(1 - \frac{\eta_t}{2}\right) \mathbb{E}[h_t] + \frac{\eta_t^2 \beta}{2} \min\left\{1, \frac{5\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4}} \mathbb{E}[h_t^{1/4}], \frac{3\sqrt{2}}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} \mathbb{E}[h_t^{1/2}]\right\} + \eta_t \xi_t \\ &\leq \left(1 - \frac{\eta_t}{2}\right) \mathbb{E}[h_t] + \frac{\eta_t^2 \beta}{2} \min\left\{1, \frac{5\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4}} \mathbb{E}[h_t]^{1/4}, \frac{3\sqrt{2}}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} \mathbb{E}[h_t]^{1/2}\right\} + \eta_t \xi_t,\end{aligned}$$

where the first inequality follows since the function  $f(x, y, z) = \min\{x, y, z\}$  is concave, and thus the inequality follows from applying Jensen's inequality. Similarly, the second inequality follows since both functions  $g(x) = x^{1/4}$ ,  $q(x) = x^{1/2}$  are also concave on  $(0, \infty)$ .  $\square$

## D Proof of Theorem 1

We can now turn to prove our main theorem, Theorem 1. The proof follows from deriving each one of the convergence rates in the theorem independently using the result of Lemma 7. This is done in the following Lemmas 8,9, 10. We then show that there exists a choice of step-size sequence and error-tolerance bounds for the eigenvector computations that satisfy all lemmas at once, and thus the theorem is obtained.

**Lemma 8.** *Let  $C, t_0$  be non-negative scalars that satisfy:*

$$C \geq 18, \quad \frac{C}{2} - 1 \geq t_0 \geq \frac{C}{6} - 1.$$

*Then if for all  $t \geq 1$  we define  $\eta_t = \frac{C}{3(t+t_0)}$ , and we set  $\xi_0 = \beta$  and  $\forall t \geq 1: \xi_t = \frac{\beta C}{6(t+t_0)}$ , it follows that all iterates of Algorithm 2 are feasible, and*

$$\forall t \geq 1: \quad \mathbb{E}[h_t] \leq \frac{\beta C}{t + t_0}.$$

*Proof.* From Lemma 7 we have that for all  $t \geq 1$ ,

$$\forall t \geq 1: \quad \mathbb{E}[h_{t+1}] \leq \left(1 - \frac{\eta_t}{2}\right) \mathbb{E}[h_t] + \frac{\eta_t^2 \beta}{2} + \eta_t \xi_t.$$

We are going to assume throughout the proof that  $\xi_t \leq \mathbb{E}[h_t]/6$ . It thus follows that

$$\forall t \geq 1: \quad \mathbb{E}[h_{t+1}] \leq \left(1 - \frac{\eta_t}{3}\right) \mathbb{E}[h_t] + \frac{\eta_t^2 \beta}{2}. \quad (27)$$

For all  $t \geq 1$ , define  $v_t := \beta^{-1} \mathbb{E}[h_t]$ . Dividing both sides of Eq. (27) by  $\beta$ , we have that

$$\forall t \geq 1: \quad v_{t+1} \leq \left(1 - \frac{\eta_t}{3}\right) v_t + \frac{\eta_t^2}{2}. \quad (28)$$

We are going to prove by induction on  $t$  that  $v_t \leq \frac{C}{t+t_0}$  for suitable value of  $C, t_0$  and a sequence of step-sizes  $\{\eta_t\}_{t \geq 1}$ . Obviously for the base case  $t = 1$  to hold, we must restrict  $\frac{C}{t_0+1} \geq v_1$ .

Let us assume now that the induction hypothesis holds for some  $t \geq 1$ .

Setting  $\eta_t = \frac{C}{3(t+t_0)}$  in Eq. (28) we have that

$$\begin{aligned}v_{t+1} &\leq v_t \left(1 - \frac{C}{9(t+t_0)}\right) + \frac{C^2}{18(t+t_0)^2} \leq \frac{C}{t+t_0} \left(1 - \frac{C}{9(t+t_0)}\right) + \frac{C^2}{18(t+t_0)^2} \\ &= \frac{C}{t+t_0} \left(1 - \frac{C}{18(t+t_0)}\right) = \frac{C}{t+t_0+1} \left(1 + \frac{1}{t+t_0}\right) \left(1 - \frac{C}{18(t+t_0)}\right).\end{aligned}$$

Thus, choosing  $C \geq 18$  gives:

$$v_{t+1} \leq \frac{C}{t+1+t_0} \left(1 + \frac{1}{t+t_0}\right) \left(1 - \frac{1}{t+t_0}\right) < \frac{C}{t+1+t_0}$$

as needed.

We can now set values for  $C, t_0$  under the constraints that

$$i. C \geq 18, \quad ii. \frac{C}{t_0+1} \geq v_1, \quad iii. \forall t \geq 1 : \eta_t = \frac{C}{3(t+t_0)} \in [0, 2]. \quad (29)$$

In order for our choice of step-sizes to satisfy the conditions of Observation 1, it must hold that  $\{\eta_t\}_{t \geq 1} \subset [0, 2]$ . Since by definition this sequence is monotonic decreasing it suffices to show it for  $\eta_1$ . Thus we must require that  $\frac{C}{3(1+t_0)} \leq 2$ , which gives us the constraint  $t_0 \geq \frac{C}{6} - 1$ .

It remains to deal with base case of the induction, i.e., we need to show that  $v_1 = \beta^{-1}h_1 \leq \frac{C}{1+t_0}$  for our choice of  $C, t_0$ .

Recall that according to Algorithm 2 it holds that  $\mathbf{X}_1 = \mathbf{x}_1\mathbf{x}_1^\top$ , such that  $\mathbf{x}_1 = \mathbf{E}\mathbf{V}(\nabla f(\mathbf{x}_0\mathbf{x}_0^\top))$ , where  $\mathbf{x}_0$  is some unit vector. Using the smoothness of  $f$  we have that

$$\begin{aligned} h_1 &= f(\mathbf{x}_1\mathbf{x}_1^\top) - f(\mathbf{X}^*) = f(\mathbf{x}_0\mathbf{x}_0^\top + \mathbf{x}_1\mathbf{x}_1^\top - \mathbf{x}_0\mathbf{x}_0^\top) - f(\mathbf{X}^*) \\ &\leq f(\mathbf{x}_0\mathbf{x}_0^\top) - f(\mathbf{X}^*) + (\mathbf{x}_1\mathbf{x}_1^\top - \mathbf{x}_0\mathbf{x}_0^\top) \bullet \nabla f(\mathbf{x}_0\mathbf{x}_0^\top) + \frac{\beta}{2} \|\mathbf{x}_1\mathbf{x}_1^\top - \mathbf{x}_0\mathbf{x}_0^\top\|_F^2 \\ &\leq f(\mathbf{x}_0\mathbf{x}_0^\top) - f(\mathbf{X}^*) + (\mathbf{X}^* - \mathbf{x}_0\mathbf{x}_0^\top) \bullet \nabla f(\mathbf{x}_0\mathbf{x}_0^\top) + \frac{\beta}{2} \|\mathbf{x}_1\mathbf{x}_1^\top - \mathbf{x}_0\mathbf{x}_0^\top\|_F^2 + \xi_0 \\ &\leq \frac{\beta}{2} \|\mathbf{x}_1\mathbf{x}_1^\top - \mathbf{x}_0\mathbf{x}_0^\top\|_F^2 + \xi_0 \leq \beta + \xi_0, \end{aligned} \quad (30)$$

where the second inequality follows from the choice of  $\mathbf{x}_1$ , and the third inequality follows from the convexity of  $f(\mathbf{X})$ .

Setting  $\xi_0 = \beta$ , it follows that

$$v_1 \leq \beta^{-1} \cdot 2\beta = 2.$$

Thus we must require that  $\frac{C}{1+t_0} \geq 2$ , which gives us the constraint  $t_0 \leq \frac{C}{2} - 1$ .

Thus, the conditions in Eq. (29) boils down to the following constraints:

$$C \geq 18, \quad \frac{C}{2} - 1 \geq t_0 \geq \frac{C}{6} - 1.$$

For  $C, t_0$  that indeed satisfy these constraints we can thus conclude that

$$\forall t \geq 1 : \mathbb{E}[h_t] \leq \beta v_t \leq \frac{\beta C}{t+t_0}.$$

□

**Lemma 9.** *Let  $C, t_0$  be non-negative scalars that satisfy:*

$$C \geq 30^{4/3}, \quad C^{3/4} - 1 \geq t_0 \geq \frac{C^{3/4}}{6} - 1.$$

*Then if for all  $t \geq 1$  we define  $\eta_t = \frac{C^{3/4}}{3(t+t_0)}$ , and set  $\xi_0 = \beta$ ,  $\forall t \geq 1 : \xi_t = \frac{1}{6} \left( \frac{5C^{3/4}\beta\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4}(t+t_0)} \right)^{4/3}$ , it follows that all iterates of Algorithm 2 are feasible, and*

$$\forall t \geq 1 : \mathbb{E}[h_t] \leq \left( \frac{5C^{3/4}\beta\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}}{\alpha^{1/4}(t+t_0)} \right)^{4/3}.$$

*Proof.* From Lemma 7 we have that for all  $t \geq 1$ ,

$$\forall t \geq 1: \quad \mathbb{E}[h_{t+1}] \leq \left(1 - \frac{\eta_t}{2}\right) \mathbb{E}[h_t] + \frac{5\eta_t^2 \beta \sqrt{\sqrt{2} \text{rank}(\mathbf{X}^*)}}{2\alpha^{1/4}} \mathbb{E}[h_t]^{1/4} + \eta_t \xi_t.$$

We are going to assume throughout the proof that  $\xi_t \leq \mathbb{E}[h_t]/6$ . It thus follows that

$$\forall t \geq 1: \quad \mathbb{E}[h_{t+1}] \leq \left(1 - \frac{\eta_t}{3}\right) \mathbb{E}[h_t] + \frac{5\eta_t^2 \beta \sqrt{\sqrt{2} \text{rank}(\mathbf{X}^*)}}{2\alpha^{1/4}} \mathbb{E}[h_t]^{1/4}. \quad (31)$$

For all  $t \geq 1$ , define  $v_t := \left(\frac{5\sqrt{\sqrt{2} \text{rank}(\mathbf{X}^*)\beta}}{\alpha^{1/4}}\right)^{-4/3} \mathbb{E}[h_t]$ . Dividing both sides of Eq. (31) by  $\left(\frac{5\sqrt{\sqrt{2} \text{rank}(\mathbf{X}^*)\beta}}{\alpha^{1/4}}\right)^{4/3}$ , we have that

$$\forall t \geq 1: \quad v_{t+1} \leq \left(1 - \frac{\eta_t}{3}\right) v_t + \frac{\eta_t^2}{2} v_t^{1/4}. \quad (32)$$

We are going to prove by induction on  $t$  that  $v_t \leq \frac{C}{(t+t_0)^{4/3}}$  for suitable value of  $C, t_0$  and a sequence of step-sizes  $\{\eta_t\}_{t \geq 1}$ . Obviously for the base case  $t = 1$  to hold, we must restrict  $\frac{C}{(t_0+1)^{4/3}} \geq v_1$ .

Let us assume now that the induction hypothesis holds for some  $t \geq 1$ .

Setting  $\eta_t = \frac{C^{3/4}}{3(t+t_0)}$  in Eq. (32) we have that

$$\begin{aligned} v_{t+1} &\leq v_t \left(1 - \frac{C^{3/4}}{9(t+t_0)}\right) + \frac{C^{3/2}}{18(t+t_0)^2} v_t^{1/4} \\ &\leq \frac{C}{(t+t_0)^{4/3}} \left(1 - \frac{C^{3/4}}{9(t+t_0)}\right) + \frac{C^{7/4}}{18(t+t_0)^{7/3}} \\ &= \frac{C}{(t+t_0)^{4/3}} \left(1 - \frac{C^{3/4}}{18(t+t_0)}\right) \\ &= \frac{C}{(t+1+t_0)^{4/3}} \left(1 + \frac{1}{t+t_0}\right)^{4/3} \left(1 - \frac{C^{3/4}}{18(t+t_0)}\right) \\ &= \frac{C}{(t+1+t_0)^{4/3}} \left(1 + \frac{1}{t+t_0}\right) \left(1 + \frac{1}{t+t_0}\right)^{1/3} \left(1 - \frac{C^{3/4}}{18(t+t_0)}\right) \end{aligned}$$

The single variable function  $g(x) = x^{1/3}$  is concave on  $(0, \infty)$ , and thus,  $g(1+x) \leq g(1) + g'(1) \cdot x = 1 + \frac{x}{3}$ . Using this fact, we have that

$$\begin{aligned} v_{t+1} &\leq \frac{C}{(t+1+t_0)^{4/3}} \left(1 + \frac{1}{t+t_0}\right) \left(1 + \frac{1}{3(t+t_0)}\right) \left(1 - \frac{C^{3/4}}{18(t+t_0)}\right) \\ &< \frac{C}{(t+1+t_0)^{4/3}} \left(1 + \frac{5}{3(t+t_0)}\right) \left(1 - \frac{C^{3/4}}{18(t+t_0)}\right). \end{aligned}$$

Thus, choosing  $C \geq (90/3)^{4/3}$  gives:

$$v_{t+1} \leq \frac{C}{(t+1+t_0)^{4/3}} \left(1 + \frac{5}{3(t+t_0)}\right) \left(1 - \frac{5}{3(t+t_0)}\right) < \frac{C}{(t+1+t_0)^{4/3}},$$

as needed.

We can now set values for  $C, t_0$  under the constraints that

$$i. C \geq 30^{4/3}, \quad ii. \frac{C}{(t_0+1)^{4/3}} \geq v_1, \quad iii. \forall t \geq 1: \eta_t = \frac{C^{3/4}}{3(t+t_0)} \in [0, 2]. \quad (33)$$

As in the proof of Lemma 8 it follows that constraining  $t_0 \geq \frac{C^{3/4}}{6} - 1$ , will result in step-sizes that satisfy the conditions of Observation 1.

Moving to deal with the base case of the induction, again similarly to Lemma 8, we have that

$$\begin{aligned} v_1 &= \left( \frac{\alpha^{1/4}}{5\beta\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}} \right)^{4/3} h_1 \leq \left( \frac{\alpha^{1/4}}{5\beta\sqrt{\sqrt{2}\text{rank}(\mathbf{X}^*)}} \right)^{4/3} \cdot 2\beta \\ &= \left( \frac{\sqrt{2}\alpha^{1/4}}{5\beta^{1/4}\sqrt{\text{rank}(\mathbf{X}^*)}} \right)^{4/3} < 1, \end{aligned}$$

where the inequality follows since  $\alpha \leq \beta$ . Thus we must require that  $\frac{C}{(1+t_0)^{4/3}} \geq 1$ , which gives us the constraint  $t_0 \leq C^{3/4} - 1$ .

Thus, the conditions in Eq. (33) boils down to the following constraints:

$$C \geq 30^{4/3}, \quad C^{3/4} - 1 \geq t_0 \geq \frac{C^{3/4}}{6} - 1.$$

For  $C, t_0$  that indeed satisfy these constraints we can thus conclude that

$$\forall t \geq 1: \quad \mathbb{E}[h_t] \leq \left( \frac{5\beta\sqrt{\text{rank}(\mathbf{X}^*)}}{2^{3/4}\alpha^{1/4}} \right)^{4/3} v_t \leq \left( \frac{5C^{3/4}\beta\sqrt{\text{rank}(\mathbf{X}^*)}}{2^{3/4}\alpha^{1/4}(t+t_0)} \right)^{4/3}.$$

□

**Lemma 10.** *Let  $C, t_0$  be non-negative scalars that satisfy:*

$$C \geq 2916, \quad C^{1/2} - 1 \geq t_0 \geq \frac{C^{1/2}}{6} - 1.$$

*Then if for all  $t \geq 1$  we define  $\eta_t = \frac{C^{1/2}}{3(t+t_0)}$  and  $\xi_0 = \beta$ ,  $\forall t \geq 1: \xi_t = \frac{1}{6} \left( \frac{3\sqrt{2C}\beta}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)(t+t_0)} \right)^2$ , it follows that all iterates of Algorithm 2 are feasible, and*

$$\forall t \geq 1: \quad \mathbb{E}[h_t] \leq \left( \frac{3\sqrt{2C}\beta}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)(t+t_0)} \right)^2.$$

*Proof.* From Lemma 7 we have that for all  $t \geq 1$ ,

$$\forall t \geq 1: \quad \mathbb{E}[h_{t+1}] \leq \left(1 - \frac{\eta_t}{2}\right) \mathbb{E}[h_t] + \frac{3\sqrt{2}\eta_t^2\beta}{2\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} \mathbb{E}[h_t]^{1/2} + \eta_t\xi_t.$$

We are going to assume throughout the proof that  $\xi_t \leq \mathbb{E}[h_t]/6$ . It thus follows that

$$\forall t \geq 1: \quad \mathbb{E}[h_{t+1}] \leq \left(1 - \frac{\eta_t}{3}\right) \mathbb{E}[h_t] + \frac{3\sqrt{2}\eta_t^2\beta}{2\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} \mathbb{E}[h_t]^{1/2}. \quad (34)$$

For all  $t \geq 1$ , define  $v_t := \left( \frac{3\sqrt{2}\beta}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} \right)^{-2} \mathbb{E}[h_t]$ . Dividing both sides of Eq. (31) by  $\left( \frac{3\sqrt{2}\beta}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)} \right)^2$ , we have that

$$\forall t \geq 1: \quad v_{t+1} \leq \left(1 - \frac{\eta_t}{3}\right) v_t + \frac{\eta_t^2}{2} v_t^{1/2}. \quad (35)$$

We are going to prove by induction on  $t$  that  $v_t \leq \frac{C}{(t+t_0)^2}$  for suitable value of  $C, t_0$  and a sequence of step-sizes  $\{\eta_t\}_{t \geq 1}$ . Obviously for the base case  $t = 1$  to hold, we must restrict  $\frac{C}{(t_0+1)^2} \geq v_1$ .

Let us assume now that the induction hypothesis holds for some  $t \geq 1$ .

Setting  $\eta_t = \frac{C^{1/2}}{3(t+t_0)}$  in Eq. (32) we have that

$$\begin{aligned}
v_{t+1} &\leq v_t \left(1 - \frac{C^{1/2}}{9(t+t_0)}\right) + \frac{C}{18(t+t_0)^2} v_t^{1/2} \\
&\leq \frac{C}{(t+t_0)^2} \left(1 - \frac{C^{1/2}}{9(t+t_0)}\right) + \frac{C^{3/2}}{18(t+t_0)^3} \\
&= \frac{C}{(t+1+t_0)^2} \left(1 + \frac{1}{t+t_0}\right)^2 \left(1 - \frac{C^{1/2}}{18(t+t_0)}\right) \\
&\leq \frac{C}{(t+1+t_0)^2} \left(1 + \frac{3}{t+t_0}\right) \left(1 - \frac{C^{1/2}}{18(t+t_0)}\right)
\end{aligned}$$

Thus, choosing  $C \geq 2916$  gives:

$$v_{t+1} \leq \frac{C}{(t+1+t_0)^2} \left(1 + \frac{3}{t+t_0}\right) \left(1 - \frac{3}{t+t_0}\right) < \frac{C}{(t+1+t_0)^2},$$

as needed.

We can now set values for  $C, t_0$  under the constraints that

$$i. C \geq 2916 \quad ii. \frac{C}{(t_0+1)^2} \geq v_1, \quad iii. \forall t \geq 1 : \eta_t = \frac{C^{1/2}}{3(t+t_0)} \in [0, 2]. \quad (36)$$

As in Lemma 8, it follows that in order for our step-sizes satisfy the conditions of Observation 1, we need to require that  $t_0 \geq \frac{C^{1/2}}{6} - 1$ .

Also, for the base case of the induction, also similarly to Lemma 8, it holds that

$$v_1 \leq \left(\frac{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)}{3\sqrt{2}\beta}\right)^2 \cdot 2\beta = \left(\frac{\sqrt{2\alpha}\lambda_{\min}(\mathbf{X}^*)}{3\sqrt{2}\sqrt{\beta}}\right)^2 < 1$$

where the second inequality follows since  $\alpha \leq \beta$  and  $\lambda_{\min}(\mathbf{X}^*) \leq 1$ . Thus in order to satisfy the constraint  $v_1 \leq \frac{C}{(t_0+1)^2}$ , it suffices to require  $t_0 \leq \sqrt{C} - 1$ .

Thus, the conditions in Eq. (33) boils down to the following constraints:

$$C \geq 2916, \quad C^{1/2} - 1 \geq t_0 \geq \frac{C^{1/2}}{6} - 1.$$

For  $C, t_0$  that indeed satisfy these constraints we can thus conclude that We can thus conclude that

$$\forall t \geq 1 : \mathbb{E}[h_t] \leq \left(\frac{3\sqrt{2}\beta}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)}\right)^2 v_t \leq \left(\frac{3\sqrt{2C}\beta}{\sqrt{\alpha}\lambda_{\min}(\mathbf{X}^*)(t+t_0)}\right)^2.$$

□

We can now finally wrap-up the proof of Theorem 1.

*Proof.* The proof is an immediate consequence of Lemmas 8, 9, 10, and the observation that the step-size  $\eta_t = \frac{54}{3(t+8)} = \frac{18}{t+8}$ , which implicitly sets  $t_0 = 8$  in all of above lemmas and corresponds to setting  $C = 54$  for Lemma 8,  $C = 54^{4/3}$  in Lemma 9, and  $C = 2916$  in Lemma 10, satisfies all lemmas together. □