

Improved Complexities of Conditional Gradient-Type Methods with Applications to Robust Matrix Recovery Problems

Dan Garber, Atara Kaplan and Shoham Sabach

Received: date / Accepted: date

Abstract Motivated by robust matrix recovery problems such as Robust Principal Component Analysis, we consider a general optimization problem of minimizing a smooth and strongly convex loss function applied to the sum of two blocks of variables, where each block of variables is constrained or regularized individually. We study a Conditional Gradient-Type method which is able to leverage the special structure of the problem to obtain faster convergence rates than those attainable via standard methods, under a variety of assumptions. In particular, our method is appealing for matrix problems in which one of the blocks corresponds to a low-rank matrix since it avoids prohibitive full-rank singular value decompositions required by most standard methods. While our initial motivation comes from problems which originated in statistics, our analysis does not impose any statistical assumptions on the data.

Keywords conditional gradient method · Frank-Wolfe algorithm · convex optimization · robust PCA · low-rank matrix recovery · low-rank optimization · semidefinite programming · nuclear norm minimization

Mathematics Subject Classification (2000) 90C22 · 90C06 · 68W27 · 68W20

Dan Garber
Faculty of Industrial Engineering and Management, Technion, 32000 Haifa, Israel
E-mail: dangar@technion.ac.il

Atara Kaplan
Faculty of Mathematics, Technion, 32000 Haifa, Israel
E-mail: ataragold@campus.technion.ac.il

Shoham Sabach
Department of Industrial Engineering and Management, Technion, 32000 Haifa, Israel
E-mail: ssabach@ie.technion.ac.il

1 Introduction

In this paper we consider the following general convex optimization problem

$$\min \{f(\mathbf{X}, \mathbf{Y}) := g(\mathbf{X} + \mathbf{Y}) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}) : \mathbf{X}, \mathbf{Y} \in \mathbf{E}\}, \quad (1)$$

where \mathbf{E} is a finite-dimensional normed vector space over the reals, $g : \mathbf{E} \rightarrow \mathbb{R}$ is assumed to be continuously differentiable and strongly convex, while $\mathcal{R}_{\mathcal{X}} : \mathbf{E} \rightarrow (-\infty, +\infty]$ and $\mathcal{R}_{\mathcal{Y}} : \mathbf{E} \rightarrow (-\infty, +\infty]$ are proper, lower semicontinuous and convex functions which can be thought of either as regularization functions, or indicator functions¹ of certain closed and convex feasible sets \mathcal{X} and \mathcal{Y} .

Problem (1) captures several important problems of interest, perhaps the most well-studied is that of *Robust Principal Component Analysis (PCA)* [3, 14, 11], in which the goal is to (approximately) decompose an $m \times n$ input matrix \mathbf{M} into the sum of a low-rank matrix \mathbf{X} and a sparse matrix \mathbf{Y} . The underlying optimization problem for Robust PCA can be written as (see, for instance, [11])

$$\min \left\{ \frac{1}{2} \|\mathbf{X} + \mathbf{Y} - \mathbf{M}\|_F^2 : \|\mathbf{X}\|_{\text{nuc}} \leq \tau, \|\mathbf{Y}\|_1 \leq s, \mathbf{X}, \mathbf{Y} \in \mathbb{R}^{m \times n} \right\}, \quad (2)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, $\|\cdot\|_{\text{nuc}}$ denotes the nuclear norm, i.e., the sum of singular values, which is a highly popular convex surrogate for low-rank penalty, and $\|\cdot\|_1$ is the entry-wise ℓ_1 -norm, which is a well-known convex surrogate for entry-wise sparsity.

Other variants of interest of Problem (2) are when the data matrix \mathbf{M} is a corrupted covariance matrix, in which case it is reasonable to further constrain \mathbf{X} to be positive semidefinite, i.e., use the constraints $\mathbf{X} \succeq \mathbf{0}$ and $\text{Tr}(\mathbf{X}) \leq \tau$. In the case that \mathbf{M} is assumed to have several fully corrupted rows or columns, a popular alternative to the ℓ_1 -norm regularizer on the variable \mathbf{Y} is to use either the norm $\|\cdot\|_{1,2}$ (sum of ℓ_2 -norm of rows) in case of corrupted rows, or the norm $\|\cdot\|_{2,1}$ (sum of ℓ_2 -norm of columns) in case of corrupted columns, as a regularizer/constraint [15]. Finally, moving beyond Robust PCA, a different choice of interest for the loss $g(\cdot)$ could be $g(\mathbf{Z}) := (1/2) \|\mathcal{A}\mathbf{Z} - \mathbf{M}\|_F^2$, where \mathcal{A} is a linear sensing operator such that $\mathcal{A}^T \mathcal{A}$ is positive definite (so $g(\cdot)$ is strongly convex).

In this paper we present an algorithm and analyses that build on the special structure of Problem (1), which improve upon state-of-the-art complexity bounds, under several different assumptions. A common key to all of our results is the ability to exploit the strong convexity of $g(\cdot)$ to obtain improved complexity bounds. Here it should be noted that while $g(\cdot)$ is assumed to be strongly convex, Problem (1) is in general not strongly convex in (\mathbf{X}, \mathbf{Y}) . This can already be observed when choosing $g(\mathbf{z}) := \frac{1}{2} \|\mathbf{z}\|_2^2$, and

¹ An indicator function of a set is defined to be 0 in the set and $+\infty$ outside.

$\mathcal{R}_{\mathcal{X}}(\cdot) = \mathcal{R}_{\mathcal{Y}}(\cdot) = 0$, where $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. In this case, denoting the overall objective as $f(\mathbf{x}, \mathbf{y}) := \frac{1}{2} \|\mathbf{x} + \mathbf{y}\|_2^2$, it is easily observed that the Hessian matrix of $f(\cdot, \cdot)$ is given by $\nabla^2 f(\mathbf{x}, \mathbf{y}) = (\mathbf{I} \ \mathbf{I})^\top (\mathbf{I} \ \mathbf{I})$, and hence is not full-rank.

The fastest known convergence rate for first-order methods applicable to Problem (1), is achievable by accelerated gradient methods such as Nesterov’s optimal method [12] and FISTA [2], which converge at a rate of $O(1/t^2)$, where t denoted the number of performed iterations. However, in the context of low-rank matrix optimization problems such as Robust PCA, these methods require to compute a full-rank singular value decomposition on each iteration to update the low-rank component, which is often prohibitive for large scale instances. A different type of first-order methods is the Conditional Gradient (CG) Method (a.k.a Frank-Wolfe algorithm) and variants of [6–10, 17]. In the context of low-rank matrix optimization, the CG method simply requires to compute an approximate leading singular vector pair of the negative gradient at each iteration, i.e., a rank-one SVD. Hence, in this case, the CG method is much more scalable, than projection/proximal based methods. However, the rate of convergence is slower, e.g., if both $\mathcal{R}_{\mathcal{X}}(\cdot)$ and $\mathcal{R}_{\mathcal{Y}}(\cdot)$ are indicator functions of certain closed and convex sets \mathcal{X} and \mathcal{Y} , then the convergence rate of the conditional gradient method is of the form $O((D_{\mathcal{X}}^2 + D_{\mathcal{Y}}^2)/t)$, where $D_{\mathcal{X}}$ and $D_{\mathcal{Y}}$ denote the Euclidean diameter of the corresponding feasible sets \mathcal{X} and \mathcal{Y} , where the diameter of a subset C of \mathbb{R}^d is defined by $D_C = \max_{\mathbf{x}_1, \mathbf{x}_2 \in C} \|\mathbf{x}_1 - \mathbf{x}_2\|_2$.

Recently, two variants of the conditional gradient method for low-rank matrix optimization were suggested, which enjoy faster convergence rates when the optimal solution has low rank (which is indeed a key implicit assumption in such problems), while requiring to compute only a single low-rank SVD on each iteration [5, 1]. However, both of these new methods require the objective function to be strongly convex, which as we discussed above, does not hold in our case. Nevertheless, both our algorithm and our analysis are inspired by these two works. In particular, we generalize the low-rank SVD approach of [1] to non-strongly-convex problems of the form of Problem (1), which include arbitrary regularizers or constraints.

In another recent related work [11], which also serves as a motivation for this current work, the authors considered a variant of the conditional gradient method tailored for low-rank and robust matrix recovery problems such as Problem (2), which combines standard conditional gradient updates of the low-rank variable (i.e., rank-one SVD) and proximal gradient updates for the sparse noisy component. However, both the worst-case convergence rate and running time do not improve over the standard conditional gradient method. Combining conditional gradient and proximal gradient updates for low-rank models was also considered in [4] for solving a convex optimization problem related to temporal recommendation systems.

Finally, it should be noted that while developing efficient *non-convex optimization*-based algorithms for Robust PCA with provable guarantees is an

active subject (see, e.g., [13,16]), such works fall short in two aspects: (a) they are not flexible as the general model (1), which allows for instance to impose a PSD constraint on the low-rank component or to consider various sparsity-promoting regularizers for the sparse component \mathbf{Y} , and (b) all provable guarantees are heavily based on assumptions on the input matrix \mathbf{M} (such as incoherence of the singular value decomposition of the low-rank component or assuming certain patterns of the sparse component), which can be quite limiting in practice. This work, on the other hand, is completely free of such assumptions.

To overcome the shortcomings of previous methods applicable to Problem (1), in this paper we present a first-order method, which combines two well-known ideas, for tackling Problem (1). In particular we show that under several assumptions of interest, despite the fact that the objective in Problem (1) is in general not strongly convex, it is possible to leverage the strong convexity of $g(\cdot)$ towards obtaining better complexity results, while applying update steps that are scalable to large scale problems. Informally speaking, our main improved complexity bounds are as follows:

1. In the case that both $\mathcal{R}_{\mathcal{X}}(\cdot)$ and $\mathcal{R}_{\mathcal{Y}}(\cdot)$ are indicators of compact and convex sets (as in Problem (2)), we obtain convergence rate of $O(\min\{D_{\mathcal{X}}^2, D_{\mathcal{Y}}^2\}/t)$. In particular when \mathbf{X} is constrained, for example, via a low-rank promoting constraint, such as the nuclear-norm, our method requires on each iteration only a SVD computation of $\text{rank}=\text{rank}(\mathbf{X}^*)$, where \mathbf{X}^* is part of certain optimal solution $(\mathbf{X}^*, \mathbf{Y}^*)$. This result improves (in terms of running time), in a wide regime of parameters, mainly when $\min\{D_{\mathcal{X}}^2, D_{\mathcal{Y}}^2\} \ll \max\{D_{\mathcal{X}}^2, D_{\mathcal{Y}}^2\}$, over the conditional gradient method which converges with rate of $O(\max\{D_{\mathcal{X}}^2, D_{\mathcal{Y}}^2\}/t)$, and over accelerated gradient methods which require, in the context of low-rank matrix optimization problems, a full-rank SVD computation on each iteration.
2. In the case that $\mathcal{R}_{\mathcal{Y}}(\cdot)$ is an indicator of a strongly convex set (e.g., an ℓ_p -norm ball for $p \in (1, 2]$), our method achieves a fast convergence rate of $O(1/t^2)$. As in the previous case, if \mathbf{X} is constrained/regularized via the nuclear norm, then our method only requires a SVD computation of $\text{rank}=\text{rank}(\mathbf{X}^*)$. To the best of our knowledge, this is the first result that combines an $O(1/t^2)$ convergence rate and low-rank SVD computations in this setting. In particular, in the context of Robust PCA, such a result allows us to replace a traditional sparsity-promoting constraint of the form $\|\mathbf{Y}\|_1 \leq \tau$ with $\|\mathbf{Y}\|_{1+\delta} \leq \tau'$, for some small constant δ . Using the $\ell_{1+\delta}$ -norm instead of the ℓ_1 -norm gives rise to a strongly convex feasible set and, as we demonstrate empirically in Section 3.2, may provide a satisfactory approximation to the ℓ_1 -norm constraint in terms of sparsity.
3. In the case that either $\mathcal{R}_{\mathcal{X}}(\cdot)$ or $\mathcal{R}_{\mathcal{Y}}(\cdot)$ are strongly convex (though not necessarily differentiable), our method achieves a linear convergence rate. In fact, we show that even if only one of the variables is regularized by a strongly convex function, then the entire objective of Problem (1) becomes

setting	Cond. Grad.[10]		FISTA [2]		Algorithm 1	
	rate	SVD rank	rate	SVD rank	rate	SVD rank
$\tau \gg s$ (“high SNR regime”)	τ^2/t	1	τ^2/t^2	n	s^2/t	$\text{rank}(\mathbf{X}^*)$
$\tau \ll s$ (“low SNR regime”)	s^2/t	1	s^2/t^2	n	τ^2/t	1
$\mathcal{Y} := \{\mathbf{Y} : \ \mathbf{Y}\ _{1+\delta} \leq s\}$	$\frac{\max\{s,\tau\}^2}{t}$	1	$\frac{\max\{s,\tau\}^2}{t^2}$	n	$\frac{s^2 n^{\frac{1-\delta}{1+\delta}}}{t^2}$	$\text{rank}(\mathbf{X}^*)$
$\mathcal{R}_{\mathcal{Y}}(\mathbf{Y}) = \lambda_1 \ \mathbf{Y}\ _1 + \lambda_2 \ \mathbf{Y}\ _F^2$	$1/t$	1	$e^{-\Theta(\sqrt{\lambda_2}t)}$	n	$e^{-\Theta(\lambda_2 t)}$	$\text{rank}(\mathbf{X}^*)$

Table 1: Comparison of convergence rates and iteration complexity bounds for Robust PCA problems (see Problem (2)) with a $n \times n$ input matrix \mathbf{M} . For all methods the computational bottleneck is a single SVD computation to update the variable \mathbf{X} , hence we focus on the rank of the required SVD. For clarity of presentation the results are given in simplified form. The dependence on n in the rate for Algorithm 1 in the third row comes from the strong convexity parameter of the set \mathcal{Y} .

strongly convex in (\mathbf{X}, \mathbf{Y}) . Here also, in the case of a nuclear norm constraint/regularization on one of the variables, we are able to leverage the use of only low-rank SVD computations. In the context of Robust PCA such a natural strongly convex regularizer may arise by replacing the ℓ_1 -norm regularization on \mathbf{Y} with the elastic net regularizer, which combines both the ℓ_1 -norm and the squared ℓ_2 -norm, and serves as a popular alternative to the ℓ_1 -norm regularizer in LASSO.

A quick summary of the above results in the context of Robust PCA problems, such as Problem (2), is given in Table 1. See Section 3.2 in the sequel for a detailed discussion.

2 Preliminaries

Throughout the paper we let \mathbf{E} denote an arbitrary finite-dimensional normed vector space over \mathbb{R} where $\|\cdot\|$ and $\|\cdot\|_*$ denote the primal and dual norms over \mathbf{E} , respectively.

2.1 Smoothness and strong convexity of functions and sets

Definition 1 (smooth function) Let $f : \mathbf{E} \rightarrow \mathbb{R}$ be a continuously differentiable function over a convex set $\mathcal{K} \subseteq \mathbf{E}$. We say that f is β -smooth over \mathcal{K} with respect to $\|\cdot\|$, if for all $\mathbf{x}, \mathbf{y} \in \mathcal{K}$ it holds that $f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \mathbf{y} - \mathbf{x}, \nabla f(\mathbf{x}) \rangle + (\beta/2) \|\mathbf{x} - \mathbf{y}\|^2$.

Definition 2 (strongly convex function) Let $f : \mathbf{E} \rightarrow \mathbb{R}$ be a continuously differentiable function over a convex set $\mathcal{K} \subseteq \mathbf{E}$. We say that f is α -strongly convex over \mathcal{K} with respect to $\|\cdot\|$, if it satisfies for all $\mathbf{x}, \mathbf{y} \in \mathcal{K}$ that $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \mathbf{y} - \mathbf{x}, \nabla f(\mathbf{x}) \rangle + (\alpha/2) \|\mathbf{x} - \mathbf{y}\|^2$.

The above definition combined with the first-order optimality condition implies that for a continuously differentiable and α -strongly convex function f , if $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{K}} f(\mathbf{x})$, then for any $\mathbf{x} \in \mathcal{K}$ it holds that $f(\mathbf{x}) - f(\mathbf{x}^*) \geq (\alpha/2) \|\mathbf{x} - \mathbf{x}^*\|^2$.

This last inequality further implies that the magnitude of the gradient of f at a point \mathbf{x} , $\|\nabla f(\mathbf{x})\|_*$ is at least of the order of the square-root of the objective value approximation error at \mathbf{x} , that is, $f(\mathbf{x}) - f(\mathbf{x}^*)$. Indeed, this follows since

$$\begin{aligned} \sqrt{\frac{2}{\alpha}} (f(\mathbf{x}) - f(\mathbf{x}^*)) \cdot \|\nabla f(\mathbf{x})\|_* &\geq \|\mathbf{x} - \mathbf{x}^*\| \cdot \|\nabla f(\mathbf{x})\|_* \\ &\geq \langle \mathbf{x} - \mathbf{x}^*, \nabla f(\mathbf{x}) \rangle \\ &\geq f(\mathbf{x}) - f(\mathbf{x}^*), \end{aligned}$$

where the second inequality follows from Holder's inequality and the third from the convexity of f . Thus, at any point $\mathbf{x} \in \mathcal{K}$, it holds that

$$\|\nabla f(\mathbf{x})\|_* \geq \sqrt{\frac{\alpha}{2}} \cdot \sqrt{f(\mathbf{x}) - f(\mathbf{x}^*)}. \quad (3)$$

Definition 3 (strongly convex set) We say that a convex set $\mathcal{K} \subset \mathbf{E}$ is α -strongly convex with respect to $\|\cdot\|$ if for any $\mathbf{x}, \mathbf{y} \in \mathcal{K}$, any $\gamma \in [0, 1]$ and any vector $\mathbf{z} \in \mathbf{E}$ such that $\|\mathbf{z}\| = 1$, it holds that $\gamma\mathbf{x} + (1-\gamma)\mathbf{y} + \gamma(1-\gamma)(\alpha/2)\|\mathbf{x} - \mathbf{y}\|^2 \mathbf{z} \in \mathcal{K}$. That is, \mathcal{K} contains a ball of radius $\gamma(1-\gamma)(\alpha/2)\|\mathbf{x} - \mathbf{y}\|^2$ induced by the norm $\|\cdot\|$ centered at $\gamma\mathbf{x} + (1-\gamma)\mathbf{y}$.

For more details on strongly convex sets, examples and connections to optimization, we refer the reader to [6].

3 Algorithm and Results

As discussed in the introduction, in this paper we study efficient algorithms for the minimization model (1), where, throughout the paper, our blanket assumption is as follows

Assumption 1 – $g : \mathbf{E} \rightarrow \mathbb{R}$ is β -smooth and α -strongly convex.
– $\mathcal{R}_{\mathcal{X}} : \mathbf{E} \rightarrow (-\infty, +\infty]$ and $\mathcal{R}_{\mathcal{Y}} : \mathbf{E} \rightarrow (-\infty, +\infty]$ are proper, lower semi-continuous and convex functions.

It should be noted that since $\mathcal{R}_{\mathcal{X}}(\cdot)$ (similarly for $\mathcal{R}_{\mathcal{Y}}(\cdot)$) is assumed to be extended-valued function, it allows the inclusion of a constraint through the indicator function of the corresponding constraint set. Indeed, in this case one will consider $\mathcal{R}_{\mathcal{X}}(\mathbf{X}) := \iota_{\mathcal{X}}(\mathbf{X})$, where $\mathcal{X} \subset \mathbf{E}$ is a nonempty, closed and convex.

We now present the main algorithmic framework, which will be used to derive all of our results.

Algorithm 1 Alternating Conditional Gradient Proximal Gradient Method

```

1: input:  $\{\eta_t\}_{t \geq 1} \subset [0, 1]$  - a sequence of step-sizes.
2:  $\mathbf{X}_1$  is an arbitrary point in  $\text{dom } \mathcal{R}_X$ ,  $\mathbf{Y}_1$  is an arbitrary point in  $\text{dom } \mathcal{R}_Y$ .
3: for  $t = 1, 2, \dots$  do
4:    $\mathbf{W}_t = \arg \min_{\mathbf{W} \in \mathbf{E}} \{ \mathcal{R}_Y(\mathbf{W}) + \langle \mathbf{W}, \nabla g(\mathbf{X}_t + \mathbf{Y}_t) \rangle \}$ ,
5:    $\mathbf{V}_t = \arg \min_{\mathbf{V} \in \mathbf{E}} \left\{ \phi_t(\mathbf{V}) := \mathcal{R}_X(\mathbf{V}) + \langle \mathbf{V}, \nabla g(\mathbf{X}_t + \mathbf{Y}_t) \rangle + \frac{\eta_t \beta}{2} \|\mathbf{V} + \mathbf{W}_t - (\mathbf{X}_t + \mathbf{Y}_t)\|^2 \right\}$ ,
      {in fact it suffices that  $\phi_t(\mathbf{V}_t) \leq \phi_t(\mathbf{X}^*)$  for some optimal solution  $(\mathbf{X}^*, \mathbf{Y}^*)$ }
6:    $(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}) = (1 - \eta_t)(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t(\mathbf{V}_t, \mathbf{W}_t)$ ,
7: end for

```

Algorithm 1 is based on three well-known corner stones in continuous optimization: alternating minimization, conditional gradient, and proximal gradient. Since Problem (1) involves two variables \mathbf{X} and \mathbf{Y} , we update each one of them separately and differently in an alternating fashion. Indeed, the \mathbf{Y} variable is first updated using a conditional gradient step (see step (4)) and then the alternating idea comes into a play and we use the updated information in order to update the \mathbf{X} variable using a proximal gradient step (see step (5))².

3.1 Outline of the main results

Let us denote by f^* the optimal value of the optimization Problem (1). In the sequel we prove the following three theorems on the performance of Algorithm 1. For clarity, below we present a concise and simplified version of the results. In section 4, in which we provide complete proofs for these theorems, we also restate them with complete detail. In all three theorems we assume that Assumption 1 holds true, and we bound the convergence rate of the sequence $\{(\mathbf{X}_t, \mathbf{Y}_t)\}_{t \geq 1}$ produced by Algorithm 1 with a suitable choice of step-sizes $\{\eta_t\}_{t \geq 1}$.

Theorem 1 *Assume that $\mathcal{R}_Y := \iota_Y$ where \mathcal{Y} is a nonempty, closed and convex subset of \mathbf{E} . There exists a choice of step-sizes such that Algorithm 1 converges with a rate of $O(\beta D_Y^2/t)$.*

Remark 1 Note that since \mathbf{X} and \mathbf{Y} are in principle interchangeable, Theorem 1 implies a rate of $O(\beta \min\{D_X^2, D_Y^2\}/t)$. This improves over the rate of $O(\beta \max\{D_X^2, D_Y^2\}/t)$ achieved by standard analyses of projected/proximal gradient methods and the conditional gradient method.

Theorem 2 *Assume $\mathcal{R}_X := \iota_X$ where \mathcal{X} is a nonempty, closed and convex subset of \mathbf{E} and $\mathcal{R}_Y := \iota_Y$, where \mathcal{Y} is a strongly convex and closed*

² We note that a practical implementation of Algorithm 1 for a specific problem, such as Problem (2), may require to account for approximation errors in the computation of \mathbf{W}_t or \mathbf{V}_t , since exact computation is not always practically feasible. Such considerations which can be easily incorporated both into Algorithm 1 and our corresponding analyses (see examples in [10, 5, 1]), are beyond the scope of this current paper, and for the simplicity and clarity of presentation, we assume all such computations are precise.

subset of \mathbf{E} . There exists a choice of step-sizes such that Algorithm 1 converges with a rate of $O(1/t^2)$. Moreover, if there exists $G > 0$ such that $\min_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \|\nabla g(\mathbf{X} + \mathbf{Y})\|_* \geq G$, then using a fixed step-size, Algorithm 1 converges with a rate of $O(\exp(-\Theta(t)))$.

Remark 2 While a rate of $O(1/t^2)$ for the conditional gradient method over strongly convex sets was recently showed to hold in [6], it should be noted that it does not apply in the case of Theorem 2, since only the set \mathcal{Y} is assumed to be strongly convex. In particular, both the set of sums $\mathcal{X} + \mathcal{Y} \subset \mathbf{E}$ and the product set $\mathcal{X} \times \mathcal{Y} \subset \mathbf{E} \times \mathbf{E}$ need not be strongly convex.

Theorem 3 *Assume that $\mathcal{R}_{\mathcal{Y}}(\cdot)$ is strongly convex. Then, there exists a fixed step-size such that Algorithm 1 converges with a rate of $O(\exp(-\Theta(t)))$.*

3.2 Putting our results in the context of Robust PCA problems

As discussed in the Introduction, this work is mostly motivated by low-rank matrix optimization problems such as Robust PCA (see Problem (2)). Thus, towards better understanding of our results for this setting, we now briefly detail the applications to Problem (2). As often standard in such problems, we assume that there exists an optimal solution $(\mathbf{X}^*, \mathbf{Y}^*)$ such that the signal matrix \mathbf{X}^* is of rank at most r^* , where $r^* \ll \min\{m, n\}$ ³.

3.2.1 Using low-rank SVD computations

Note that the computation of \mathbf{V}_t in Algorithm 1, which is used to update the estimate \mathbf{X}_{t+1} , simply requires that \mathbf{V}_t satisfies $\|\mathbf{V}_t\|_{\text{nuc}} \leq \tau$ and

$$\left\| \mathbf{V}_t - \left(\mathbf{X}_t + \mathbf{Y}_t - \mathbf{W}_t - \frac{1}{\eta_t} \nabla_t \right) \right\|_F^2 \leq \left\| \mathbf{X}^* - \left(\mathbf{X}_t + \mathbf{Y}_t - \mathbf{W}_t - \frac{1}{\eta_t} \nabla_t \right) \right\|_F^2, \quad (4)$$

where we use the short notation $\nabla_t := \nabla g(\mathbf{X}_t + \mathbf{Y}_t)$. Since \mathbf{X}^* is assumed to have rank at most r^* , it follows that

$$\text{RHS of (4)} \geq \min_{\mathbf{X} \in C} \left\| \mathbf{X} - \left(\mathbf{X}_t + \mathbf{Y}_t - \mathbf{W}_t - \frac{1}{\eta_t} \nabla_t \right) \right\|_F^2, \quad (5)$$

where $C := \{\mathbf{X} : \|\mathbf{X}\|_{\text{nuc}} \leq \tau, \text{rank}(\mathbf{X}) \leq r^*\}$. The solution to the minimization problem on the RHS of (5) is given simply by computing the rank- r^* singular value decomposition of the matrix $\mathbf{A}_t = (\mathbf{X}_t + \mathbf{Y}_t - \mathbf{W}_t - (1/\eta_t) \nabla_t)$, and projecting the resulted vector of singular values onto the ℓ_1 -norm ball of radius τ (which can be done in $O(r^* \log(r^*))$ time). Thus, indeed the time to compute the update for \mathbf{X}_{t+1} on each iteration of Algorithm 1 is dominated by

³ Our results could be easily extended to the case in which $(\mathbf{X}^*, \mathbf{Y}^*)$ is nearly of rank r^* , i.e., of distance much smaller than the required approximation accuracy ε to a rank r^* matrix, however for the sake of clarity we simply assume that $(\mathbf{X}^*, \mathbf{Y}^*)$ is of low-rank.

a single rank- r^* SVD computation. This observation holds for all the following discussions in this section as well. This low-rank SVD approach was already suggested in the recent work [1], that studied smooth and strongly convex minimization over the nuclear-norm ball (which differs from our setting).

3.2.2 Improved complexities for low/high SNR regimes

In case that $(\mathbf{X}^*, \mathbf{Y}^*)$ is an (say, unique) optimal solution to Problem (2), which satisfies $\|\mathbf{Y}^*\|_F^2 \ll \|\mathbf{X}^*\|_F^2$, i.e., a high signal-to-noise ratio regime, we expect that $D_{\mathcal{X}} \gg D_{\mathcal{Y}}$, where $D_{\mathcal{X}}$ and $D_{\mathcal{Y}}$ are the Euclidean diameters of the nuclear norm ball and the ℓ_1 -norm ball, respectively. In this case, the result of Theorem 1 is appealing since the convergence rate depends only on $D_{\mathcal{Y}}^2$, and not on $D_{\mathcal{X}}^2 + D_{\mathcal{Y}}^2$ as standard algorithms/analyses. In the opposite case, i.e., $D_{\mathcal{X}} \ll D_{\mathcal{Y}}$, which naturally corresponds to a low signal-to-noise ratio regime, since \mathbf{X} and \mathbf{Y} are interchangeable in our setting, we can reverse their roles in the optimization and get via Theorem 1 dependency only on $D_{\mathcal{X}}$. Moreover, now the nuclear-norm constrained variable (assuming the role of \mathbf{Y} in Algorithm 1) is only updated via a conditional gradient update, i.e., requires only a rank-one SVD computation on each iteration. In particular, statistical recovery results such as the seminal work [3], show that under suitable assumptions on the data, exact recovery is possible in both of these cases, even for instance, when $\|\mathbf{Y}^*\|_F^2 / \|\mathbf{X}^*\|_F^2 = \text{poly}(n)$.

3.2.3 Replacing the ℓ_1 constraint with an $\ell_{1+\delta}$ constraint

The ℓ_1 -norm is traditionally used in Robust PCA to constrain/regularize the sparse noisy component. The standard geometric intuition is that since the boundary of the ℓ_1 -norm ball becomes sharp near the axes, this choice promotes sparse solutions. This property also holds for an ℓ_p -norm ball where p is sufficiently close to 1. Thus, it might be reasonable to replace the ℓ_1 -norm constraint on \mathbf{Y} with an $\ell_{1+\delta}$ -norm constraint for some small constant δ , which results in a strongly convex feasible set for the variable \mathbf{Y} (see [6]). Using Theorem 2, we will obtain an improved convergence rate of $O(1/t^2)$ instead of $O(1/t)$, practically without increasing the computational complexity per iteration (since \mathbf{Y} is updated via a conditional gradient update and linear optimization over a ℓ_p -norm ball can be carried-out in linear time [6]).

In order to demonstrate the plausibility of using the $\ell_{1+\delta}$ -norm instead of ℓ_1 -norm, in Table 2 we present results on synthetic data (similar to those done in [3]), which show that already for a moderate value of $\delta = 0.2$ we obtain quite satisfactory recovery results.

3.2.4 Replacing the ℓ_1 -norm regularizer with an elastic net regularizer

In certain cases it may be beneficial to replace the ℓ_1 -norm constraint (or regularizer) of the variable \mathbf{Y} in Problem (2) with an elastic net regularizer,

	$\delta = 0.05$	$\delta = 0.1$	$\delta = 0.2$	$\delta = 0.3$	$\delta = 0.4$
$\ \mathbf{X} - \mathbf{X}^*\ _F^2 / \ \mathbf{X}^*\ _F^2$	9.2×10^{-5}	6.0×10^{-4}	4.3×10^{-2}	0.25	0.61
$\ \mathbf{Y} - \mathbf{Y}^*\ _F^2 / \ \mathbf{Y}^*\ _F^2$	1.5×10^{-6}	6.4×10^{-6}	4.2×10^{-4}	2.5×10^{-3}	6.1×10^{-3}

Table 2: Empirical results for solving Problem (2) with $\ell_{1+\delta}$ -norm constraint on \mathbf{Y} instead of ℓ_1 -norm. The input matrix is $\mathbf{M} = \mathbf{X}^* + \mathbf{Y}^*$, where $\mathbf{X}^* = \mathbf{U}\mathbf{V}^\top$ for $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{n \times r}$ with entries sampled i.i.d. from $\mathcal{N}(0, 1/n)$ with $n = 100$ and $r = 10$. Every entry in \mathbf{Y}^* is set i.i.d. to 0 w.p. 0.9 and to either +1 or -1 w.p. 0.05. For simplicity we set exact bounds $\tau = \|\mathbf{X}^*\|_{\text{nuc}}$ and $s = \|\mathbf{Y}^*\|_{1+\delta}$. First row gives the relative recovery error of the low-rank signal and the second row gives the relative recovery error of the sparse noise component, averaged over 10 i.i.d. experiments.

i.e., to take $\mathcal{R}_{\mathbf{y}}(\mathbf{Y}) = \lambda_1 \|\mathbf{Y}\|_1 + \lambda_2 \|\mathbf{Y}\|_F^2$, for some $\lambda_1, \lambda_2 > 0$. The elastic net is a popular alternative to the standard ℓ_1 -norm regularizer for problems such as LASSO (see, for instance [18]). As opposed to the ℓ_1 -norm regularizer, the elastic net is strongly convex (though not differentiable). Thus, with such a choice for $\mathcal{R}_{\mathbf{y}}(\cdot)$, by invoking Theorem 3, Algorithm 1 guarantees a linear convergence rate. We note that when using the elastic net regularizer, the computation of \mathbf{W}_t on each iteration of Algorithm 1 requires to solve the optimization problem:

$$\arg \min_{\mathbf{W} \in \mathbf{E}} \left\{ \langle \mathbf{W}, \nabla_t \rangle + \lambda_1 \|\mathbf{W}\|_1 + \lambda_2 \|\mathbf{W}\|_F^2 \right\} = \arg \min_{\mathbf{W} \in \mathbf{E}} \left\{ \lambda_1 \|\mathbf{W}\|_1 + \lambda_2 \left\| \mathbf{W} + \frac{1}{2\lambda_2} \nabla_t \right\|_F^2 \right\},$$

where we again use the short notation $\nabla_t = \nabla g(\mathbf{X}_t + \mathbf{Y}_t)$. In the optimization problem above, the RHS admits a well-known closed-form solution given by the *shrinkage/soft-thresholding operator*, which can be computed in linear time (i.e., $O(mn)$ time), see for instance [2].

4 Rate of Convergence Analysis

In this section we provide the proofs for Theorems 1, 2, and 3. Throughout this section and for the simplicity of the yet to come developments we denote, for all $t \geq 1$, $\mathbf{Z}_t := \mathbf{X}_t + \mathbf{Y}_t$, $\mathbf{U}_t := \mathbf{V}_t + \mathbf{W}_t$, and $\mathbf{Q}_t := (\mathbf{X}_t, \mathbf{Y}_t)$. Note that, using these notations we obviously have that $\mathbf{Z}_{t+1} = (1 - \eta_t) \mathbf{Z}_t + \eta_t \mathbf{U}_t$. Similarly, for an optimal solution $\mathbf{Q}^* := (\mathbf{X}^*, \mathbf{Y}^*)$ of Problem (1) we denote $\mathbf{Z}^* := \mathbf{X}^* + \mathbf{Y}^*$.

We will need the following technical result which forms the basis for the proofs of all stated theorems.

Proposition 1 Let $\{(\mathbf{X}_t, \mathbf{Y}_t)\}_{t \geq 1}$ be a sequence generated by Algorithm 1. Then, for all $t \geq 1$, we have that

$$\begin{aligned} f(\mathbf{Q}_{t+1}) &\leq (1 - \eta_t) f(\mathbf{Q}_t) + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) + \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_t)) \\ &\quad + \eta_t \langle \mathbf{X}^* + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \eta_t^2 \beta \left(\|\mathbf{Z}_t - \mathbf{Z}^*\|^2 + \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \right). \end{aligned} \quad (6)$$

Proof Fix $t \geq 1$. Observe that by the update rule of Algorithm 1 (see step 6 of the algorithm), it holds that

$$\mathbf{X}_{t+1} + \mathbf{Y}_{t+1} = (1 - \eta_t)(\mathbf{X}_t + \mathbf{Y}_t) + \eta_t(\mathbf{V}_t + \mathbf{W}_t).$$

Thus, since $g(\cdot)$ is β -smooth, it follows that

$$\begin{aligned} g(\mathbf{X}_{t+1} + \mathbf{Y}_{t+1}) &\leq g(\mathbf{X}_t + \mathbf{Y}_t) + \eta_t \langle (\mathbf{V}_t + \mathbf{W}_t) - (\mathbf{X}_t + \mathbf{Y}_t), \nabla g(\mathbf{X}_t + \mathbf{Y}_t) \rangle \\ &\quad + \frac{\eta_t^2 \beta}{2} \|(\mathbf{X}_t + \mathbf{Y}_t) - (\mathbf{V}_t + \mathbf{W}_t)\|^2 \\ &= g(\mathbf{Z}_t) + \eta_t \langle \mathbf{V}_t + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \frac{\eta_t^2 \beta}{2} \|\mathbf{Z}_t - \mathbf{V}_t - \mathbf{W}_t\|^2. \end{aligned}$$

Using the above inequality we can write,

$$\begin{aligned} f(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}) &= g(\mathbf{X}_{t+1} + \mathbf{Y}_{t+1}) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}_{t+1}) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}_{t+1}) \\ &\leq g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}_{t+1}) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}_{t+1}) + \eta_t \langle \mathbf{V}_t + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle \\ &\quad + \frac{\eta_t^2 \beta}{2} \|\mathbf{Z}_t - \mathbf{V}_t - \mathbf{W}_t\|^2 \\ &\stackrel{(a)}{\leq} (1 - \eta_t) (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}_t) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}_t)) \\ &\quad + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{V}_t) + \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_t)) \\ &\quad + \eta_t \langle \mathbf{V}_t + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \frac{\eta_t^2 \beta}{2} \|\mathbf{X}_t + \mathbf{Y}_t - \mathbf{V}_t - \mathbf{W}_t\|^2 \\ &= (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{V}_t) + \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_t)) \\ &\quad + \eta_t \langle \mathbf{V}_t + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \frac{\eta_t^2 \beta}{2} \|\mathbf{Z}_t - \mathbf{V}_t - \mathbf{W}_t\|^2 \\ &\stackrel{(b)}{\leq} (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) + \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_t)) \\ &\quad + \eta_t \langle \mathbf{X}^* + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \frac{\eta_t^2 \beta}{2} \|\mathbf{Z}_t - \mathbf{X}^* - \mathbf{W}_t\|^2, \end{aligned} \quad (7)$$

where (a) follows from the convexity of $\mathcal{R}_{\mathcal{X}}(\cdot)$ and $\mathcal{R}_{\mathcal{Y}}(\cdot)$, while (b) follows from the choice of \mathbf{V}_t . Using the inequality $(a + b)^2 \leq 2a^2 + 2b^2$ which holds true for all $a, b \in \mathbb{R}$, we obtain

$$\begin{aligned} \|\mathbf{Z}_t - \mathbf{X}^* + \mathbf{W}_t\|^2 &= \|\mathbf{Z}_t - \mathbf{X}^* - \mathbf{Y}^* + (\mathbf{Y}^* - \mathbf{W}_t)\|^2 \\ &\leq 2\|\mathbf{Z}_t - \mathbf{Z}^*\|^2 + 2\|\mathbf{W}_t - \mathbf{Y}^*\|^2, \end{aligned} \quad (8)$$

where the last equality follows from the definition of \mathbf{Z}^* .

Finally, plugging the RHS of Eq. (8) into the RHS of Eq. (7) completes the proof of the proposition.

We now prove Theorem 1. For convenience, we first state the theorem in full details.

Theorem 4 *Assume that $\mathcal{R}_{\mathcal{Y}} := \iota_{\mathcal{Y}}$ where \mathcal{Y} is a nonempty, closed and convex subset of \mathbf{E} . Let $\{(\mathbf{X}_t, \mathbf{Y}_t)\}_{t \geq 1}$ be a sequence generated by Algorithm 1 with the following step-sizes:*

$$\eta_t = \begin{cases} \frac{\alpha}{2\beta}, & \text{if } t \leq t_0, \\ \frac{2}{t-t_0+\frac{4\beta}{\alpha}}, & \text{if } t > t_0, \end{cases} \quad (9)$$

where $t_0 := \max\{0, \lceil 2\beta/(\alpha) \ln(2C/(\alpha D_{\mathcal{Y}}^2)) \rceil\}$, for C satisfying $C \geq f(\mathbf{X}_1, \mathbf{Y}_1) - f^*$. Then, for all $t \geq t_0 + 1$ it holds that

$$f(\mathbf{X}_t, \mathbf{Y}_t) - f(\mathbf{X}^*, \mathbf{Y}^*) \leq \frac{4\beta D_{\mathcal{Y}}^2}{t - t_0 - 1 + \frac{4\beta}{\alpha}}.$$

Proof From the choice of \mathbf{W}_t we have that

$$\langle \mathbf{W}_t, \nabla g(\mathbf{Z}_t) \rangle + \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_t) \leq \langle \mathbf{Y}^*, \nabla g(\mathbf{Z}_t) \rangle + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*). \quad (10)$$

Now, using this in Proposition 1, we get for all $t \geq 1$, that

$$\begin{aligned} f(\mathbf{Q}_{t+1}) &\leq (1 - \eta_t) f(\mathbf{Q}_t) + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*)) \\ &\quad + \eta_t \langle \mathbf{X}^* + \mathbf{Y}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \eta_t^2 \beta \left(\|\mathbf{Z}_t - \mathbf{Z}^*\|^2 + \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \right) \\ &\leq (1 - \eta_t) f(\mathbf{Q}_t) + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*)) \\ &\quad + \eta_t \langle \mathbf{X}^* + \mathbf{Y}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \eta_t^2 \beta \left(\|\mathbf{Z}_t - \mathbf{Z}^*\|^2 + D_{\mathcal{Y}}^2 \right), \end{aligned} \quad (11)$$

where the last inequality follows from the fact that $\mathbf{W}_t, \mathbf{Y}^* \in \mathcal{Y}$. On the other hand, from the strong convexity of $g(\cdot)$ we obtain that

$$g(\mathbf{Z}_t) + \langle \mathbf{X}^* + \mathbf{Y}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle \leq g(\mathbf{Z}^*) - \frac{\alpha}{2} \|\mathbf{Z}_t - \mathbf{Z}^*\|^2.$$

Therefore, by combining these two inequalities we derive that

$$f(\mathbf{Q}_{t+1}) \leq (1 - \eta_t) f(\mathbf{Q}_t) + \eta_t f(\mathbf{Q}^*) - \eta_t \left(\frac{\alpha}{2} - \eta_t \beta \right) \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 + \eta_t^2 \beta D_{\mathcal{Y}}^2.$$

Subtracting $f(\mathbf{Q}^*)$ from both sides of the inequality and by denoting $h_t := f(\mathbf{Q}_t) - f(\mathbf{Q}^*)$, we obtain that $h_{t+1} \leq (1 - \eta_t) h_t + \eta_t^2 \beta D_{\mathcal{Y}}^2$ holds true for all $0 < \eta_t \leq \alpha/(2\beta)$. The result now follows from simple induction arguments and the choice of step-sizes detailed in the theorem (for details see Lemma 1 in the appendix below).

Before proving Theorem 2 we would like to comment about the constant C , which was used in the result above and appears in the step-size.

Remark 3 The constant C , even though appears in the step-size of the algorithm, can be easily bounded from above as we describe now. Suppose, we are setting the points \mathbf{W}_1 and \mathbf{V}_1 to be used in our algorithm as follows:

$$\mathbf{W}_1 = \arg \min_{\mathbf{W} \in \mathbf{E}} \{ \mathcal{R}_{\mathcal{Y}}(\mathbf{W}) + \langle \mathbf{W}, \nabla g(\mathbf{X}_1 + \mathbf{Y}_1) \rangle \},$$

and

$$\mathbf{V}_1 = \arg \min_{\mathbf{V} \in \mathbf{E}} \{ \mathcal{R}_{\mathcal{X}}(\mathbf{V}) + \langle \mathbf{V}, \nabla g(\mathbf{X}_1 + \mathbf{Y}_1) \rangle \}.$$

For these choices we obviously have (using optimality conditions) that

$$\begin{aligned} \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_1) + \mathcal{R}_{\mathcal{X}}(\mathbf{V}_1) + \langle \mathbf{W}_1 + \mathbf{V}_1, \nabla g(\mathbf{X}_1 + \mathbf{Y}_1) \rangle &\leq \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) \\ &\quad + \langle \mathbf{Y}^* + \mathbf{X}^*, \nabla g(\mathbf{X}_1 + \mathbf{Y}_1) \rangle. \end{aligned}$$

Hence, using the gradient inequality on the function g , yields that

$$\begin{aligned} f(\mathbf{X}_1, \mathbf{Y}_1) - f^* &= g(\mathbf{X}_1, \mathbf{Y}_1) - g(\mathbf{X}^*, \mathbf{Y}^*) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}_1) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}_1) - \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*) - \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) \\ &\leq \langle (\mathbf{X}_1 + \mathbf{Y}_1) - (\mathbf{Y}^* + \mathbf{X}^*), \nabla g(\mathbf{X}_1 + \mathbf{Y}_1) \rangle + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}_1) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}_1) \\ &\quad - \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_1) - \mathcal{R}_{\mathcal{X}}(\mathbf{V}_1) + \langle (\mathbf{Y}^* + \mathbf{X}^*) - (\mathbf{W}_1 + \mathbf{V}_1), \nabla g(\mathbf{X}_1 + \mathbf{Y}_1) \rangle \\ &= \langle (\mathbf{X}_1 + \mathbf{Y}_1) - (\mathbf{W}_1 + \mathbf{V}_1), \nabla g(\mathbf{X}_1 + \mathbf{Y}_1) \rangle + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}_1) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}_1) \\ &\quad - \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_1) - \mathcal{R}_{\mathcal{X}}(\mathbf{V}_1). \end{aligned}$$

The obtained bound does not depend on the optimal solution and therefore can be computed explicitly.

It should be noted that in the case of Robust PCA (e.g., Problem (2)), we have that $\mathcal{R}_{\mathcal{X}}(\mathbf{X}) = \iota_{\|\cdot\|_{\text{nuc}} \leq \tau}(\mathbf{X})$ and $\mathcal{R}_{\mathcal{Y}}(\mathbf{Y}) = \iota_{\|\cdot\|_1 \leq s}(\mathbf{Y})$. In this case, computing the matrices \mathbf{W}_1 and \mathbf{V}_1 is computationally very efficient, since it requires to compute a single leading singular vectors pair, and solving a single linear minimization problem over an ℓ_1 -ball, respectively.

Now, we turn to prove Theorem 2. Again, we first state the theorem in full details.

Theorem 5 *Assume that $\mathcal{R}_{\mathcal{X}} := \iota_{\mathcal{X}}$ where \mathcal{X} is a nonempty, closed and convex subset of \mathbf{E} and $\mathcal{R}_{\mathcal{Y}} := \iota_{\mathcal{Y}}$, where \mathcal{Y} is an γ -strongly convex and closed subset of \mathbf{E} . Let $\{(\mathbf{X}_t, \mathbf{Y}_t)\}_{t \geq 1}$ be a sequence generated by Algorithm 1 using the step-size $\eta_t = (t - 1 + 6\beta/\alpha)^{-1}$ for all $t \geq 1$. Then, for all $t \geq 1$ it holds that*

$$f(\mathbf{X}_t, \mathbf{Y}_t) - f^* \leq \frac{9 \max \left\{ \frac{128\beta^2}{\alpha\gamma^2}, \frac{4\beta^2}{\alpha^2} (f(\mathbf{X}_1, \mathbf{Y}_1) - f^*) \right\}}{\left(t - 1 + 6\frac{\beta}{\alpha} \right)^2}.$$

Moreover, if there exists $G > 0$ such that $\min_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \|\nabla g(\mathbf{X} + \mathbf{Y})\|_* \geq G$, then using a fixed step-size $\eta_t = \min\{\alpha/(2\beta), \gamma G/(8\beta)\}$ for all $t \geq 1$, guarantees that

$$f(\mathbf{X}_t, \mathbf{Y}_t) - f^* \leq (f(\mathbf{X}_1, \mathbf{Y}_1) - f^*) \cdot \exp \left(- \min \left\{ \frac{\alpha}{2\beta}, \frac{\gamma G}{8\beta} \right\} (t - 1) \right).$$

Proof Fix some iteration $t \geq 1$ and define the point $\tilde{\mathbf{W}}_t := \frac{1}{2}(\mathbf{W}_t + \mathbf{Y}^*) - (\gamma/8)\|\mathbf{W}_t - \mathbf{Y}^*\|^2 \mathbf{P}_t$ where $\mathbf{P}_t \in \arg \min_{\mathbf{P} \in \mathbf{E}, \|\mathbf{P}\| \leq 1} \langle \mathbf{P}, \nabla g(\mathbf{Z}_t) \rangle$. Note that since \mathcal{Y} is an γ -strongly convex set, it follows from Definition 3 that $\tilde{\mathbf{W}}_t \in \mathcal{Y}$. Moreover, from the optimal choice of \mathbf{W}_t we have that $\langle \mathbf{W}_t, \nabla g(\mathbf{Z}_t) \rangle \leq \langle \tilde{\mathbf{W}}_t, \nabla g(\mathbf{Z}_t) \rangle$. Thus, we have that

$$\begin{aligned} \langle \mathbf{X}^* + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle &\leq \langle \mathbf{X}^* + \tilde{\mathbf{W}}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle \\ &= \left\langle \mathbf{X}^* - \mathbf{X}_t + \frac{\mathbf{W}_t + \mathbf{Y}^*}{2} - \frac{\gamma}{8} \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \mathbf{P}_t - \mathbf{Y}_t, \nabla g(\mathbf{Z}_t) \right\rangle \\ &\stackrel{(a)}{\leq} \left\langle \mathbf{X}^* - \mathbf{X}_t + \frac{\mathbf{Y}^* + \mathbf{Y}^*}{2} - \frac{\gamma}{8} \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \mathbf{P}_t - \mathbf{Y}_t, \nabla g(\mathbf{Z}_t) \right\rangle \\ &\stackrel{(b)}{=} \langle \mathbf{Z}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle - \frac{\gamma}{8} \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \cdot \|\nabla g(\mathbf{Z}_t)\|_*, \end{aligned} \quad (12)$$

where (a) follows from the fact that $\langle \mathbf{W}_t - \mathbf{Y}^*, \nabla g(\mathbf{Z}_t) \rangle \leq 0$, and (b) follows from the definition of \mathbf{P}_t and Holder's inequality.

Plugging Eq. (12) into Eq. (6), and recalling that $\mathbf{W}_t \in \mathcal{Y}$ (hence, $\mathcal{R}_{\mathcal{Y}}(\mathbf{W}_t) = \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*) = \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) = 0$), we have that

$$\begin{aligned} f(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}) &\leq (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t g(\mathbf{Z}_t) + \eta_t \langle \mathbf{X}^* + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle \\ &\quad + \eta_t^2 \beta \left(\|\mathbf{Z}_t - \mathbf{Z}^*\|^2 + \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \right) \\ &\leq (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t g(\mathbf{Z}_t) + \eta_t \langle \mathbf{Z}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle \\ &\quad + \eta_t^2 \beta \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\gamma}{8} \|\nabla g(\mathbf{Z}_t)\|_* - \eta_t \beta \right) \\ &\stackrel{(a)}{\leq} (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t g(\mathbf{Z}^*) - \eta_t \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 \left(\frac{\alpha}{2} - \eta_t \beta \right) \\ &\quad - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\gamma}{8} \|\nabla g(\mathbf{Z}_t)\|_* - \eta_t \beta \right) \\ &= (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t f(\mathbf{X}^*, \mathbf{Y}^*) - \eta_t \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 \left(\frac{\alpha}{2} - \eta_t \beta \right) \\ &\quad - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\gamma}{8} \|\nabla g(\mathbf{Z}_t)\|_* - \eta_t \beta \right), \end{aligned} \quad (13)$$

where (a) follows from the strong convexity of $g(\cdot)$ since we have that

$$\langle \mathbf{Z}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle \leq g(\mathbf{Z}^*) - g(\mathbf{Z}_t) - \frac{\alpha}{2} \|\mathbf{Z}_t - \mathbf{Z}^*\|^2.$$

Using again the strong convexity of $g(\cdot)$, we have from Eq. (3) that

$$\begin{aligned} \|\nabla g(\mathbf{Z}_t)\|_* &\geq \sqrt{\frac{\alpha}{2}} \cdot \sqrt{g(\mathbf{Z}_t) - g(\mathbf{Z}^*)} \stackrel{(a)}{=} \sqrt{\frac{\alpha}{2}} \cdot \sqrt{f(\mathbf{X}_t, \mathbf{Y}_t) - f(\mathbf{X}^*, \mathbf{Y}^*)} \\ &= \sqrt{\frac{\alpha}{2}} \cdot \sqrt{h_t}, \end{aligned}$$

where $h_t := f(\mathbf{X}_t, \mathbf{Y}_t) - f(\mathbf{X}^*, \mathbf{Y}^*)$, and (a) follows since $\mathcal{R}_{\mathcal{X}}(\mathbf{X}_t) = \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}_t) = \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) = \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*) = 0$. Therefore, by subtracting $f(\mathbf{X}^*, \mathbf{Y}^*)$ from both sides of (13), we get that

$$h_{t+1} \leq (1 - \eta_t) h_t - \eta_t \left(\frac{\alpha}{2} - \eta_t \beta \right) \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\gamma}{8} \sqrt{\frac{\alpha}{2}} \cdot \sqrt{h_t} - \eta_t \beta \right).$$

Thus, we obtain the recursion: $h_{t+1} \leq (1 - \eta_t) h_t$ for all $\eta_t \leq \min \{ \gamma \sqrt{\alpha} \sqrt{h_t} / (8\sqrt{2}\beta), \alpha / (2\beta) \}$.

In particular, setting η_t as stated in the theorem yields the stated convergence rate via a simple induction argument, given in Lemma 2 (see appendix for a proof).

In order to prove the second part of the theorem, i.e., a linear rate of convergence in the case that the gradients are bounded from below, we observe that plugging the bound G on the magnitude of the gradients into the RHS of Eq. (13), directly gives

$$f(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}) \leq (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t f(\mathbf{X}^*, \mathbf{Y}^*) - \eta_t \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 \left(\frac{\alpha}{2} - \eta_t \beta \right) - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\gamma G}{8} - \eta_t \beta \right).$$

Thus, for any $\eta_t \leq \min \{ \alpha / (2\beta), \gamma G / (8\beta) \}$, by subtracting $f(\mathbf{X}^*, \mathbf{Y}^*)$ from both sides, we obtain $h_{t+1} \leq (1 - \eta_t) h_t$. In particular, setting $\eta_t = \min \{ \alpha / (2\beta), \gamma G / (8\beta) \}$ for all $t \geq 1$ and using elementary manipulations, gives the linear rate stated in the theorem.

Before stating in details Theorem 3 and proving it, we would like to prove that the additional assumption made in this result, i.e., that $\mathcal{R}_{\mathcal{X}}(\cdot)$ or $\mathcal{R}_{\mathcal{Y}}(\cdot)$ is δ -strongly convex, actually guarantees that the whole objective function $f(\cdot, \cdot)$ is also strongly convex. The following result is valid when the function g is strongly convex with respect to the Euclidean norm.

Proposition 2 *Assume that $\mathcal{R}_{\mathcal{X}}$ or $\mathcal{R}_{\mathcal{Y}}$ is δ -strongly convex with respect to the Euclidean norm. Then, the objective function $f(\cdot, \cdot)$ of Problem (1), is τ -strongly convex with respect to the Euclidean norm, where $\tau = (\delta + 2\alpha - \sqrt{\delta^2 + 4\alpha^2}) / 2$.*

Proof Throughout the proof we let $\|\cdot\|$ denote the Euclidean norm over \mathbf{E} . Without the loss of generality we assume that $\mathcal{R}_{\mathcal{X}}$ is δ -strongly convex. Let $\mathbf{Q}_1 = (\mathbf{X}_1, \mathbf{Y}_1)$ and $\mathbf{Q}_2 = (\mathbf{X}_2, \mathbf{Y}_2)$ be two points in $\mathbf{E} \times \mathbf{E}$ and $\lambda \in [0, 1]$. Then, by the definition of strong convexity, we have that

$$\mathcal{R}_{\mathcal{X}}(\lambda \mathbf{X}_1 + (1 - \lambda) \mathbf{X}_2) \leq \lambda \mathcal{R}_{\mathcal{X}}(\mathbf{X}_1) + (1 - \lambda) \mathcal{R}_{\mathcal{X}}(\mathbf{X}_2) - \frac{\lambda(1 - \lambda)\delta}{2} \|\mathbf{X}_1 - \mathbf{X}_2\|^2,$$

and

$$g(\mathbf{P}) \leq \lambda g(\mathbf{X}_1 + \mathbf{Y}_1) + (1 - \lambda) g(\mathbf{X}_2 + \mathbf{Y}_2) - \frac{\lambda(1 - \lambda)\alpha}{2} \|\mathbf{X}_1 + \mathbf{Y}_1 - \mathbf{X}_2 - \mathbf{Y}_2\|^2,$$

where $\mathbf{P} = \lambda(\mathbf{X}_1 + \mathbf{Y}_1) + (1 - \lambda)(\mathbf{X}_2 + \mathbf{Y}_2)$. On the other hand, for any $s > 0$, we have that

$$\|\mathbf{X}_1 + \mathbf{Y}_1 - \mathbf{X}_2 - \mathbf{Y}_2\|^2 \geq (1 - s)\|\mathbf{X}_1 - \mathbf{X}_2\|^2 + \frac{s-1}{s}\|\mathbf{Y}_1 - \mathbf{Y}_2\|^2,$$

where we have used the fact that $(a + b)^2 \geq (1 - s)a^2 + (s - 1)b^2/s$ for all $a, b \in \mathbb{R}$, and that the norm is the Euclidean norm. Combining all these facts and using the fact that $\mathcal{R}_Y(\cdot)$ is convex yields that

$$\begin{aligned} f(\lambda\mathbf{Q}_1 + (1 - \lambda)\mathbf{Q}_2) &= g(\mathbf{P}) + \mathcal{R}_X(\lambda\mathbf{X}_1 + (1 - \lambda)\mathbf{X}_2) + \mathcal{R}_Y(\lambda\mathbf{Y}_1 + (1 - \lambda)\mathbf{Y}_2) \\ &\leq \lambda f(\mathbf{Q}_1) + (1 - \lambda)f(\mathbf{Q}_2) - \frac{\lambda(1 - \lambda)\delta}{2}\|\mathbf{X}_1 - \mathbf{X}_2\|^2 \\ &\quad - \frac{\lambda(1 - \lambda)\alpha(1 - s)}{2}\|\mathbf{X}_1 - \mathbf{X}_2\|^2 - \frac{\lambda(1 - \lambda)\alpha(s - 1)}{2s}\|\mathbf{Y}_1 - \mathbf{Y}_2\|^2 \\ &= \lambda f(\mathbf{Q}_1) + (1 - \lambda)f(\mathbf{Q}_2) - \frac{\lambda(1 - \lambda)}{2}(\delta + \alpha(1 - s))\|\mathbf{X}_1 - \mathbf{X}_2\|^2 \\ &\quad - \frac{\lambda(1 - \lambda)\alpha(s - 1)}{2s}\|\mathbf{Y}_1 - \mathbf{Y}_2\|^2 \\ &\leq \lambda f(\mathbf{Q}_1) + (1 - \lambda)f(\mathbf{Q}_2) - \frac{\lambda(1 - \lambda)\tau(s)}{2}\|\mathbf{Q}_1 - \mathbf{Q}_2\|^2, \end{aligned}$$

where $\tau(s) = \min\{\delta + \alpha(1 - s), \alpha(s - 1)/s\}$ and the last inequality follows from the definitions of \mathbf{Q}_1 and \mathbf{Q}_2 . It is easy to check that $\tau(s)$ gets its maximum with respect to s when $s = (\delta + \sqrt{\delta^2 + 4\alpha^2})/(2\alpha)$. Therefore we get that $f(\cdot, \cdot)$ is strongly convex with parameter $(\delta + 2\alpha - \sqrt{\delta^2 + 4\alpha^2})/2$.

Thanks to Proposition 2, Problem (1) becomes an unconstrained minimization of a strongly convex function. Therefore, we can expect to achieve a linear rate of convergence of Algorithm 1 as we prove below. We now state first Theorem 3 in full details and then prove it.

Theorem 6 *Assume that $\mathcal{R}_Y(\cdot)$ is δ -strongly convex. Let $\{(\mathbf{X}_t, \mathbf{Y}_t)\}_{t \geq 1}$ be a sequence generated by Algorithm 1 using the fixed step-size $\eta_t = \min\{\alpha, \delta\}/(2\beta)$ for all $t \geq 1$. Then, for all $t \geq 1$, we have that*

$$f(\mathbf{X}_t, \mathbf{Y}_t) - f^* \leq (f(\mathbf{X}_1, \mathbf{Y}_1) - f^*) \cdot \exp\left(-\frac{\min\{\alpha, \delta\}}{2\beta}(t - 1)\right).$$

Proof Fix some iteration $t \geq 1$. Using the optimal choice of \mathbf{W}_t , from the strong convexity of $\mathbf{W} \rightarrow \langle \mathbf{W}, \nabla g(\mathbf{Z}_t) \rangle + \mathcal{R}_Y(W)$, we have that

$$\langle \mathbf{W}_t, \nabla g(\mathbf{Z}_t) \rangle + \mathcal{R}_Y(\mathbf{W}_t) \leq \langle \mathbf{Y}^*, \nabla g(\mathbf{Z}_t) \rangle + \mathcal{R}_Y(\mathbf{Y}^*) - \frac{\delta}{2}\|\mathbf{W}_t - \mathbf{Y}^*\|^2.$$

Plugging the above inequality into Proposition 1, we have that

$$\begin{aligned}
f(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}) &\leq (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) + \mathcal{R}_{\mathcal{Y}}(\mathbf{W}_t)) \\
&\quad + \eta_t \langle \mathbf{X}^* + \mathbf{W}_t - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \eta_t^2 \beta \left(\|\mathbf{Z}_t - \mathbf{Z}^*\|^2 + \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \right) \\
&\leq (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t (g(\mathbf{Z}_t) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*)) \\
&\quad + \eta_t \langle \mathbf{Z}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle + \eta_t^2 \beta \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\delta}{2} - \eta_t \beta \right) \\
&\stackrel{(a)}{\leq} (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t (g(\mathbf{Z}^*) + \mathcal{R}_{\mathcal{X}}(\mathbf{X}^*) + \mathcal{R}_{\mathcal{Y}}(\mathbf{Y}^*)) \\
&\quad - \eta_t \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 \left(\frac{\alpha}{2} - \eta_t \beta \right) - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\delta}{2} - \eta_t \beta \right) \\
&= (1 - \eta_t) f(\mathbf{X}_t, \mathbf{Y}_t) + \eta_t f(\mathbf{X}^*, \mathbf{Y}^*) \\
&\quad - \eta_t \|\mathbf{Z}_t - \mathbf{Z}^*\|^2 \left(\frac{\alpha}{2} - \eta_t \beta \right) - \eta_t \|\mathbf{W}_t - \mathbf{Y}^*\|^2 \left(\frac{\delta}{2} - \eta_t \beta \right),
\end{aligned}$$

where (a) follows from the strong convexity of $g(\cdot)$, which implies that

$$\langle \mathbf{Z}^* - \mathbf{Z}_t, \nabla g(\mathbf{Z}_t) \rangle \leq g(\mathbf{Z}^*) - g(\mathbf{Z}_t) - \frac{\alpha}{2} \|\mathbf{Z}_t - \mathbf{Z}^*\|^2.$$

Thus, subtracting $f(\mathbf{X}^*, \mathbf{Y}^*)$ from both sides and using the notation $h_t := f(\mathbf{X}_t, \mathbf{Y}_t) - f(\mathbf{X}^*, \mathbf{Y}^*)$, we conclude that

$$\forall 0 < \eta_t \leq \frac{\min\{\alpha, \delta\}}{2\beta} : \quad h_{t+1} \leq (1 - \eta_t) h_t.$$

The theorem now follows from choosing $\eta_t = \min\{\alpha, \delta\} / (2\beta)$ and using standard manipulations.

5 Numerical Results

In this section we present evidence for the empirical performance of Algorithm 1 on the Robust PCA problem in the constrained formulation given in Problem (2). Focusing on first-order methods that are scalable to high-dimensional problems involving optimization problems with a nuclear-norm constraint, we compare our method to other competing projection-free first-order methods that avoid using high-rank SVD computations.

We tested Algorithm 1 and compare to two methods: the standard conditional gradient method and the conditional gradient variant proposed in [11], which adds an additional proximal step to update the sparse component \mathbf{Y} , on top of the standard CG method. See Table 3 for description of the algorithms.

For all algorithms, we apply a line-search procedure to compute the optimal convex combination taken in the conditional gradient-like step on each iteration. This implementation is straightforward for the standard conditional gradient method (CGCG) and the proposed variant of [11] (CGCG-P). For

abbrv.	description
ALT-CGPG	Algorithm 1, sparse component \mathbf{Y} updated via CG, low-rank component \mathbf{X} updated via low-rank PG (see Section 3.2.1 for details)
ALT-PGCG	Algorithm 1, sparse component \mathbf{Y} updated via PG, low-rank component \mathbf{X} updated via CG
CGCG	Both blocks \mathbf{X} and \mathbf{Y} updated via CG
CGCG-P	Algorithm FW-P from [11] - both \mathbf{X}, \mathbf{Y} updated via CG, followed by an additional PG update to sparse component \mathbf{Y} only

Table 3: Description of the tested algorithms. CG is a short notation for a conditional gradient update, and PG is a short notation for a proximal gradient update.

our proposed methods (ALT-CGPG, ALT-PGCG), we set the step-size for the proximal gradient update on each iteration t to $\eta_t = \frac{2}{t+1}$ (this seems as a very good and practical approximation to the choice in Theorem 4 once we neglect the first short phase with a constant step-size, and start immediately with the second regime of step-sizes). Once we have computed the proximal gradient update with this step-size (Line 5 of Algorithm 1), we use a line-search to set the optimal convex combination parameter (used in Line 6 of Algorithm 1). It is straightforward to argue that performing such a line-search instead of using the pre-fixed value of η_t used for the proximal gradient update, does not hurt any of the guarantees specified in Theorems 1, 2 and 3, but can be quite important from a practical point of view.

5.1 Experiments

We generate a synthetic data as follows. We set the dimensions in all experiments to $m = n = 1000$. We generate the low-rank component by taking $\mathbf{L} = 10\mathbf{U}\mathbf{V}^\top$, where \mathbf{U} is an $m \times r$ matrix and \mathbf{V} is an $n \times r$ matrix, where r varies, and the entries of \mathbf{U} and \mathbf{V} are i.i.d. standard Gaussian random variables. The sparse component \mathbf{S} is generated by setting $\mathbf{S} = 10\mathbf{N}$, where \mathbf{N} is a matrix with i.i.d. standard Gaussian entries, and each entry in \mathbf{S} is set to zero with probability $1 - p$ (independently of all other entries), where p varies. We then set the observed matrix to $\mathbf{M} = \mathbf{L} + \mathbf{S}$. For each value of (r, p) we have generated 15 i.i.d. experiments. See Table 4 for a quick summary.

All algorithms were used with the exact parameters $\tau = \|\mathbf{L}\|_*$ and $s = \|\mathbf{S}\|_1$, and with the same initialization point $(\mathbf{X}_1, \mathbf{Y}_1) = (\mathbf{0}, \mathbf{0})$. The algorithms were implemented in Matlab with the `svds` command used to compute the low-rank SVD updates. For our algorithm ALT-CGPG we have used a rank- r SVD to compute the low-rank proximal gradient update (were r is the rank of the low-rank data component \mathbf{L}). For all experiments we measured the function value (as given in Problem (2)) both as a function of the number of iterations executed and the overall runtime.

From the results in Figure 1 and Figure 2 it is clear that in each one of the six scenarios tested (different values of r and p), at least one of the

config.	fig.	rank of \mathbf{L} (r)	sampling freq. in \mathbf{S} (p)	avg. $\ \mathbf{L}\ _*$	avg. $\ \mathbf{S}\ _1$
1	1	5	0.001	4.9926e+04	7.9669e+04
2	1	5	0.003	4.9682e+04	2.3806e+05
3	1	25	0.001	2.4872e+05	7.9837e+04
4	1	25	0.003	2.4853e+05	2.3675e+05
5	2	25	0.03	2.4830e+05	2.3914e+06
6	2	130	0.01	1.2589e+06	7.9836e+05

Table 4: Description of data used. The forth and fifth columns give the values $\|\mathbf{L}\|_*$ and $\|\mathbf{S}\|_1$ averaged over the 15 i.i.d. experiments.

variants of Algorithm 1 (either ALT-CGPG or ALT-PGCG) clearly outperforms CGCG and CGCG-P. In particular, when examining the results and the corresponding norm bounds τ and s (as recorded in Table 4), we can see evidence for the improved complexity achieved in Theorem 1, which presents improved convergence bound in low/high Signal-to-Noise regimes (see also discussion in Section 3.2.2). We can see that, as a rule of thumb, indeed when $\|\mathbf{S}\|_1 \gg \|\mathbf{L}\|_*$, updating the sparse component \mathbf{Y} via a proximal gradient update in Algorithm 1, results in significantly faster performances than when a conditional gradient update is used. Similarly, when $\|\mathbf{S}\|_1 \gg \|\mathbf{L}\|_*$, we can see that updating the low-rank component \mathbf{X} via a (low-rank) proximal-update in Algorithm 1, results in significantly faster performances of our algorithm. Perhaps surprisingly, it also seems that the standard conditional gradient method (CGCG) outperforms the variant recently proposed in [11] (CGCG-P), with configuration 5 ($r = 25$, $p = 0.03$) being the exception.

Acknowledgments Dan Garber is supported by the ISRAEL SCIENCE FOUNDATION (grant No. 1108/18).

6 Appendix

Lemma 1 Consider a sequence $\{h_t\}_{t \geq 1} \subset \mathbb{R}_+$ satisfying the recursion:

$$\forall t \geq 1 \forall \eta_t \in \left(0, \frac{\alpha}{2\beta}\right] : \quad h_{t+1} \leq (1 - \eta_t) h_t + \eta_t^2 \beta D_y^2.$$

Then, setting the step-size η_t according to:

$$\eta_t = \begin{cases} \frac{\alpha}{2\beta} & \text{if } t \leq t_0, \\ \frac{2}{t - t_0 + \frac{4\beta}{\alpha}} & \text{if } t > t_0, \end{cases}$$

where $t_0 := \max\{0, \lceil 2\beta/(\alpha) \ln(2C/(\alpha D_y^2)) \rceil\}$, for C satisfying $C \geq h_1$, guarantees, for all $t \geq t_0 + 1$ that

$$h_t \leq \frac{4\beta D_y^2}{t - t_0 - 1 + \frac{4\beta}{\alpha}}.$$









(a) $r = 5$ and $p = 0.001$ (b) $r = 5$ and $p = 0.001$ (c) $r = 5$ and $p = 0.003$ (d) $r = 5$ and $p = 0.003$ (e) $r = 25$ and $p = 0.001$ (f) $r = 25$ and $p = 0.001$ (g) $r = 25$ and $p = 0.003$ (h) $r = 25$ and $p = 0.003$

Fig. 1: Results for configurations 1-4. Each configuration presents the average over 15 i.i.d. runs.

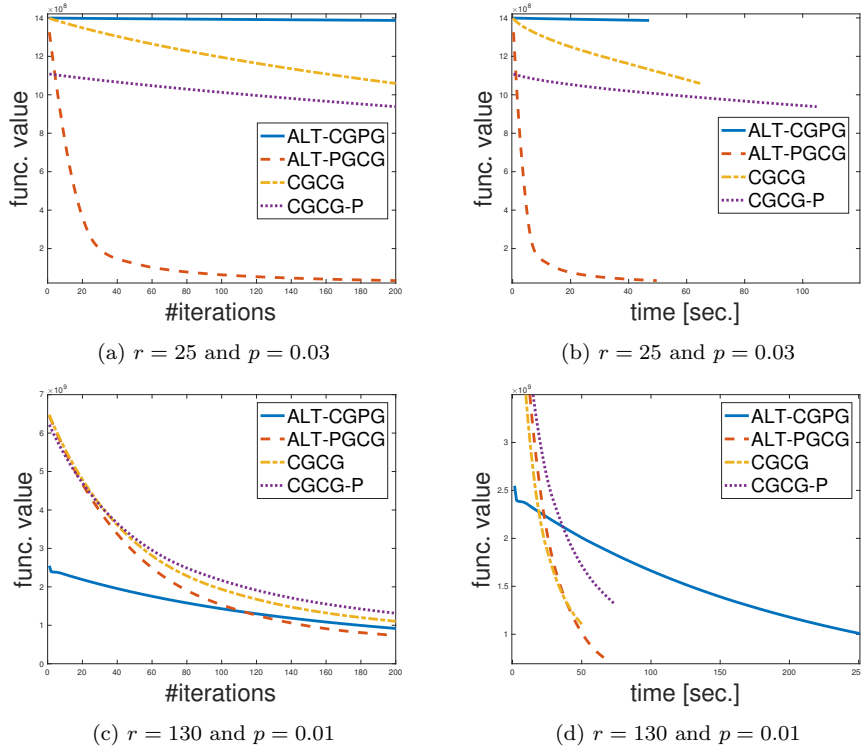


Fig. 2: Results for configurations 5-6. Each configuration presents the average over 15 i.i.d. runs.

Proof Let us define $v_t := h_t / (2\beta D_y^2)$ for all $t \geq 1$. Dividing both sides of the recursion in the lemma by $2\beta D_y^2$, we obtain the recursion

$$\forall 0 < \eta_t \leq \frac{\alpha}{2\beta} : \quad v_{t+1} \leq (1 - \eta_t) v_t + \frac{\eta_t^2}{2}. \quad (14)$$

Let C, t_0 and $\{\eta_t\}_{t \geq 1}$ be as defined in the lemma, and recall that $\eta_t = \alpha / (2\beta)$, for all $t \leq t_0$. Using Eq. (14), we have that

$$\begin{aligned} v_{t_0+1} &\leq (1 - \eta_0)^{t_0} v_1 + \frac{\eta_0^2}{2} \sum_{t=1}^{t_0} (1 - \eta_0)^{t-1} \leq \frac{C}{2\beta D_y^2} \cdot e^{-\eta_0 t_0} + \frac{\eta_0^2}{2} \cdot \frac{1}{\eta_0} \\ &= \frac{C}{2\beta D_y^2} \cdot e^{-\frac{\alpha t_0}{2\beta}} + \frac{\alpha}{4\beta}. \end{aligned}$$

Thus, for $t_0 := \max\{0, \lceil 2\beta / (\alpha \ln(2C / (\alpha D_y^2))) \rceil\}$, we obtain that $v_{t_0+1} \leq \alpha / (2\beta)$.

We now show that for all $t \geq t_0 + 1$, it holds that $v_t \leq 2 / (t - t_0 + 1 + c_0)$ for $c_0 := 4\beta / \alpha - 2$. For the base case $t = t_0 + 1$, we indeed have already

showed that $v_{t_0+1} \leq \alpha/(2\beta)$, as needed. Note that using the step-size choice $\eta_t := 2/(t - t_0 + 2 + c_0)$ for all $t \geq t_0 + 1$, as defined in the lemma, we have that $\eta_t \leq \eta_{t_0+1} = 2/(3 + c_0) < 2/c_0 = \alpha/(2\beta)$, and hence we can apply the recursion (14) for all $t \geq t_0 + 1$. Thus, assuming the induction holds for some $t \geq t_0 + 1$, using the recursion (14), the induction hypothesis, and our step-size choice, we have that

$$\begin{aligned}
v_{t+1} &\leq v_t(1 - \eta_t) + \frac{\eta_t^2}{2} \leq \frac{2}{t - t_0 + 1 + c_0} \left(1 - \frac{2}{t - t_0 + 2 + c_0}\right) \\
&\quad + \frac{2}{(t - t_0 + 2 + c_0)^2} \\
&= \frac{2}{t - t_0 + 2 + c_0} \left(1 + \frac{1}{t - t_0 + 1 + c_0}\right) \left(1 - \frac{2}{t - t_0 + 2 + c_0}\right) \\
&\quad + \frac{2}{(t - t_0 + 2 + c_0)^2} \\
&= \frac{2}{t - t_0 + 2 + c_0} \left(1 + \frac{(t - t_0 + 2 + c_0) - 2(t - t_0 + 1 + c_0) - 2 + (t - t_0 + 1 + c_0)}{(t - t_0 + 1 + c_0)(t - t_0 + 2 + c_0)}\right) \\
&= \frac{2}{t - t_0 + 2 + c_0} \left(1 - \frac{1}{(t - t_0 + 1 + c_0)(t - t_0 + 2 + c_0)}\right) \\
&< \frac{2}{t - t_0 + 2 + c_0}.
\end{aligned}$$

Hence, the induction implies that $v_t \leq 2/(t - t_0 - 1 + 4\beta/\alpha)$ for all $t \geq t_0 + 1$. The proof is completed by recalling that $h_t = 2\beta D_y^2 v_t$.

Lemma 2 Consider a sequence $\{h_t\}_{t \geq 1} \subset \mathbb{R}_+$ satisfying the recursion:

$$\forall t \geq 1: \quad h_{t+1} \leq h_t - \eta_t h_t,$$

where $0 < \eta_t \leq \min\{c_1\sqrt{h_t}, c_2\}$, $c_1 > 0$ and $0 < c_2 \leq 1$. Then, setting $\eta_t = 3/(t - 1 + 3c_2^{-1})$ for all $t \geq 1$, yields that $h_t \leq 9 \max\{c_1^{-2}, c_2^{-2}h_1\}/(t - 1 + 3c_2^{-1})^2$.

Proof We prove via induction that for suitably chosen positive constants a and b that $h_t \leq a/(t + b)^2$ for all $t \geq 1$.

Fix some iteration $t \geq 1$ and suppose the claim holds for h_t . We consider now two cases. First, if $h_t \leq a/(t + b + 1)^2$, then, since by the recursion in the lemma it holds that $h_{t+1} \leq h_t$, the claim clearly holds in this case for h_{t+1} .

For the second case, we assume $h_t \geq a/(t+b+1)^2$. Using this assumption, the recursion, and the induction hypothesis, we have that

$$\begin{aligned} h_{t+1} &\leq h_t - \eta_t h_t \\ &\leq \frac{a}{(t+b)^2} - \eta_t \frac{a}{(t+b+1)^2} \\ &= \frac{a}{(t+b+1)^2} \left(\left(\frac{t+b+1}{t+b} \right)^2 - \eta_t \right) \\ &\leq \frac{a}{(t+b+1)^2} \left(1 + \frac{3}{t+b} - \eta_t \right). \end{aligned}$$

Thus, for $\eta_t = 3/(t+b)$, the induction clearly holds.

Since, for the recursion to hold, it is also required that $\eta_t \leq \min\{c_1\sqrt{h_t}, c_2\} \leq \min\{c_1\sqrt{a}/(t+b), c_2\}$, this brings us to the following conditions on a and b which should be valid for all $t \geq 1$:

$$\frac{3}{t+b} \leq c_2 \quad \text{and} \quad \frac{3}{t+b} \leq \frac{c_1\sqrt{a}}{t+b}.$$

Thus, we get the requirements $b \geq 3c_2^{-1} - 1$ and $a \geq 9c_1^{-2}$.

Finally, since for the base case $t = 1$ it needs to hold that $h_1 \leq a/(b+1)^2$, we can choose $b = 3c_2^{-1} - 1$ and $a = 9 \max\{c_1^{-2}, c_2^{-2}h_1\}$, which guarantee that.

References

1. Zeyuan Allen-Zhu, Elad Hazan, Wei Hu, and Yuanzhi Li. Linear convergence of a frank-wolfe type algorithm over trace-norm balls. In *Advances in Neural Information Processing Systems*, pages 6192–6201, 2017.
2. Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
3. Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.
4. Nan Du, Yichen Wang, Niao He, Jimeng Sun, and Le Song. Time-sensitive recommendation from recurrent user activities. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3492–3500. Curran Associates, Inc., 2015.
5. Dan Garber. Faster projection-free convex optimization over the spectrahedron. In *Advances in Neural Information Processing Systems*, pages 874–882, 2016.
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. Gauthier Gidel, Fabian Pedregosa, and Simon Lacoste-Julien. Frank-wolfe splitting via augmented lagrangian method. In *AISTATS*, 2018.
9. Xiangru Huang, Ian En-Hsu Yen, Ruohan Zhang, Qixing Huang, Pradeep Ravikumar, and Inderjit Dhillon. Greedy Direction Method of Multiplier for MAP Inference of Large Output Domain. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1550–1559, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR.

10. Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning, ICML*, 2013.
11. Cun Mu, Yuqian Zhang, John Wright, and Donald Goldfarb. Scalable robust matrix recovery: Frank–wolfe meets proximal methods. *SIAM Journal on Scientific Computing*, 38(5):A3291–A3317, 2016.
12. Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
13. Praneeth Netrapalli, UN Niranjan, Sujay Sanghavi, Animashree Anandkumar, and Prateek Jain. Non-convex robust pca. In *Advances in Neural Information Processing Systems*, pages 1107–1115, 2014.
14. John Wright, Arvind Ganesh, Shankar Rao, Yigang Peng, and Yi Ma. Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization. In *Advances in neural information processing systems*, pages 2080–2088, 2009.
15. Huan Xu, Constantine Caramanis, and Sujay Sanghavi. Robust pca via outlier pursuit. In *Advances in Neural Information Processing Systems*, pages 2496–2504, 2010.
16. Xinyang Yi, Dohyung Park, Yudong Chen, and Constantine Caramanis. Fast algorithms for robust pca via gradient descent. In *Advances in neural information processing systems*, pages 4152–4160, 2016.
17. Alp Yurtsever, Olivier Fercoq, Francesco Locatello, and Volkan Cevher. A conditional gradient framework for composite convex minimization with applications to semidefinite programming. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5727–5736, Stockholmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
18. Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.