

Efficient Optimization Algorithms for Robust Principal Component Analysis and Its Variants

Shiqian Ma, and Necdet Serhat Aybat

Abstract—Robust PCA has drawn significant attention in the last decade due to its success in numerous application domains, ranging from bio-informatics, statistics, and machine learning to image and video processing in computer vision. Robust PCA and its variants such as sparse PCA and stable PCA can be formulated as optimization problems with exploitable special structures. Many specialized efficient optimization methods have been proposed to solve robust PCA and related problems. In this paper we review existing optimization methods for solving convex and nonconvex relaxations/variants of robust PCA, discuss their advantages and disadvantages, and elaborate on their convergence behaviors. We also provide some insights for possible future research directions including new algorithmic frameworks that might be suitable for implementing on multi-processor setting to handle large-scale problems.

Index Terms—PCA, Robust PCA, Convex Optimization, Non-convex Optimization, Iteration Complexity, Convergence Rate, ϵ -Stationary Solution

I. INTRODUCTION

PRINCIPAL component analysis (PCA) is a fundamental tool in statistics and data science. It obtains a low-dimensional expression for high-dimensional data in an ℓ_2 sense. However, it is known that the classical PCA is sensitive to gross errors. Robust PCA (RPCA) has been proposed to remove the effect of sparse gross errors. For a given data matrix $M \in \mathbb{R}^{m \times n}$, RPCA seeks to decompose it into two parts $M := L^\circ + S^\circ$ where L° is a low-rank matrix and S° is a sparse matrix. That is, RPCA assumes that M is a superposition of L° and S° . As a result, the gross errors will be captured by the sparse matrix S° so that the low-rank matrix L° can still approximate M well. RPCA does not only provide a low-dimensional approximation which is robust to outliers, it also finds vast applications in a variety of real applications such as computer vision [1], image alignment [2], subspace recovery [3], clustering [4] and so on.

Mathematically, [1], [5], [6], [7] investigated the conditions on the low-rank and sparse components (L°, S°) so that the inverse problem of recovering unknown (L°, S°) given M is well defined. One particular formulation of RPCA can be stated as follows:

$$\min_{L, S \in \mathbb{R}^{m \times n}} \text{rank}(L) + \rho \|S\|_0 \text{ s.t. } L + S = M, \quad (1)$$

S. Ma is with the Department of Mathematics at University of California, Davis, CA, 95616 USA e-mail: (sqma@math.ucdavis.edu).

N. S. Aybat is with the Department of Industrial and Manufacturing Engineering at Penn State, University Park, PA, 16802 USA e-mail: (nsa10@psu.edu).

Manuscript received January xx, 2018; revised xxx xx, 2018.

where $\|S\|_0$ is called the ℓ_0 -norm¹ of S and counts the number of nonzero entries of S , and $\rho > 0$ is a tradeoff parameter. It is known that (1) is NP-hard and thus numerically intractable. Later, it was shown in [1], [5], [6] that under certain conditions, (1) is equivalent to the following convex program with high probability:

$$\min_{L, S \in \mathbb{R}^{m \times n}} \|L\|_* + \rho \|S\|_1 \text{ s.t. } L + S = M, \quad (2)$$

where $\|L\|_*$ is called the nuclear norm of L and equal to the sum of the singular values of L , and $\|S\|_1 := \sum_{ij} |S_{ij}|$ is called the ℓ_1 norm of S . The optimization problem in (2) is called robust principal component pursuit (RPCP), and it can be reformulated as a semidefinite program (SDP) [8] and solved by an interior point method for SDPs. However, RPCA problems arising in practice are usually of very large scale, and interior point methods do not scale well for these problems. More efficient algorithms that solve (2) and its variants by exploiting the structure in these problems were studied extensively in the literature. One variant of (2) deals with an additional dense noise component. In particular, when M contains also a dense noise component N° such that $\|N^\circ\|_F \leq \sigma$ for some noise level $\sigma > 0$, i.e., $M = L^\circ + S^\circ + N^\circ$, instead of RPCP formulation in (2), the following so-called stable PCP (SPCP) problem is solved:

$$\min_{L, S \in \mathbb{R}^{m \times n}} \|L\|_* + \rho \|S\|_1 \text{ s.t. } \|L + S - M\|_F \leq \sigma. \quad (3)$$

It is proved in [9] that, under certain conditions on M , solving (3) gives a stable estimate of L° and S° with high probability in the sense that $\|\hat{L} - L^\circ\|_F^2 + \|\hat{S} - S^\circ\|_F^2 \leq \mathcal{O}(mn\sigma^2)$ where (\hat{L}, \hat{S}) denotes the optimal solution to (3). Since (3) satisfies the Slater's condition, it is equivalent to the following unconstrained problem for an appropriately chosen penalty parameter $\mu > 0$ depending on σ :

$$\min_{L, S \in \mathbb{R}^{m \times n}} \|L\|_* + \rho \|S\|_1 + \frac{\mu}{2} \|L + S - M\|_F^2. \quad (4)$$

Note that if M is only partially observed, that is, if we only have observations on M_{ij} for some indices (i, j) from a subset Ω , then (2), (3) and (4) can be respectively reformulated as

$$\min_{L, S \in \mathbb{R}^{m \times n}} \|L\|_* + \rho \|S\|_1 \text{ s.t. } P_\Omega(L + S - M) = 0, \quad (5)$$

$$\min_{L, S \in \mathbb{R}^{m \times n}} \|L\|_* + \rho \|S\|_1 \text{ s.t. } \|P_\Omega(L + S - M)\|_F \leq \sigma, \quad (6)$$

¹Technically, it is not a norm because it is not homogeneous; but, we still call it a norm following the convention.

$$\min_{L, S \in \mathbb{R}^{m \times n}} \|L\|_* + \rho \|S\|_1 + \frac{\mu}{2} \|P_\Omega(L + S - M)\|_F^2, \quad (7)$$

where the operator $P_\Omega : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ is defined as $[P_\Omega(M)]_{ij} = M_{ij}$, if $(i, j) \in \Omega$, and $[P_\Omega(M)]_{ij} = 0$ otherwise. Most algorithms we discuss in this paper for solving (2), (3) and (4) can be used to solve (5), (6) and (7) directly or with very little modification. For brevity, we will only describe algorithms for solving (2), (3) and (4) in this paper.

For the sake of completeness, we here briefly describe the results in [1], [9]. Assume that given data matrix $M \in \mathbb{R}^{m \times n}$ is a superposition of the *unknown* component matrices L^o , S^o , and N^o , i.e., $M = L^o + S^o + N^o$, such that L^o is low-rank ($r := \mathbf{rank}(L^o) \ll \min\{m, n\}$), S^o is sparse ($s := \|S^o\|_0 \ll mn$), and $\|N^o\| \leq \sigma$ for some $\sigma > 0$. Robust/stable PCA is an inverse problem with the objective of recovering L^o and S^o from the data matrix M . Clearly, if S^o is low-rank and/or L^o is sparse, the recovery is hopeless. To avoid these pathological instances, [1], [9] consider (L^o, S^o) pairs coming from a particular class satisfying some incoherence and randomness conditions. Suppose the singular value decomposition (SVD) of L^o is given by $L^o = \sum_{i=1}^r \sigma_i u_i v_i^\top$, where $U = [u_1, \dots, u_r]$ and $V = [v_1, \dots, v_r]$ are formed by its left- and right-singular vectors. The incoherence conditions assure that there exists a parameter $\delta > 0$ such that

$$\max_i \|U^\top e_i\|^2 \leq \delta r/m, \quad \max_i \|V^\top e_i\|^2 \leq \delta r/n, \quad (8a)$$

$$\|UV^\top\|_\infty \leq \sqrt{\frac{\delta r}{mn}}, \quad (8b)$$

where $\|Z\|_\infty := \max_{ij} |Z_{ij}|$ and e_i denotes the i -th unit vector. For S^o , it is assumed that the set of indices for the nonzero entries is random and follows a uniform distribution among all the subsets of cardinality s . Roughly speaking, these conditions assure that the low-rank matrix L^o is not sparse, and the sparse matrix S^o is not low-rank. Under these assumptions it is shown in [1] that when $N^o = \mathbf{0}$, i.e., $\sigma = 0$, solving the convex program (2) with $\rho = \frac{1}{\sqrt{\max\{m, n\}}}$ recovers the optimal solution of (1), (L^*, S^*) , with high probability and $(L^*, S^*) = (L^o, S^o)$ provided that L^o is sufficiently low-rank and S^o is sufficiently sparse comparing to the matrix size – see also [6]. These results extend to the case where M is partially observed; indeed, Candès et al. [1] show that solving (5) recovers L^o under similar conditions. Moreover, in [9], the authors showed that under certain conditions, solving the convex problem (3) again with $\rho = \frac{1}{\sqrt{\max\{m, n\}}}$ generates a low-rank and sparse decomposition (L^*, S^*) such that $\|L^* - L^o\|_F^2 + \|S^* - S^o\|_F^2 \leq Cmn\sigma^2$ for some constant $C > 0$ (independent of m , n and σ) with high probability – note when $N = \mathbf{0}$, the recovery is exact with high probability.

Recently, there are works that further study statistical guarantees of different RPCA models. Zhang, Zhou and Liang [10] provide a refined analysis of RPCA which allows the support of the sparse error matrix to be generated with non-uniform sampling, i.e., entries of the low-rank matrix are corrupted with different probabilities – hence, one can model the scenario where some entries are more prone to corruption than the others. A nonconvex model of RPCA is studied in [11] and a gradient descent method with proper initialization is

shown to be able to reduce the computational complexity comparing with existing methods. Zhang and Yang [12] consider a nonconvex optimization formulation with manifold constraint for RPCA. Two algorithms for manifold optimization are proposed in [12] and it is shown that they can reduce the dependence on the condition number of the underlying low-rank matrix theoretically. Netrapalli et al. [13] consider another nonconvex formulation of RPCA and analyze the iteration complexity of the proposed alternating projection method.

There are also recent survey papers [14], [15] that discuss algorithms for solving RPCA, but these papers mainly focus on its convex relaxations. In this paper, we aim to review existing algorithms for both convex and nonconvex relaxations/variants of RPCA models and point out a few possible future directions.

II. ALGORITHMS FOR CONVEX RELAXATIONS/VARIANTS OF RPCA

The earliest first-order methods for solving the convex RPCP problem are given in [16], [17]. In [16], the authors proposed an accelerated proximal gradient method (APGM) [33], [34], [35] for solving (4) in which each iteration involves computing the proximal mappings of the nuclear norm $\|L\|_*$ and the ℓ_1 norm $\|S\|_1$. In particular, the non-accelerated proximal gradient method (PGM) for solving (4) simply updates L and S as

$$\begin{aligned} G^k &:= \mu(L^k + S^k - M) \\ L^{k+1} &:= \operatorname{argmin}_L \|L\|_* + \frac{1}{2\tau} \|L - (L^k - \tau G^k)\|_F^2 \\ S^{k+1} &:= \operatorname{argmin}_S \rho \|S\|_1 + \frac{1}{2\tau} \|S - (S^k - \tau G^k)\|_F^2. \end{aligned} \quad (9)$$

Note that G^k is the gradient of the quadratic penalty function in (4) and $\tau > 0$ denotes a step size. The two subproblems in (9) both admit easy closed-form optimal solutions. Specifically, the solution of the L -subproblem corresponds to the proximal mapping of the nuclear norm, which is given by

$$L^{k+1} = \mathcal{S}_\tau(L^k - \tau G^k),$$

where the matrix shrinkage operation is defined as

$$\mathcal{S}_\nu(Z) = U \operatorname{diag}((\sigma - \nu)_+) V^\top, \quad (10)$$

where $Z = U \operatorname{diag}(\sigma) V^\top$ is the SVD of Z , and $z_+ := \max(0, z)$. The solution of the S -subproblem corresponds to the proximal mapping of the ℓ_1 norm, which is given by

$$S^{k+1} = s_{\rho\tau}(S^k - \tau G^k),$$

where the vector shrinkage operation is defined as

$$[s_\nu(Z)]_{ij} = \operatorname{sign}(Z_{ij}) \circ \max\{0, |Z_{ij}| - \nu\}, \quad (11)$$

where $\operatorname{sign}(a)$ denotes the sign of a , and \circ denotes the Hadamard product.

APGM incorporates Nesterov's acceleration technique and updates the variables as follows starting with $t_{-1} = t_0 = 1$.

$$\begin{aligned} \bar{L}^k &:= L^k + \frac{t_{k-1}-1}{t_k} (L^k - L^{k-1}) \\ \bar{S}^k &:= S^k + \frac{t_{k-1}-1}{t_k} (S^k - S^{k-1}) \\ \bar{G}^k &:= \mu(\bar{L}^k + \bar{S}^k - M) \\ L^{k+1} &:= \operatorname{argmin}_L \|L\|_* + \frac{1}{2\tau} \|L - (\bar{L}^k - \tau \bar{G}^k)\|_F^2 \\ S^{k+1} &:= \operatorname{argmin}_S \rho \|S\|_1 + \frac{1}{2\tau} \|S - (\bar{S}^k - \tau \bar{G}^k)\|_F^2 \\ t_{k+1} &:= (1 + \sqrt{1 + 4t_k^2})/2. \end{aligned} \quad (12)$$

TABLE I
Convergence rates of different algorithms.

| Algorithm | Problem | ϵ -optimality measure | Convergence rate |
|--|----------------|--------------------------------|---|
| Algorithms for Convex Models | | | |
| PGM [16] | (4) | objective value error | $\mathcal{O}(1/k)$ |
| APGM [16] | (4) | objective value error | $\mathcal{O}(1/k^2)$ |
| IALM [17] | (2) | — | convergence, no rate given |
| ADMM [18] | (2) | — | convergence, no rate given |
| ALM [19] | (17) | objective value error | $\mathcal{O}(1/k)$ |
| FALM [19] | (17) | objective value error | $\mathcal{O}(1/k^2)$ |
| ASALM [20] | (19) | — | convergence unclear, no rate given |
| VASALM [20] | (19) | — | convergence, no rate given |
| PSPG [21] | (3) | objective value error | $\mathcal{O}(1/k)$ |
| ADMIP [22] | (3) | objective value error | $\mathcal{O}(1/k)$ |
| Quasi-Newton method (fastRPCA) [23] | (26) | — | convergence, no rate given |
| 3-block ADMM [24] | (28) | — | convergence, no rate given |
| Frank-Wolfe [25] | (30) | objective value error | $\mathcal{O}(1/k)$ |
| Algorithms for Nonconvex Models | | | |
| GoDec [26] | (33) | — | local convergence, no rate given |
| GreBsmo [27] | (36) | — | convergence unclear, no rate given |
| Alternating Minimization (R2PCP) [28] | (35) | — | local convergence, no rate given |
| Gradient Descent (GD) [11] | \approx (36) | — | linear convergence with proper initialization and incoherence assumption |
| Alternating Minimization [29] | (37) | — | local convergence with proper initialization and incoherence and RIP assumptions |
| Stochastic alg. [30] | (39) | — | convergence if the iterates are always full rank matrices, no rate given |
| LMafit [31] | (44) | — | convergence if difference between two consecutive iterates tends to zero, no rate given |
| Conditional Gradient [32] | (48) | perturbed KKT | $\mathcal{O}(1/\sqrt{k})$ |
| ADMM [32] | (51) | perturbed KKT | $\mathcal{O}(1/\sqrt{k})$ |
| Proximal BCD [32] | (54) | perturbed KKT | $\mathcal{O}(1/\sqrt{k})$ |

[†]Note: Some of these algorithms solve different problems and the ϵ -optimality measures are also different, so the convergence rates are not directly comparable with each other. [11] has no explicit optimization formulation but the objective is similar to (36). Moreover, global convergence is usually guaranteed for convex solvers, but only local convergence is usually guaranteed for nonconvex solvers, unless certain very strong assumptions are made.

Results in [33], [34], [35] show that the proximal gradient method (9) and the accelerated proximal gradient method (12) find an ϵ -optimal solution to (4) in no more than $\mathcal{O}(1/\epsilon)$ and $\mathcal{O}(1/\sqrt{\epsilon})$ iterations, respectively.

When there is no noise, i.e., $\sigma = 0$, the problem of interest is (2). The drawback of the above approach for solving the unconstrained version in (4) is that (4) is equivalent to (2) only when $\mu \rightarrow +\infty$. Therefore, for any fixed $\mu > 0$, there is always a residual term which does not go to zero. To remedy this, the same group of authors [17] considered the augmented Lagrangian method (ALM) for solving (2). By associating a Lagrange multiplier Λ to the linear equality constraint, the augmented Lagrangian function of (2) can be written as

$$\mathcal{L}_\beta(L, S; \Lambda) := \|L\|_* + \rho \|S\|_1 - \langle \Lambda, L + S - M \rangle + \frac{\beta}{2} \|L + S - M\|_F^2,$$

where $\beta > 0$ is a penalty parameter. A typical iteration of ALM iterates the updates as follows:

$$(L^{k+1}, S^{k+1}) := \operatorname{argmin}_{L, S} \mathcal{L}_\beta(L, S; \Lambda^k) \quad (13a)$$

$$\Lambda^{k+1} := \Lambda^k - \beta(L^{k+1} + S^{k+1} - M). \quad (13b)$$

Note that the first step in (13) requires to minimize the augmented Lagrangian function with respect to L and S simultaneously, which usually is computationally very expensive and almost as hard as solving the original problem in (2). In [17], the authors proposed both exact and inexact versions of ALM, where the former one solves the subproblems (almost

exactly and the latter one solves the subproblems inexactly according to a particular subproblem termination criterion. Both the exact ALM and inexact ALM (IALM) employ some iterative algorithm for minimizing the augmented Lagrangian function until certain overall stopping criterion is met, which may require many iterations and thus time consuming. Around the same time when [16], [17] appeared, the alternating direction method of multipliers (ADMM) was revisited and found very successful in solving signal processing and image processing problems [36], [37], [38], [39]. It was then found that RPCP in (2) can be nicely solved by ADMM due to its special separable structure [18], [19]. The ADMM iterations for solving (2) take the following form:

$$\begin{aligned} L^{k+1} &:= \operatorname{argmin}_L \mathcal{L}_\beta(L, S^k; \Lambda^k) \\ S^{k+1} &:= \operatorname{argmin}_S \mathcal{L}_\beta(L^{k+1}, S; \Lambda^k) \\ \Lambda^{k+1} &:= \Lambda^k - \beta(L^{k+1} + S^{k+1} - M). \end{aligned} \quad (14)$$

Comparing to ALM in (13), it is noted that ADMM splits the subproblem in (13a) into two smaller subproblems that correspond to computing proximal mappings of $\|L\|_*$ and $\|S\|_1$, respectively. The ADMM (14) is known as two-block ADMM as there are two block variables L and S and hence two subproblems are solved in each iteration of the algorithm. It is now widely known that the two-block ADMM is a special case of the so-called Douglas-Rachford operator splitting method [40], [41], [42], [43] applied to the dual problem, and the two-block ADMM for solving convex problems globally

converges for any penalty parameter $\beta > 0$ [44] and converges with a sublinear rate $O(1/k)$ (see, e.g., [45], [46], [47]).

The alternating linearization method (ALM) proposed by Goldfarb, Ma and Scheinberg [19] is shown to be equivalent to a symmetric version of ADMM (14) with either $\|L\|_*$ or $\|S\|_1$ replaced with some suitable smooth approximation. For instance, given $\nu > 0$, define $g_\nu : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ such that

$$g_\nu(S) = \max_{Z \in \mathbb{R}^{m \times n}} \left\{ \langle S, Z \rangle - \frac{\nu}{2} \|Z\|_F^2 : \|Z\|_\infty \leq \rho \right\}, \quad (15)$$

and let $g(S) = \rho \|S\|_1$. Clearly, $g_\nu \rightarrow g$ uniformly as $\nu \searrow 0$. Moreover, g_ν is a differentiable convex function such that ∇g_ν is Lipschitz continuous. Indeed, given $S \in \mathbb{R}^{m \times n}$, let $Z_\nu(S)$ be the maximizer for (15), which in closed form can be written as $Z_\nu(S) = \text{sign}(S) \circ \max\{\frac{1}{\nu}|S|, \rho \mathbf{1}_{m \times n}\}$, and $\nabla g_\nu(S) = Z_\nu(S)$ is Lipschitz continuous with constant $C_{g_\nu} = \frac{1}{\nu}$. Similarly, given $\mu > 0$, define $f_\mu : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ such that

$$f_\mu(L) = \max_{W \in \mathbb{R}^{m \times n}} \left\{ \langle L, W \rangle - \frac{\mu}{2} \|W\|_F^2 : \|W\| \leq 1 \right\}, \quad (16)$$

where $\|\cdot\|$ denotes the spectral norm, and let $f(L) = \lambda \|L\|_*$. Clearly, $f_\mu \rightarrow f$ uniformly as $\mu \searrow 0$. Moreover, f_μ is a differentiable convex function such that ∇f_μ is Lipschitz continuous. Indeed, given $L \in \mathbb{R}^{m \times n}$, let $W_\mu(L)$ be the maximizer for (16), which in closed form can be written as $W_\mu(L) = U \text{diag}(\max\{\frac{\sigma}{\mu} - 1, \mathbf{0}\}) V^\top$, where $L = U \text{diag}(\sigma) V^\top$ is the singular value decomposition of L with $\sigma \in \mathbb{R}_{++}^r$ denoting the vector of singular values; moreover, $\nabla f_\mu(S) = W_\mu(L)$ is Lipschitz continuous with constant $C_{f_\mu} = \frac{1}{\mu}$.

The alternating linearization method in [19] can be applied to solve the following problem, which is a smoothed version of (2):

$$\min \|L\|_* + g_\nu(S), \text{ s.t., } L + S = M. \quad (17)$$

Denote the augmented Lagrangian function of (17) as $\tilde{\mathcal{L}}_\beta(L, S; \Lambda)$, the alternating linearization method in [19] iterates the updates as follows:

$$\begin{aligned} L^{k+1} &:= \text{argmin}_L \tilde{\mathcal{L}}_\beta(L, S^k; \Lambda^k), \\ \Lambda^{k+\frac{1}{2}} &:= \Lambda^k - \beta(L^{k+1} + S^k - M), \\ S^{k+1} &:= \text{argmin}_S \tilde{\mathcal{L}}_\beta(L^{k+1}, S; \Lambda^{k+\frac{1}{2}}), \\ \Lambda^{k+1} &:= \Lambda^{k+\frac{1}{2}} - \beta(L^{k+1} + S^{k+1} - M). \end{aligned} \quad (18)$$

The authors in [19] proved that this method has a sublinear convergence rate $O(1/k)$. They also proposed an accelerated version of (18) (FALM) in [19] by adopting Nesterov's acceleration technique, and proved that the accelerated alternating linearization method has a better sublinear convergence rate $O(1/k^2)$.

Based on the success of two-block ADMM for solving (2), it is then very natural to apply ADMM to solve SPCP in (3). To do so, one has to introduce a new variable N , and rewrite (3) equivalently as

$$\min_{L, S, N \in \mathbb{R}^{m \times n}} \|L\|_* + \rho \|S\|_1 \quad (19) \\ \text{s.t., } L + S + N = M, \|N\|_F \leq \sigma.$$

The ADMM for solving (19) iterates as follows with three block variables:

$$\begin{aligned} L^{k+1} &:= \text{argmin}_L \mathcal{L}_\beta(L, S^k, N^k; \Lambda^k) \\ S^{k+1} &:= \text{argmin}_S \mathcal{L}_\beta(L^{k+1}, S, N^k; \Lambda^k) \\ N^{k+1} &:= \text{argmin}_N \mathcal{L}_\beta(L^{k+1}, S^{k+1}, N; \Lambda^k) \\ \Lambda^{k+1} &:= \Lambda^k - \beta(L^{k+1} + S^{k+1} + N^{k+1} - M), \end{aligned} \quad (20)$$

where the augmented Lagrangian function for (19) is defined as

$$\begin{aligned} \mathcal{L}_\beta(L, S, N; \Lambda) &:= \|L\|_* + \rho \|S\|_1 + \mathbf{1}(N \mid \|N\|_F \leq \sigma) \\ &\quad - \langle \Lambda, L + S + N - M \rangle + \frac{\beta}{2} \|L + S + N - M\|_F^2, \end{aligned}$$

where $\mathbf{1}(N \mid \mathcal{N})$ denotes the indicator function of the set $\{N \in \mathcal{N}\}$, i.e., $\mathbf{1}(N \mid \mathcal{N}) = 0$ if $N \in \mathcal{N}$ and $\mathbf{1}(N \mid \mathcal{N}) = \infty$ otherwise. Note that the three subproblems in (20) all have closed-form solutions. In particular, the L -subproblem corresponds to the proximal mapping of $\|L\|_*$, the S -subproblem corresponds to proximal mapping of $\|S\|_1$, and the N -subproblem corresponds to projection onto the set $\{N \mid \|N\|_F \leq \sigma\}$. Similar idea was used in [2] for robust image alignments. In practice, this three-block ADMM usually works very well. However, it was later discovered that the ADMM with more than two block variables is not necessarily convergent in general [48]. Note that although (19) contains three block variables, it can be viewed as a two-block problem, if we group S and N as one (larger) block variable. One of the earliest methods for solving SPCP in (3) and (19) is a three-block ADMM algorithm, ASALM, proposed by Tao and Yuan [20], and although it does not have any convergence guarantees, it works well in practice; and slightly changing the update rule in ASALM leads to VASALM, of which iterate sequence converges to an optimal solution; but this comes at the cost of degradation in practical convergence speed when compared to ASALM – indeed, VASALM [20] can be seen as a linearized version of two-block ADMM to solve (3) with a convergence guarantee without any convergence rate result. To remedy the shortcoming associated with the theoretical convergence of three-block ADMM, several other alternatives based on two-block ADMM were proposed [49], [22], [21].

Aybat, Goldfarb and Ma [21] proposed an accelerated proximal gradient method, PSPG, for solving SPCP in (3). First, (3) is reformulated with a partially smooth objective. In particular, the nuclear norm is smoothed according to (16):

$$\min_{L, S \in \mathbb{R}^{m \times n}} f_\mu(L) + \rho \|S\|_1 \text{ s.t. } (L, S) \in \chi, \quad (21)$$

$$\chi := \{(L, S) \mid \|L + S - M\|_F \leq \sigma\}, \quad (22)$$

where $\mu > 0$ is a given smoothing parameter. An accelerated proximal gradient method such as [34], [33] can be applied to solve (21), because it was shown in [21] that the following subproblem is easy to solve:

$$\min_{L, S} \|S\|_1 + \frac{1}{2\xi} \|L - \tilde{L}\|_F^2 \text{ s.t. } (L, S) \in \chi, \quad (23)$$

where $\xi > 0$ denotes a step size of the proximal gradient step and \tilde{L} denotes some known matrix. This operation requires one sorting which has $O(mn \log(mn))$ complexity.

For any $\epsilon > 0$, setting $\mu = \Omega(\epsilon)$, PSPG proposed in [21] can compute an ϵ -optimal solution to (3) within $O(1/\epsilon)$ iterations,

and its computational complexity per iteration is comparable to the work per iteration required by ASALM and VASALM, which is mainly determined by an SVD computation. On the other hand, it is also important to emphasize that PSPG iterate sequences do not converge to an optimal solution to the SPCP problem in (3). In particular, since within PSPG the smoothing parameter μ is fixed, depending on the approximation parameter ϵ for solving (21), further iterations after reaching an ϵ -optimal solution in $\mathcal{O}(1/\epsilon)$ iterations do not necessarily improve the solution quality.

In [49], [22], the *variable penalty* ADMM algorithm ADMIP (Alternating Direction Method with Increasing Penalty) is proposed to solve the following equivalent formulation for (3) using the variable splitting trick:

$$\min_{\hat{L}, L, S \in \mathbb{R}^{m \times n}} \|\hat{L}\|_* + \rho \|S\|_1 \text{ s.t. } (L, S) \in \chi, \quad L = \hat{L}. \quad (24)$$

The augmented Lagrangian function of (24) can be written as

$$\mathcal{L}_\beta(\hat{L}, L, S; \Lambda) := \|\hat{L}\|_* + \rho \|S\|_1 - \langle \Lambda, \hat{L} - L \rangle + \frac{\beta}{2} \|\hat{L} - L\|_F^2.$$

Given a nondecreasing penalty parameter sequence $\{\beta^k\}_{k \in \mathbb{Z}_+}$, ADMIP updates the variables as follows:

$$\hat{L}^{k+1} := \underset{\hat{L}}{\operatorname{argmin}} \mathcal{L}_{\beta^k}(\hat{L}, L^k, S^k; \Lambda^k) \quad (25a)$$

$$(L^{k+1}, S^{k+1}) := \underset{(L, S) \in \chi}{\operatorname{argmin}} \mathcal{L}_{\beta^k}(\hat{L}^{k+1}, L, S; \Lambda^k) \quad (25b)$$

$$\Lambda^{k+1} := \Lambda^k - \beta^k (\hat{L}^{k+1} - L^{k+1}). \quad (25c)$$

The step in (25a) requires computing a soft thresholding on the singular values of an $m \times n$ matrix and the step in (25b) requires an operation given in (23).

Under mild conditions on the penalty parameter sequence, Aybat and Iyengar show that the primal-dual ADMIP iterate sequence converges to an optimal primal-dual solution to the SPCP problem in (24) – hence, $\{(L^k, S^k)\}_{k \in \mathbb{Z}_+}$ converges to an optimal solution to (3), and when constant penalty parameter is used as a special case, it can compute an ϵ -optimal solution within $\mathcal{O}(1/\epsilon)$ iterations, of which complexity is determined by an SVD. In particular, one needs the penalty parameter sequence $\{\beta^k\}_{k \in \mathbb{Z}_+}$ to be non-decreasing and to satisfy $\sum_k (\beta^k)^{-1} = +\infty$. The main advantages of adopting an increasing sequence of penalties are as follows:

- 1) The algorithm is robust in the sense that there is no need to search for problem data dependent β^* that works well in practice.
- 2) The algorithm is likely to achieve primal feasibility faster.
- 3) The complexity of initial (transient) iterations can be controlled through controlling $\{\beta^k\}$. The main computational bottleneck in ADMIP is the SVD computation in (25a). Since the optimal L^* is of low-rank, and $L_k \rightarrow L^*$, eventually the SVD computations are likely to be very efficient. However, since the initial iterates in the transient phase of the algorithm may have large rank, the complexity of the SVD in the initial iterations can be quite large. To compute the solution to the subproblem in (25a), one does not need to compute

singular values smaller than $1/\beta^k$; hence, initializing ADMIP with a *small* $\beta^0 > 0$ will significantly decrease the complexity of initial iterations through employing *partial* SVD computations, e.g., Lanczos-based methods such as PROPACK [50].

In [22], Aybat and Iyengar compared ADMIP against ASALM on both randomly generated synthetic problems and surveillance video foreground extraction problems. According to numerical results reported in [22], on the synthetic problems ASALM requires about *twice* as many iterations for convergence, while the total runtime for ASALM is considerably larger.

Aravkin et al. [23] proposed solving

$$\min \psi(L, S) \quad \text{s.t.} \quad \phi(\mathcal{A}(L, S) - M) \leq \sigma, \quad (26)$$

where $\mathcal{A} : \mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ is a linear operator, $\phi : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is a smooth convex loss-function, and ψ can be set to either one of the following functions:

$$\begin{aligned} \psi_{\text{sum}}(L, S) &:= \|L\|_* + \rho \|S\|_1, \\ \psi_{\text{max}}(L, S) &:= \max\{\|L\|_*, \rho_{\text{max}} \|S\|_1\}, \end{aligned}$$

$\rho, \rho_{\text{max}} > 0$ are some given function parameters. Note that setting $\psi = \psi_{\text{sum}}$, $\rho(\cdot) = \|\cdot\|_F^2$, and $\mathcal{A}(L, S) = \pi_\Omega L + S$ in (26), one obtains the SPCP problem in (6). This approach offers advantages over the original SPCP formulation in terms of practical parameter selection. The authors make a case that although setting $\rho = \frac{1}{\sqrt{\max\{m, n\}}}$ in (3) has theoretical justification as briefly discussed in the introduction section, many practical problems may violate the underlying assumptions in (8); in those cases one needs to tune ρ via cross validation, and selecting ρ_{max} in ψ_{max} might be easier than selecting ρ in ψ_{sum} . Instead of solving (26) directly, a convex variational framework, accelerated with a “quasi-Newton” method, is proposed. In particular, Newton’s method is used to find a root of the value function:

$$v(\tau) := \min_{L, S \in \mathbb{R}^{m \times n}} \phi(\mathcal{A}(L, S) - A) - \sigma \quad \text{s.t.} \quad \psi(L, S) \leq \tau, \quad (27)$$

i.e., given $\sigma > 0$ compute $\tau^* = \tau(\sigma)$ such that $v(\tau^*) = 0$. According to results in [51], if the constraint in (26) is tight at an optimal solution, then there exists $\tau^* = \tau(\sigma)$ such that $v(\tau^*) = 0$ and the corresponding optimal solution to (27) is also optimal to (26). Within Newton’s method for root finding, to compute the next iterate τ^{k+1} , one can compute the derivative of the value function at the current iterate τ^k as follows $v'(\tau_k) = -\psi^\circ(\mathcal{A}^\top \nabla \phi(\mathcal{A}(L_k, S_k) - A))$, where ψ° denotes the polar gauge to ψ and (L_k, S_k) denotes the optimal solution to (27) at $\tau = \tau_k$ – Aravkin et al. proposed a projected “Quasi-Newton” method to solve (27). According to numerical tests reported in [51], QN-max, the quasi-Newton method running on (26) with $\psi = \psi_{\text{max}}$ and $\phi(\cdot) = \|\cdot\|_F^2$, is competitive with the state-of-the-art codes, ASALM [20], PSPG [21], and ADMIP [22].

In a recent work [24], Lin, Ma and Zhang considered the penalty formulation of the SPCP problem, which is equivalent to solving (3) for certain noise level $\sigma > 0$:

$$\begin{aligned} \min \quad & \|L\|_* + \rho \|S\|_1 + \mu \|N\|_F^2 \\ \text{s.t.} \quad & L + S + N = M, \end{aligned} \quad (28)$$

where $\rho > 0$ is the sparsity tradeoff parameter and $\mu > 0$ is a suitable penalty parameter depending on the noise level $\sigma > 0$. The authors showed that the following 3-block ADMM for solving (28) globally converges for any penalty parameter $\beta > 0$.

$$\begin{aligned} L^{k+1} &:= \operatorname{argmin}_L \mathcal{L}_\beta(L, S^k, N^k; \Lambda^k) \\ S^{k+1} &:= \operatorname{argmin}_S \mathcal{L}_\beta(L^{k+1}, S, N^k; \Lambda^k) \\ N^{k+1} &:= \operatorname{argmin}_N \mathcal{L}_\beta(L^{k+1}, S^{k+1}, N; \Lambda^k) \\ \Lambda^{k+1} &:= \Lambda^k - \beta(L^{k+1} + S^{k+1} + N^{k+1} - M), \end{aligned} \quad (29)$$

where the augmented Lagrangian function is

$$\begin{aligned} \mathcal{L}_\beta(L, S, N; \Lambda) &:= \|L\|_* + \rho \|S\|_1 + \mu \|N\|_F^2 \\ &\quad - \langle \Lambda, L + S + N - M \rangle + \frac{\beta}{2} \|L + S + N - M\|_F^2. \end{aligned}$$

Note that the three subproblems in (29) are all easy to solve. Specifically, the L -subproblem corresponds to the proximal mapping of nuclear norm $\|L\|_*$, the S -subproblem corresponds to the proximal mapping of ℓ_1 norm, and the N -subproblem admits a very easy analytical solution.

The Frank-Wolfe method (aka conditional gradient method) [52] was revisited recently for solving large-scale machine learning problems [53], [54]. RPCA is a representative example that is suitable for Frank-Wolfe method. Note that algorithms discussed above usually involve computing the proximal mapping of the nuclear norm, which is given by an SVD in (10). Computing full SVD for a large matrix in every iteration can be very time consuming. In contrast, the Frank-Wolfe method deals with nuclear norm in a much simpler manner, which only computes the largest singular value of a matrix in each iteration. The Frank-Wolfe method for solving RPCA was proposed by Mu et al. in [25]. The authors in [25] considered the penalized variant of RPCA (4). However, this problem cannot be directly solved by the Frank-Wolfe method, because the Frank-Wolfe method requires a bounded constraint set. Therefore, the authors further reformulated (4) to the following problem for properly chosen constants λ_L , λ_S , U_L and U_S :

$$\begin{aligned} \min \quad & \frac{1}{2} \|L + S - M\|_F^2 + \lambda_L t_L + \lambda_S t_S, \\ \text{s.t.}, \quad & \|L\|_* \leq t_L \leq U_L, \|S\|_1 \leq t_S \leq U_S. \end{aligned} \quad (30)$$

The Frank-Wolfe method iterates the updates as follows:

$$\begin{aligned} G^k &:= L^k + S^k - M \\ (d_L^k, d_{t_L}^k) &:= \operatorname{argmin}_{\|L\|_* \leq t_L \leq U_L} \langle G^k, L \rangle + \lambda_L t_L \\ (d_S^k, d_{t_S}^k) &:= \operatorname{argmin}_{\|S\|_1 \leq t_S \leq U_S} \langle G^k, S \rangle + \lambda_S t_S \\ \gamma^k &:= 2/(k+2) \\ L^{k+1} &:= (1 - \gamma^k)L^k + \gamma^k d_L^k \\ t_L^{k+1} &:= (1 - \gamma^k)t_L^k + \gamma^k d_{t_L}^k \\ S^{k+1} &:= (1 - \gamma^k)S^k + \gamma^k d_S^k \\ t_S^{k+1} &:= (1 - \gamma^k)t_S^k + \gamma^k d_{t_S}^k. \end{aligned} \quad (31)$$

It was shown in [25] that the two minimization subproblems in (31) are easy to solve. In particular, solving the subproblem for (d_L, d_{t_L}) requires only to compute the largest singular value and its corresponding singular vector of an $m \times n$ matrix. This is a big saving compared with computing the full SVD as required for computing the proximal mapping of the nuclear norm. As a result, Frank-Wolfe method has better per-iteration

complexity than the proximal gradient method and ADMM algorithms discussed above, and thus may have better scalability for very large-scale problems. On the other hand, as pointed out in [25], one clear disadvantage of Frank-Wolfe method on (30) is that at every iteration only one entry of the sparse component is updated. This leads to very slow convergence in practice. Hence, Mu et al. [25] proposed combining Frank-Wolfe iterations with an additional proximal gradient step in S -block. In particular, they proposed after Frank-Wolfe iterate $(L_{k+1}, \tilde{S}_{k+1})$ is computed, an extra proximal gradient step is computed and S -block is updated again. Moreover, the authors also showed that this hybrid method obtained by combining Frank-Wolfe and proximal gradient steps enjoys a sublinear convergence rate $O(1/k)$ similar to Frank-Wolfe method given in (31).

As a special case of RPCA, one can consider that all columns of the low-rank matrix L are identical. That is, the given matrix M is a superposition of a special rank-one matrix L and a sparse matrix S . This special RPCA finds many interesting applications in practice such as video processing [55], [56] and bioinformatics [57]. For instance, in the background extraction of surveillance video, if the background is static, then the low-rank matrix L that corresponds to the background should have identical columns. As a result, the background and foreground can be separated by solving the following convex program:

$$\begin{aligned} \min_{x, S} \quad & \|S\|_1 \\ \text{s.t.} \quad & [x, x, \dots, x] \circ \mathbf{E} + S = M, \end{aligned} \quad (32)$$

where $[x, x, \dots, x]$ denotes the $m \times n$ matrix with all columns being x , \mathbf{E} denotes the $m \times n$ matrix with all ones. Note that the optimal x of (32) corresponds to the static background for all frames and S corresponds to the moving foreground. The advantage of (32) is that it does not involve nuclear norm. As a result, SVD can be avoided when designing algorithms for solving it which makes the resulting algorithms very efficient. Yang, Pong and Chen [58] adopted the similar idea and designed variants of ADMM algorithm for solving a more general model where the sparsity function of S is allowed to be a nonconvex function. Convergence of the proposed ADMM was proved under the assumption of KL property [59], [60] being satisfied. We will discuss these topics in more details in the next section.

III. ALGORITHMS FOR NONCONVEX RELAXATIONS/VARIANTS OF RPCA

In this section, we discuss nonconvex relaxations and variants of RPCA given in (1) and algorithms for solving them. Some researchers aim to (approximately) solve RPCA in (1) directly without convexifying the rank function and/or the ℓ_0 norm. In [26], Zhou and Tao considered a variant of (1):

$$\min \|L + S - M\|_F^2 \quad \text{s.t.} \quad \mathbf{rank}(L) \leq \tau_r, \|S\|_0 \leq \tau_s, \quad (33)$$

where τ_r and τ_s are given parameters to control the rank of L and sparsity of S . The authors proposed the GoDec algorithm which alternately minimizes the objective function in one variable while fixing the other, which is a special case of

alternating projection method analyzed in [61]. In particular, a naive version of GoDec algorithm iterates as follows:

$$\begin{aligned} L^{k+1} &:= \operatorname{argmin}_L \|L + S^k - M\|_F^2 \quad \text{s.t.} \quad \mathbf{rank}(L) \leq \tau_r, \\ S^{k+1} &:= \operatorname{argmin}_S \|L^{k+1} + S - M\|_F^2 \quad \text{s.t.} \quad \|S\|_0 \leq \tau_s. \end{aligned} \quad (34)$$

The two subproblems correspond to two projections. Although the projection for S is easy, the projection for L requires computing a partial SVD, which may be time consuming when the matrix size is large. The authors proposed to use a low-rank approximation based on bilateral random projections to approximate this projection operation which can significantly speed up the computation. The authors showed that the iterate sequence converges to a local minimum provided that the initial point is close to some point in the intersection of the two manifolds $\{L \mid \mathbf{rank}(L) \leq \tau_r\}$ and $\{S \mid \|S\|_0 \leq \tau_s\}$. The convergence of GoDec follows from the results in [61].

In [28], Hintermüller and Wu considered a regularized version of (33):

$$\begin{aligned} \min_{L, S \in \mathbb{R}^{m \times n}} & \|L + S - M\|_F^2 + \frac{\rho}{2} \|L\|_F^2 \\ \text{s.t.} & \quad \mathbf{rank}(L) \leq \tau_r, \quad \|S\|_0 \leq \tau_s, \end{aligned} \quad (35)$$

where $\tau_r, \tau_s > 0$ are given model parameters as in (33), and $0 \leq \rho \ll 1$ is a given regularization parameter. An *inexact alternating minimization method (R2PCP)* on matrix manifolds is proposed to solve (35). The iterates L_{k+1} and S_{k+1} are computed as “inexact” solutions to subproblems $\min_L \{\|L + S_k - M\|_F^2 + \rho \|L\|_F^2 : \mathbf{rank}(L) \leq \tau_r\}$ and $\min_S \{\|L_{k+1} + S - M\|_F^2 : \|S\|_0 \leq \tau_s\}$, respectively. Provided that a limit point of the iterate sequence exists, under some further restrictive technical assumptions, it is shown that first-order necessary optimality conditions are satisfied.

Note that the convex relaxation in (2) involves the nuclear norm $\|\cdot\|_*$ in the objective. Algorithms dealing with nuclear norm (like the ones discussed in Section II) usually require to compute its proximal mapping, which then require an SVD. This can be very time consuming when the problem size is large, even when $\min\{m, n\}$ is in the order of thousands. This has motivated researchers to consider nonconvex relaxations of RPCA that avoid SVD calculations. One way to achieve SVD-free methods is to factorize the low rank matrix $L \in \mathbb{R}^{m \times n}$ as a product of two low-rank matrices, i.e., factorize $L = UV^\top$, where $U \in \mathbb{R}^{m \times r}$, $V \in \mathbb{R}^{n \times r}$ and $r \ll \min\{m, n\}$ such that r is an upper bound on $\mathbf{rank}(L^\circ)$. This leads to many different nonconvex relaxations of RPCA.

In [27], Zhou and Tao considered a regularized version of (33):

$$\begin{aligned} \min_{U, V, S} & \rho_1 \|S\|_1 + \|UV^\top + S - M\|_F^2 \\ \text{s.t.} & \quad \mathbf{rank}(U) = \mathbf{rank}(V) \leq \tau_r, \end{aligned} \quad (36)$$

where $\rho_1 > 0$ and $\tau_r \in \mathbb{Z}_+$ such that $\tau_r \geq \mathbf{rank}(L^\circ)$. The authors propose a three-block alternating minimization algorithm, GreBsmo, for solving (36). The proposed algorithm lacks theoretical convergence guarantees; but, on the other hand, according to numerical results reported in [27], GreBsmo performs considerably better than both GoDec [26]

and inexact ALM method [17] (around 30-100 times faster than both) when applied to foreground extraction problems.

A nonconvex model of RPCA, similar to one in [27], is studied in [11] and a gradient descent (GD) method with proper initialization is proposed to solve it. The algorithm proposed in [11] has two phases and in both phases the objective is to reduce the function $Q(U, V, S) := \|UV^\top + S - M\|_F^2$. In the first phase, a sorting-based sparse estimator is used to generate a rough initial estimate S_0 to the unknown sparse target matrix S° , and then U_0 and V_0 are generated via an SVD of $M - S_0$ such that $U_0 V_0^\top$ forms a rough initial estimate to the unknown low-rank target matrix L° . In the second phase, the algorithm alternately performs two operations: taking gradient steps for U and V , and computing a sparse estimator to adjust S . The sparse estimator is to guarantee that the fraction of nonzero entries in each column and row of S is bounded above so that the nonzero entries are spread out in S . The authors showed that the proposed two-phase algorithm recovers the target decomposition and linear convergence is achieved with proper initialization and step size, under the incoherence assumptions similar to (8). For more detailed description of the assumptions and the results, see [11].

In [29], assuming that the data matrix M is observed indirectly through compressive measurements, Gu, Wang and Liu considered the following variant of RPCA:

$$\min H(U, V, S) := \|A(UV^\top + S) - M\|_F^2 \quad \text{s.t.} \quad \|S\|_0 \leq \tau_s, \quad (37)$$

where A is a sensing matrix. The alternating minimization algorithm proposed in [29] iterates as follows:

$$\begin{aligned} U^{k+1} &:= \operatorname{argmin}_U H(U, V^k, S^k) \\ V^{k+1} &:= \operatorname{argmin}_V H(U^{k+1}, V, S^k) \\ S^{k+1} &:= \operatorname{argmin}_S H(U^{k+1}, V^{k+1}, S) \quad \text{s.t.} \quad \|S\|_0 \leq \tau_s. \end{aligned} \quad (38)$$

It is noted that the U and V subproblems in (38) correspond to solving linear systems and the S -subproblem admits an easily computable closed-form solution. The authors showed that under incoherence assumption on L° and A satisfying restricted isometry property (RIP), (38) converges globally. However, note that $A = I$ does not satisfy the RIP condition, and therefore the convergence is not guaranteed for RPCA problem in (1) for which $A = I$. Similar idea was also investigated in [62], assuming the RIP condition on the sensing matrix A , and thus does not apply to the RPCA problem either.

In [30], the authors considered the scenario such that the columns of the data matrix M are observed in an online fashion. This is suitable for many real applications, e.g., in surveillance video background separation. To handle this problem, the authors proposed a stochastic algorithm, which solves a nonconvex variant of RPCA:

$$\min_{U, V, S} \frac{1}{2} \|UV^\top + S - M\|_F^2 + \frac{\rho_1}{2} (\|U\|_F^2 + \|V\|_F^2) + \rho_2 \|S\|_1, \quad (39)$$

where ρ_1 and ρ_2 are some weight parameters. The formulation in (39) exploits the representation of the nuclear norm est-

blished in [8]. In particular, for any given $L \in \mathbb{R}^{m \times n}$ such that $\text{rank}(L) \leq r$, $\|L\|_*$ can be computed as follows:

$$\|L\|_* := \inf_{U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}} \left\{ \frac{1}{2} \|U\|_F^2 + \frac{1}{2} \|V\|_F^2 : UV^\top = L \right\}. \quad (40)$$

Form (40) we know that (4) is equivalent to

$$\min_{U, V, S} \frac{\mu}{2} \|UV^\top + S - M\|_F^2 + \frac{1}{2} (\|U\|_F^2 + \|V\|_F^2) + \rho \|S\|_1. \quad (41)$$

As a result, (39) is a nonconvex reformulation of the penalized variant of RPCA in (4). For given matrix $M = [M_1, \dots, M_n] \in \mathbb{R}^{m \times n}$, solving (39) is the same as the following empirical risk minimization problem:

$$\min_U \frac{1}{n} \sum_{i=1}^n \ell(M_i, U) + \rho_1 \|U\|_F^2, \quad (42)$$

where $\ell(M_i, U)$ is defined as

$$\ell(M_i, U) := \min_{V_i \in \mathbb{R}^r, S_i \in \mathbb{R}^m} \frac{1}{2} \|UV_i + S_i - M_i\|_2^2 + \frac{\rho_1}{2} \|V_i\|_2^2 + \rho_2 \|S_i\|_1. \quad (43)$$

The empirical risk minimization (42) favors stochastic gradient descent algorithm. Of course every time to compute the gradient of $\ell(M_i, U)$, another minimization problem in (43) needs to be solved. Therefore, the algorithm proposed in [30] is an alternating minimization method with subproblem for U being solved using stochastic gradient descent. The authors showed that the proposed method converges to the correct low-dimensional subspace asymptotically under certain assumptions.

The following nonconvex variant of RPCA was proposed by Shen, Wen and Zhang in [31]:

$$\min_{U, V} \|UV^\top - M\|_1. \quad (44)$$

This simple reformulation can be viewed as a nonconvex reformulation of (2) but without any regularization terms on U and V . In particular, (44) can be rewritten as

$$\min \|S\|_1 \quad \text{s.t.} \quad S + UV^\top = M. \quad (45)$$

The authors in [31] proposed an ADMM algorithm (named LMafit) for solving (45). By associating a Lagrange multiplier Λ to the constraint, the augmented Lagrangian function for (45) can be written as

$$\begin{aligned} \mathcal{L}_\beta(U, V, S; \Lambda) := & \|S\|_1 - \langle \Lambda, UV^\top + S - M \rangle \\ & + \frac{\beta}{2} \|UV^\top + S - M\|_F^2, \end{aligned}$$

where $\beta > 0$ is a penalty parameter. The nonconvex ADMM for solving (45) iterates the updates as follows:

$$\begin{aligned} U^{k+1} & := \arg\min_U \mathcal{L}_\beta(U, V^k, S^k; \Lambda^k) \\ V^{k+1} & := \arg\min_V \mathcal{L}_\beta(U^{k+1}, V, S^k; \Lambda^k) \\ S^{k+1} & := \arg\min_S \mathcal{L}_\beta(U^{k+1}, V^{k+1}, S; \Lambda^k) \\ \Lambda^{k+1} & := \Lambda^k - \beta(U^{k+1}V^{k+1\top} + S^{k+1} - M). \end{aligned} \quad (46)$$

Note that all three subproblems in (46) are easy to solve. In particular, the U and V -subproblems correspond to solving linear systems, and the S -subproblem corresponds to the

soft-shrinkage operation of the ℓ_1 norm (11). However, this nonconvex ADMM lacks convergence guarantees.

In [32], the authors studied some variants of the conditional gradient method and ADMM for solving nonconvex and nonsmooth optimization problems. Consider a general nonconvex optimization problem:

$$\min_x f(x) + r(x) \quad \text{s.t.} \quad x \in \mathcal{X}, \quad (47)$$

where f is smooth and possibly nonconvex, $r(x)$ is convex and nonsmooth, \mathcal{X} is a convex compact set. Moreover, f satisfies the so-called Hölder condition:

$$f(y) \leq f(x) + \nabla f(x)^\top (y - x) + \frac{\gamma}{2} \|y - x\|_p^p, \quad \forall x, y \in \mathcal{X},$$

where $p > 1$ and $\gamma > 0$. The definition of ϵ -stationary solution given in [32] is as follows.

Definition 1: $x \in \mathcal{X}$ is called an ϵ -stationary solution ($\epsilon \geq 0$) for (47) if the following holds:

$$\Psi_{\mathcal{X}}(x) := \inf_{y \in \mathcal{X}} \{ \nabla f(x)^\top (y - x) + r(y) - r(x) \} \geq -\epsilon.$$

The authors commented that this definition is stronger than the one used by Ghadimi et al. in [63].

Now we briefly discuss how to apply the algorithms analyzed in [32] to solve nonconvex RPCA variants. Consider the nonconvex RPCA variant given in (39), the generalized conditional gradient method proposed in [32] can be customized to solve (39). Since the generalized conditional gradient method requires a compact constraint set, one can equivalently reformulate (39) in the following form:

$$\begin{aligned} \min_{U, V, S} & f(U, V, S) + \rho_2 \|S\|_1, \\ \text{s.t.} & \max\{\|U\|_F, \|V\|_F\} \leq \frac{1}{\sqrt{\rho_1}} \|M\|_F, \\ & \|S\|_1 \leq \frac{1}{2\rho_2} \|M\|_F^2, \end{aligned} \quad (48)$$

where $f(U, V, S) := \frac{1}{2} \|UV^\top + S - M\|_F^2 + \frac{\rho_1}{2} (\|U\|_F^2 + \|V\|_F^2)$ denotes the smooth part of the objective function. It is easy to see that ∇f is Lipschitz continuous – let $\gamma > 0$ denote the Lipschitz constant.

At the k -th iteration of the generalized conditional gradient method [32], implemented on (48), one needs to solve the following subproblem:

$$\begin{aligned} \min_{U, V, S} & \langle \nabla_U f(U^k, V^k, S^k), U \rangle + \langle \nabla_V f(U^k, V^k, S^k), V \rangle \\ & + \langle \nabla_S f(U^k, V^k, S^k), S \rangle + \rho_2 \|S\|_1, \\ \text{s.t.} & \max\{\|U\|_F, \|V\|_F\} \leq \frac{1}{\sqrt{\rho_1}} \|M\|_F, \\ & \|S\|_1 \leq \frac{1}{2\rho_2} \|M\|_F^2. \end{aligned} \quad (49)$$

Suppose $(\bar{U}^k, \bar{V}^k, \bar{S}^k)$ denotes the solution of (49), a typical iteration of the generalized conditional gradient method is

given as follows:

$$\left\{ \begin{array}{l} \text{Solve (49) to obtain } (\bar{U}^k, \bar{V}^k, \bar{S}^k) \\ \Delta U^k := \bar{U}^k - U^k \\ \Delta V^k := \bar{V}^k - V^k \\ \Delta S^k := \bar{S}^k - S^k \\ \Delta^k := [\Delta U^k, \Delta V^k, \Delta S^k] \\ \alpha_k := \operatorname{argmin}_{\alpha \in [0,1]} \alpha \langle \nabla f(U^k, V^k, S^k), \Delta^k \rangle \\ \quad + \frac{\alpha^2 \gamma}{2} \|\Delta^k\|_F^2 + (1 - \alpha) \rho_2 \| [U^k, V^k, S^k] \|_1 \\ \quad + \alpha \rho_2 \| [\bar{U}^k, \bar{V}^k, \bar{S}^k] \|_1 \\ U^{k+1} := (1 - \alpha_k) U^k + \alpha_k \bar{U}^k \\ V^{k+1} := (1 - \alpha_k) V^k + \alpha_k \bar{V}^k \\ S^{k+1} := (1 - \alpha_k) S^k + \alpha_k \bar{S}^k. \end{array} \right. \quad (50)$$

Note that the generalized conditional gradient method in [32] involves a line search step for computing α_k as shown in (50), which can be efficiently computed. It is shown in [32] that the generalized conditional gradient method in (50) can compute an ϵ -stationary solution of (48) in $O(1/\epsilon^2)$ iterations.

Jiang et al. [32] also proposed some ADMM variants that can solve various nonconvex RPCA formulations, and the authors provided a convergence rate analysis to compute an ϵ -stationary solution – the definition of ϵ -stationarity employed to analyze the ADMM algorithm is given in Definition 2. We here discuss the ADMM-g algorithm in [32] which can solve the following RPCA variant:

$$\begin{aligned} \min \quad & \frac{1}{2} \|L - UV^\top\|_F^2 + \frac{\rho_1}{2} (\|U\|_F^2 + \|V\|_F^2) \\ & + \rho_2 \|S\|_1 + \rho_3 \|N\|_F^2 \\ \text{s.t.} \quad & L + S + N = M. \end{aligned} \quad (51)$$

This is a nonsmooth and nonconvex problem with five block variables L , S , N , U and V , and it can be viewed as a variant of (39) with linear constraints. Treating N as the last block variable, a typical iteration of ADMM-g for solving (51) iterates as follows:

$$\begin{aligned} L^{k+1} &:= \operatorname{argmin}_L \tilde{\mathcal{L}}_{L^k}(L, U^k, V^k, S^k, N^k; \Lambda^k) \\ U^{k+1} &:= \operatorname{argmin}_U \tilde{\mathcal{L}}_{U^k}(L^{k+1}, U, V^k, S^k, N^k; \Lambda^k) \\ V^{k+1} &:= \operatorname{argmin}_V \tilde{\mathcal{L}}_{V^k}(L^{k+1}, U^{k+1}, V, S^k, N^k; \Lambda^k) \\ S^{k+1} &:= \operatorname{argmin}_S \tilde{\mathcal{L}}_{S^k}(L^{k+1}, U^{k+1}, V^{k+1}, S, N^k; \Lambda^k) \\ N^{k+1} &:= N^k - \eta \nabla_N \mathcal{L}(L^{k+1}, U^{k+1}, V^{k+1}, S^{k+1}, N; \Lambda^k) \\ \Lambda^{k+1} &:= \Lambda^k - \beta (L^{k+1} + S^{k+1} + N^{k+1} - M), \end{aligned} \quad (52)$$

where $\eta > 0$ is a step size, the augmented Lagrangian function \mathcal{L} is defined as

$$\begin{aligned} \mathcal{L}(L, U, V, S, N; \Lambda) &:= \\ & \frac{1}{2} \|L - UV^\top\|_F^2 + \frac{\rho_1}{2} (\|U\|_F^2 + \|V\|_F^2) + \rho_2 \|S\|_1 \\ & + \rho_3 \|N\|_F^2 - \langle \Lambda, L + S + N - M \rangle \\ & + \frac{\beta}{2} \|L + S + N - M\|_F^2, \end{aligned}$$

and $\tilde{\mathcal{L}}$ denotes \mathcal{L} plus a proximal term. For example, $\tilde{\mathcal{L}}_{L^k}$ is defined as

$$\begin{aligned} \tilde{\mathcal{L}}_{L^k}(L, U^k, V^k, S^k, N^k; \Lambda^k) &:= \\ & \mathcal{L}(L, U^k, V^k, S^k, N^k; \Lambda^k) + \frac{1}{2} \|L - L^k\|_F^2, \end{aligned}$$

where H denotes a pre-specified positive definite matrix which needs to satisfy certain conditions to guarantee the convergence of the method as stated in (52). It is noted that the last block variable N is treated specially. It is not updated by minimizing the augmented Lagrangian function, but by taking a gradient step on it. The results in [32] indicate that ADMM-g (52) finds an ϵ -stationary solution for (51) in no more than $O(1/\epsilon^2)$ iterations. Since (51) is a constrained problem, the definition of its ϵ -stationary solution is different from the one in Definition 1. Here we briefly discuss how it is defined for constrained problems in [32]. We consider the following constrained nonsmooth and nonconvex problem:

$$\begin{aligned} \min \quad & f(x_1, \dots, x_p) + \sum_{i=1}^{p-1} r_i(x_i) \\ \text{s.t.} \quad & \sum_{i=1}^{p-1} A_i x_i + x_p = b, \\ & x_i \in \mathcal{X}_i, i = 1, \dots, p-1, \end{aligned} \quad (53)$$

where $x_i \in \mathbb{R}^{n_i}$, f is differentiable and possibly nonconvex, each r_i is possibly nonsmooth and nonconvex, and each \mathcal{X}_i is a convex set. Note again that the last block variable x_p is treated differently, which is needed in the analysis of convergence rate. The ϵ -stationary solution to (53) is defined as follows.

Definition 2: $(x_1^*, \dots, x_p^*) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_{p-1} \times \mathbb{R}^{n_p}$ is called an ϵ -stationary solution to (53), if there exists λ^* such that the following holds for any $(x_1, \dots, x_p) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_{p-1} \times \mathbb{R}^{n_p}$:

$$\begin{aligned} (x_i - x_i^*)^\top (g_i^* + \nabla_i f(x_1^*, \dots, x_p^*) - A_i^\top \lambda^*) &\geq -\epsilon, \\ & i = 1, \dots, p-1, \\ \|\nabla_p f(x_1^*, \dots, x_{p-1}^*, x_p^*) - \lambda^*\| &\leq \epsilon, \\ \|\sum_{i=1}^{p-1} A_i x_i^* + x_p^* - b\| &\leq \epsilon, \end{aligned}$$

where g_i^* is a general subgradient of r_i at point x_i^* . This set of inequalities can be viewed as a perturbed KKT system.

It is also interesting to note that (51) is equivalent to the following unconstrained problem, and thus can be solved by block coordinate descent method (BCD).

$$\begin{aligned} \min \quad & \frac{1}{2} \|M - S - N - UV^\top\|_F^2 + \frac{\rho_1}{2} (\|U\|_F^2 + \|V\|_F^2) \\ & + \rho_2 \|S\|_1 + \rho_3 \|N\|_F^2. \end{aligned} \quad (54)$$

Most existing BCD type algorithms for solving nonconvex problems lack convergence rate analysis. In [32], the authors proposed a proximal BCD method that can find an ϵ -stationary solution in $O(1/\epsilon^2)$ iterations. Denoting the objective function in (54) as $F(U, V, S, N)$, the proximal BCD given in [32] for solving (54) updates the variables as follows:

$$\begin{aligned} U^{k+1} &:= \operatorname{argmin}_U F(U, V^k, S^k, N^k) + \frac{1}{2} \|U - U^k\|_H^2 \\ V^{k+1} &:= \operatorname{argmin}_V F(U^{k+1}, V, S^k, N^k) + \frac{1}{2} \|V - V^k\|_H^2 \\ S^{k+1} &:= \operatorname{argmin}_S F(U^{k+1}, V^{k+1}, S, N^k) + \frac{1}{2} \|S - S^k\|_H^2 \\ N^{k+1} &:= \operatorname{argmin}_N F(U^{k+1}, V^{k+1}, S^{k+1}, N) + \frac{1}{2} \|N - N^k\|_H^2, \end{aligned} \quad (55)$$

where H denotes a pre-specified positive definite matrix.

IV. PRELIMINARY NUMERICAL EXPERIMENTS

In this section we provide some elementary numerical results of different algorithms for solving RPCA. We selected eight different solvers, five for solving convex problems: IALM [17], ADM [18], ADMIP [22], fastRPCA-max [23] and fastRPCA-sum [23], and three for solving nonconvex problems: LMafit [31], R2PCP [28], GD [11]. We tested

their performance on some standard synthetic data used in many RPCA papers. The synthetic data were generated in the following manner:

- 1) $L^\circ = UV^\top$ such that $U \in \mathbb{R}^{n \times r}$, $V \in \mathbb{R}^{n \times r}$ for $r = c_r n$ and $c_r \in \{0.05, 0.1\}$. Moreover, $U_{ij} \sim \mathcal{N}(0, 1)$, $V_{ij} \sim \mathcal{N}(0, 1)$ for all i, j are independent standard Gaussian variables,
- 2) $\Omega \subset \{(i, j) : 1 \leq i, j \leq n\}$ was chosen uniformly at random such that its cardinality $|\Omega| = c_p n^2$ and $c_p \in \{0.05, 0.1\}$,
- 3) $S_{ij}^\circ \sim U[-\sqrt{8r/\pi}, \sqrt{8r/\pi}]$ for all $(i, j) \in \Omega$ are independent uniform random variables,
- 4) $N_{ij}^\circ \sim \varrho \mathcal{N}(0, 1)$ for all i, j are independent Gaussian variables, where for given signal-to-noise ratio (SNR) of M , ϱ is computed from

$$\begin{aligned} SNR(M) &= 10 \log_{10} \left(\frac{\mathbb{E}[\|L^\circ + S^\circ\|_F^2]}{\mathbb{E}[\|N^\circ\|_F^2]} \right) \\ &= 10 \log_{10} \left(\frac{c_r n + c_s \frac{8r}{3\pi}}{\varrho^2} \right), \end{aligned}$$

and $SNR(M) \in \{50dB, 100dB\}$,

- 5) The data matrix $M = L^\circ + S^\circ + N^\circ$.

Note that the non-zero entries of the sparse component and the entries of the low-rank component have approximately the same magnitude in expectation. Indeed, for $n \gg 1$, $L_{ij}^\circ \approx \sqrt{r} \mathcal{N}(0, 1)$; hence, $\mathbb{E}[L_{ij}^\circ] = \sqrt{\frac{2r}{\pi}}$ for all i, j and $\mathbb{E}[S_{ij}^\circ] = \sqrt{\frac{2r}{\pi}}$ for $(i, j) \in \Omega$.

We created 10 random problems of size $n \in \{500, 1500\}$, i.e., $M \in \mathbb{R}^{n \times n}$, for each of the two choices of $SNR(M)$, c_r and c_p using the procedure described above. We plot the figures showing the averaged relative errors of the iterates over 10 runs versus cpu times in Figure 1, where the relative error of (L, S) is defined as

$$\text{Error}(L, S) := \frac{\|L - L^\circ\|_F}{\|L^\circ\|_F} + \frac{\|S - S^\circ\|_F}{\|S^\circ\|_F}.$$

For all the eight algorithms, we used their default stopping criteria and default parameters setting if the output is of good quality; otherwise, we tuned some parameters so that the algorithm becomes competitive for our experimental setting. It is worth emphasizing that fastRPCA-sum and fastRPCA-max solve (27) with $\psi = \psi_{\text{sum}}$ and $\psi = \psi_{\text{max}}$, respectively; and ρ_{max} , the trade-off parameter for ψ_{max} , is set to $\|L^\circ\|_* / \|S^\circ\|_1$, i.e., this model needs an oracle that provides an ideal separation to tune the parameter.

Remark 4.1: We remark that comparing different algorithms for solving RPCA is not an easy task for the following reasons. (i) The algorithms are designed for solving related but different formulations. For example, IALM solves (2) and LMfit solves (45), so it is difficult to compare which algorithm is better. (ii) The performance of all the algorithms depends on the parameter settings, initial points, and data structures. For example, from Figure 1 we see that LMfit consistently outperforms fastRPCA-sum. However, this is only based on the data and parameters we tested. For other data sets, initial points, and parameter settings, fastRPCA-sum may be better than LMfit.

V. FUTURE DIRECTIONS

Although (2) is a convex problem with nice statistical properties, $\|\cdot\|_*$ causes problem for large-scale problems as standard algorithms for (2) cannot easily exploit multi-core or multi-server computing environments as (2) is not amenable to distributed optimization. Some future work in this direction are related to the non-convex formulation which uses an equivalent representation of $\|\cdot\|_*$ by Recht et al. [8] as given in (40). For instance, given some $\nu > 0$, consider $g_\nu : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ defined in (15). Given a data matrix $M \in \mathbb{R}^{m \times n}$, using g_ν , we can formulate a smooth non-convex optimization problem:

$$\min_{U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}} \Psi_\nu(U, V) \triangleq \frac{1}{2} \|U\|_F^2 + \frac{1}{2} \|V\|_F^2 + \rho g_\nu(UV^\top - M), \quad (56)$$

where $\rho > 0$ and $\mathbb{Z}_+ \ni r \geq \text{rank}(L^\circ)$ are given parameters. Here, one can use PALM algorithm [64] to generate a sequence that converges to a critical point of Ψ_ν which is a KL function – also see [65], [32] for some other related work on nonconvex optimization.

Note that instead of solving the smooth approximation given in (56), it is preferable to solve the following nonconvex formulation in (57), which is equivalent to (2).

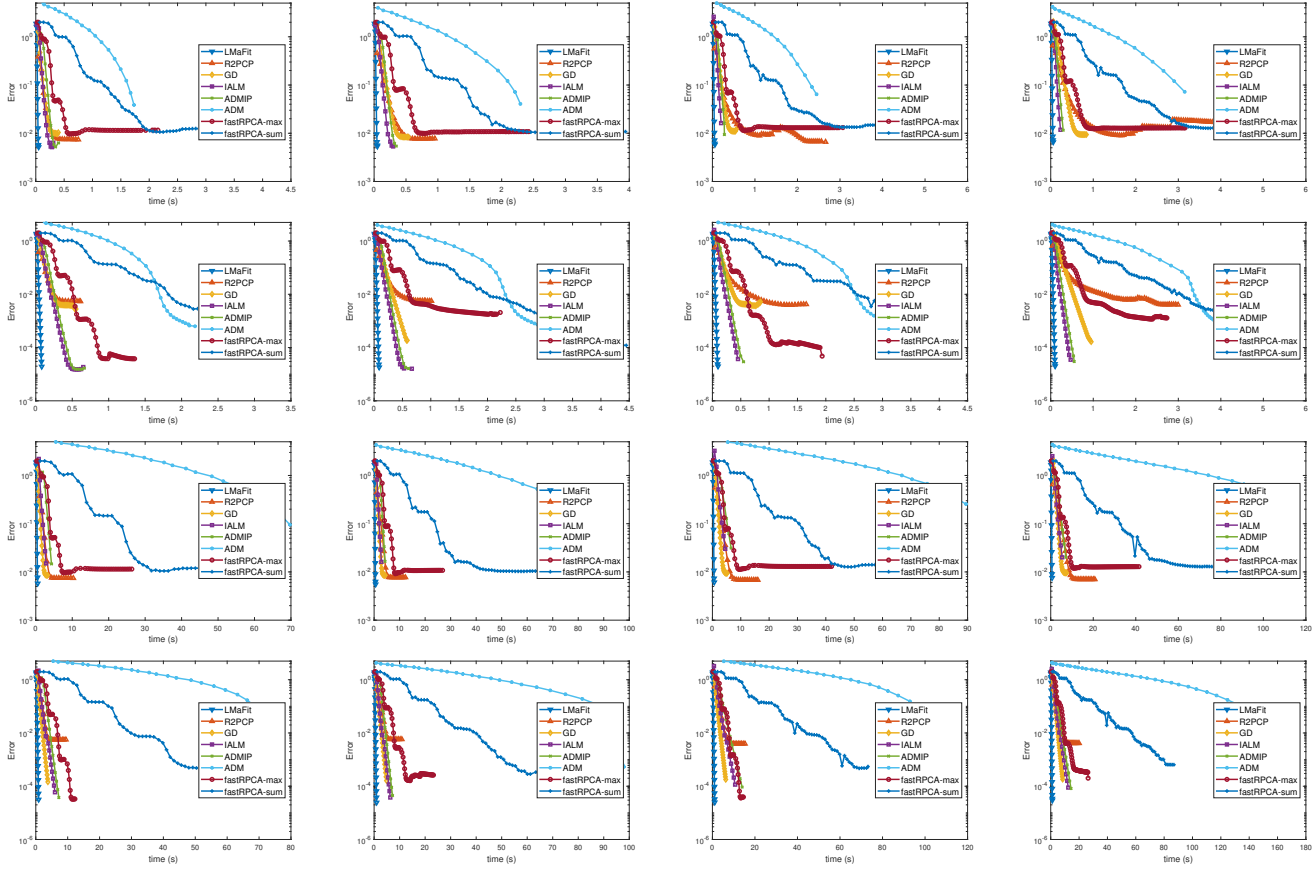
$$\min_{U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}} \frac{1}{2} \|U\|_F^2 + \frac{1}{2} \|V\|_F^2 + \rho \|UV^\top - M\|_1. \quad (57)$$

To the best of authors' knowledge, there does not exist efficient methods with strong convergence guarantees to solve (57).

Note that the third term in (57) is a composite function of the form $g(h(\cdot))$ where g is a nonsmooth convex function such that g is Lipschitz continuous, and h is a differentiable function such that its Jacobian h' is Lipschitz. Thus, one possible direction is to design trust region algorithm for (57) – See Section 7.7 in [66]. As an alternative to trust-region algorithm, one might also consider the augmented Lagrangian (AL) method. It is known that for constrained non-convex problems, provided that the second-order KKT conditions hold, the AL will have a saddle point for penalty parameter chosen sufficiently large; therefore, the duality gap encountered in Lagrangian formulations do not pose a problem for augmented Lagrangian based methods – thus, AL methods might prove useful to establish convergence to local minima when initialized sufficiently close to the local minimum [66], [67].

Due to the close relationship between low-rank matrix completion and RPCA, some algorithms for solving low-rank matrix completion problems can possibly be extended to solve variants of RPCA. New methods based on manifold optimization are recently studied for solving low-rank matrix completion problems, see e.g., [68], [69], [70], [71], [72]. It is noted that all these works consider a matrix completion variant/reformulation which is a manifold optimization problem with a smooth objective function. For example, assume that the matrix M is observed partially, i.e., only entries that in a

Fig. 1. Comparison of different algorithms. The first row shows the comparison results for $n = 500$, $SNR = 50dB$, and $(c_r, c_p) = (0.05, 0.05)$, $(0.05, 0.1)$, $(0.1, 0.05)$, $(0.1, 0.1)$, respectively from left to right. The second row corresponds to $n = 500$, $SNR = 100dB$, the third row corresponds to $n = 1500$, $SNR = 50dB$, and the fourth row corresponds to $n = 1500$, $SNR = 100dB$, all with (c_r, c_p) set the same as the first row from left to right. We remark again that these comparison results only reflect the performance of the algorithms under the current setting of parameters, stopping criteria, initial points, input data and so on. For other settings and data, the performance can be very different. See Remark 4.1.



subset Ω are observed, the low-rank matrix completion model considered in [70] is:

$$\min_{U \in \mathcal{G}(m,r), V \in \mathbb{R}^{n \times r}} \frac{1}{2} \sum_{(i,j) \in \Omega} C_{ij}^2 ((UV^\top)_{ij} - M_{ij})^2 + \frac{\mu^2}{2} (UV^\top)_{ij}^2, \quad (58)$$

where $\mathcal{G}(m,r)$ denotes Grassmann manifold, C_{ij} denotes some weighting parameter and $\mu > 0$ is a penalty parameter. It is noted that (58) is a manifold optimization problem with a smooth objective function. Many existing algorithms can be used to solve a manifold optimization problem with smooth objective, for example, Riemannian gradient method [73], Riemannian trust-region method [70] and Riemannian conjugate gradient method [72], and so on.

However, it is more challenging to design algorithms for manifold optimization reformulations of RPCA variants. The reason is that RPCA variants always involve nonsmooth terms in the objective. In fact, all RPCA variants we discussed so far involve $\|S\|_1$ in the objective. As a result, any manifold optimization reformulation of RPCA variants will involve the nonsmooth ℓ_1 term $\|S\|_1$ as well, unless one can bear with smoothing it, which can potentially degrade the sparsity of S . Algorithms for solving manifold optimization problem with nonsmooth objective function have been very limited, and

most of them lack convergence guarantees. Nonetheless, some of these algorithms can still be adopted to solve manifold optimization reformulations of RPCA variants, although their efficiency in practice needs further investigations. For instance, references [28], [74] propose optimization methods on matrix manifolds. In particular, in [74], Podosinnikova, Setzer and Hein proposed a new RPCA model by minimizing the trimmed reconstruction error, which reduces to minimizing a nonsmooth function over the Stiefel manifold. The method lacks theoretical convergence guarantees such as convergence to a critical point. That said, the authors of [74] numerically demonstrate that their method exhibits good empirical recovery and it is competitive against other nonconvex formulations and convex optimization based methods.

One simple manifold optimization reformulation of RPCA is given as follows.

$$\min_{U \in \mathcal{S}_t(m,r), V \in \mathbb{R}^{n \times r}} \|UV^\top - M\|_1, \quad (59)$$

where $\mathcal{S}_t(m,r)$ denotes Stiefel manifold. The advantages of (59) are as follows: (i) it does not involve nuclear norm and thus avoids SVD; (ii) the sizes of U and V are $m \times r$ and $n \times r$, respectively, which are much smaller than the size of M when $r \ll \min(m,n)$. One may also note that (59) differs (44) only with the Stiefel manifold constraint. The drawback

of (44) is that its optimal solution (U^*, V^*) is not unique, because $(U^*W, V^*W^{-\top})$ is also optimal for any invertible matrix $W \in \mathbb{R}^{r \times r}$. This drawback is fixed nicely in (59). There are several ways to solve (59). For example, one can reformulate (59) as the following one and then apply ADMM to solve it.

$$\begin{aligned} \min \quad & \|S\|_1 \\ \text{s.t.}, \quad & S + UV^\top = M, U \in \mathcal{S}_t(m, r). \end{aligned} \quad (60)$$

The ADMM iterates the updates as follows.

$$\begin{aligned} U^{k+1} &:= \operatorname{argmin}_U \mathcal{L}_\beta(U, V^k, S^k; \Lambda^k), \text{s.t.}, U \in \mathcal{S}_t(m, r) \\ V^{k+1} &:= \operatorname{argmin}_V \mathcal{L}_\beta(U^{k+1}, V, S^k; \Lambda^k) \\ S^{k+1} &:= \operatorname{argmin}_S \mathcal{L}_\beta(U^{k+1}, V^{k+1}, S; \Lambda^k) \\ \Lambda^{k+1} &:= \Lambda^k - \beta(U^{k+1}V^{k+1\top} + S^{k+1} - M), \end{aligned} \quad (61)$$

where the augmented Lagrangian function is defined as

$$\mathcal{L}_\beta(U, V, S; \Lambda) := \|S\|_1 - \langle UV^\top + S - M \rangle + \frac{\beta}{2} \|UV^\top + S - M\|_F^2.$$

The U -subproblem in (61) is a smooth manifold optimization problem and can be solved by existing methods [73]. This method should be very efficient, but the main issue is that under what kind of conditions it is guaranteed to converge.

Zhang, Ma and Zhang studied some ADMM variants for Riemannian manifold optimization in [75], which can be used to solve manifold optimization reformulations of some RPCA variants. We here briefly discuss this work. We consider the following RPCA variant, which minimizes a nonsmooth function over Stiefel manifold.

$$\begin{aligned} \min \quad & \frac{1}{2} \|L - UV^\top\|_F^2 + \rho \|S\|_1 + \frac{\mu}{2} \|N\|_F^2 \\ \text{s.t.} \quad & \bar{L} + S + N = M \\ & U \in \mathcal{S}_t(m, r), \end{aligned} \quad (62)$$

where $\rho > 0$, $\mu > 0$ are tradeoff parameters. One of the ADMM variants for solving (62) proposed in [75] iterates the updates as follows.

$$\begin{aligned} L^{k+1} &:= \operatorname{argmin}_L \tilde{\mathcal{L}}_{L^k}(L, U^k, V^k, S^k, N^k; \Lambda^k), \\ &\quad \text{s.t.}, U \in \mathcal{S}_t(m, r) \\ U^{k+1} &:= \operatorname{argmin}_U \tilde{\mathcal{L}}_{U^k}(L^{k+1}, U, V^k, S^k, N^k; \Lambda^k) \\ V^{k+1} &:= \operatorname{argmin}_V \tilde{\mathcal{L}}_{V^k}(L^{k+1}, U^{k+1}, V, S^k, N^k; \Lambda^k) \\ S^{k+1} &:= \operatorname{argmin}_S \tilde{\mathcal{L}}_{S^k}(L^{k+1}, U^{k+1}, V^{k+1}, S, N^k; \Lambda^k) \\ N^{k+1} &:= N^k - \eta \nabla_N \mathcal{L}(L^{k+1}, U^{k+1}, V^{k+1}, S^{k+1}, N; \Lambda^k) \\ \Lambda^{k+1} &:= \Lambda^k - \beta(L^{k+1} + S^{k+1} + N^{k+1} - M), \end{aligned} \quad (63)$$

where $\eta > 0$ is a step size, the augmented Lagrangian function \mathcal{L} is defined as

$$\begin{aligned} \mathcal{L}(L, U, V, S, N; \Lambda) &:= \\ & \frac{1}{2} \|L - UV^\top\|_F^2 + \rho_1 \|S\|_1 \\ & + \frac{\mu}{2} \|N\|_F^2 - \langle \Lambda, L + S + N - M \rangle \\ & + \frac{\beta}{2} \|L + S + N - M\|_F^2, \end{aligned}$$

and $\tilde{\mathcal{L}}$ denotes \mathcal{L} plus a proximal term. For example, $\tilde{\mathcal{L}}_{L^k}$ is defined as

$$\begin{aligned} \tilde{\mathcal{L}}_{L^k}(L, U^k, V^k, S^k, N^k; \Lambda^k) &:= \\ & \mathcal{L}(L, U^k, V^k, S^k, N^k; \Lambda^k) + \frac{1}{2} \|L - L^k\|_H^2, \end{aligned}$$

where H denotes a pre-specified positive definite matrix which needs to satisfy certain conditions to guarantee the convergence of (63). Zhang, Ma and Zhang showed in [75] that the algorithm described in (63) finds an ϵ -stationary solution to (62) in no more than $O(1/\epsilon^2)$ iterations under certain conditions on β , η and H .

One thing that we need to note is that the term $\frac{\mu}{2} \|N\|_F^2$ in (62) is very crucial here. Without this squared term, the convergence results in [75] do not apply. For example, if one considers the following RPCA variant without the noisy term N ,

$$\begin{aligned} \min \quad & \frac{1}{2} \|L - UV^\top\|_F^2 + \rho \|S\|_1 \\ \text{s.t.} \quad & \bar{L} + S = M \\ & U \in \mathcal{S}_t(m, r), \end{aligned} \quad (64)$$

then the ADMM variants proposed in [75] are not guaranteed to converge if they are applied to solve (64). How to extend and generalize the results in [75] so that they can be used to solve other manifold optimization reformulations of RPCA variants definitely deserves more investigations.

It is known that the nuclear norm minimization problem can be equivalently written as an SDP – see [8]. Though the problem size of the SDP is larger than the original nuclear norm minimization problem, it is recently found that the facial reduction technique [76] can reduce the size of the SDP significantly. This idea has been explored in low-rank matrix completion [77] and RPCA [78]. In particular, in [78] the authors showed that RPCA with partial observation

$$\min \mathbf{rank}(L) + \rho \|S\|_0, \text{ s.t.}, P_\Omega(L + S - M) = 0, \quad (65)$$

is equivalent to another optimization problem with semidefinite constraint. By applying the facial reduction technique to the semidefinite cone, the size of this reformulation can be significantly reduced, and then it can be solved very efficiently to high accuracy. Extending the facial reduction technique to other RPCA variants is an interesting topic for future research.

VI. CONCLUSIONS

In this paper, we gave a comprehensive review on algorithms for solving relaxations and variants of robust PCA. Algorithms for solving convex and nonconvex models were discussed. We elaborated in details on the applicability of the algorithms and their convergence behaviors. We also proposed several new directions in the hope that they may shed some light for future research in this area.

ACKNOWLEDGMENTS

The authors are grateful to three anonymous referees for their insightful and constructive comments that led to an improved version of the paper. The work of S. Ma was supported in part by a startup package from Department of Mathematics at UC Davis. The work of N. S. Aybat was supported in part by NSF under Grant CMMI-1400217 and Grant CMMI-1635106, and ARO grant W911NF-17-1-0298.

REFERENCES

- [1] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of ACM*, vol. 58, no. 1, pp. 1–37, 2009.
- [2] Y. Peng, A. Ganesh, J. Wright, W. Xu, and Y. Ma, "Rasl: Robust alignment by sparse and low-rank decomposition for linearly correlated images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 11, pp. 2233–2246, 2012.
- [3] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma, "Robust recovery of subspace structures by low-rank representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 171–184, 2013.
- [4] N. Shahid, V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst, "Robust principal component analysis on graphs," in *ICCV*, 2015.
- [5] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky, "Rank-sparsity incoherence for matrix decomposition," *SIAM Journal on Optimization*, vol. 21, no. 2, pp. 572–596, 2011.
- [6] J. Wright, A. Ganesh, S. Rao, Y. Peng, and Y. Ma, "Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization," in *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, Eds. Curran Associates, Inc., 2009, pp. 2080–2088. [Online]. Available: <http://papers.nips.cc/paper/3704-robust-principal-component-analysis-exact-recovery-of-corrupted-low-rank-matrices-via-convex-optimization.pdf>
- [7] H. Xu, C. Caramanis, and S. Sanghavi, "Robust PCA via outlier pursuit," in *Advances in Neural Information Processing Systems 23*, J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds. Curran Associates, Inc., 2010, pp. 2496–2504. [Online]. Available: <http://papers.nips.cc/paper/4005-robust-pca-via-outlier-pursuit.pdf>
- [8] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM review*, vol. 52, no. 3, pp. 471–501, 2010.
- [9] Z. Zhou, J. Wright, X. Li, E. J. Cands, and Y. Ma, "Stable principal component pursuit," in *Proceedings of International Symposium on Information Theory*, 2010.
- [10] H. Zhang, Y. Zhou, and Y. Liang, "Analysis of robust PCA via local incoherence," 2015.
- [11] X. Yi, D. Park, Y. Chen, and C. Caramanis, "Fast algorithms for robust PCA via gradient descent," 2016.
- [12] T. Zhang and Y. Yang, "Robust PCA by manifold optimization," <https://arxiv.org/abs/1708.00257>, 2017.
- [13] P. Netrapalli, U. N. Niranjan, S. Sanghavi, A. Anandkumar, and P. Jain, "Non-convex robust PCA," 2014.
- [14] T. Bouwmans and E. H. Zahzah, "Robust PCA via principal component pursuit: A review for a comparative evaluation in video surveillance," *Computer Vision and Image Understanding*, vol. 122, pp. 22–34, 2014.
- [15] N. S. Aybat, *Handbook of Robust Low Rank and Sparse Matrix Decomposition: Applications in Image and Video Processing*. CRC Press, Taylor and Francis Group, 2016, ch. Algorithms for Stable PCA.
- [16] Z. Lin, A. Ganesh, J. Wright, L. Wu, M. Chen, and Y. Ma, "Fast convex optimization algorithms for exact recovery of a corrupted low-rank matrix," in *International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, 2009.
- [17] Z. Lin, M. Chen, L. Wu, and Y. Ma, "The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices," UIUC Technical Report UILU-ENG-09-2215, Tech. Rep., 2009.
- [18] X. Yuan and J. Yang, "Sparse and low-rank matrix decomposition via alternating direction methods," *Pacific Journal of Optimization*, vol. 9, no. 1, pp. 167–180, 2013.
- [19] D. Goldfarb, S. Ma, and K. Scheinberg, "Fast alternating linearization methods for minimizing the sum of two convex functions," *Mathematical Programming Series A*, vol. 141, no. 1-2, pp. 349–382, 2013.
- [20] M. Tao and X. Yuan, "Recovering low-rank and sparse components of matrices from incomplete and noisy observations," *SIAM Journal on Optimization*, vol. 21, no. 1, pp. 57–81, 2011.
- [21] N. S. Aybat, D. Goldfarb, and S. Ma, "Efficient algorithms for robust and stable principal component pursuit problems," *Computational Optimization and Applications*, vol. 58, pp. 1–29, 2014.
- [22] N. S. Aybat and G. Iyengar, "An alternating direction method with increasing penalty for stable principal component pursuit," *Computational Optimization and Applications*, vol. 61, pp. 635–668, 2015.
- [23] A. Aravkin, S. Becker, V. Cevher, and P. Olsen, "A variational approach to stable principal component pursuit," in *30th Conference on Uncertainty in Artificial Intelligence (UAI) 2014*, no. EPFL-CONF-199542, 2014, preprint made available at arXiv:1406.1089 [math.OC] in June, 2014.
- [24] T. Lin, S. Ma, and S. Zhang, "Global convergence of unmodified 3-block ADMM for a class of convex minimization problems," *Journal of Scientific Computing*, 2017.
- [25] C. Mu, Y. Zhang, J. Wright, and D. Goldfarb, "Scalable robust matrix recovery: Frank-wolfe meets proximal methods," *SIAM Journal on Scientific Computing*, vol. 38, no. 5, pp. A3291–A3317, 2016.
- [26] T. Zhou and D. Tao, "Godec: Randomized low-rank and sparse matrix decomposition in noisy case," in *ICML*, 2011.
- [27] —, "Greedy bilateral sketch, completion & smoothing," in *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, C. M. Carvalho and P. Ravikumar, Eds., vol. 31. Scottsdale, Arizona, USA: PMLR, 29 Apr–01 May 2013, pp. 650–658.
- [28] M. Hintermüller and T. Wu, "Robust principal component pursuit via inexact alternating minimization on matrix manifolds," *Journal of Mathematical Imaging and Vision*, vol. 51, no. 3, pp. 361–377, 2015.
- [29] Q. Gu, Z. Wang, and H. Liu, "Low-rank and sparse structure pursuit via alternating minimization," in *AISTATS*, 2016.
- [30] J. Feng, H. Xu, and S. Yan, "Online robust PCA via stochastic optimization," in *NIPS*, 2013.
- [31] Y. Shen, Z. Wen, and Y. Zhang, "Augmented lagrangian alternating direction method for matrix separation based on low-rank factorization," *Optimization Methods and Software*, vol. 29, no. 2, pp. 239–263, 2014.
- [32] B. Jiang, T. Lin, S. Ma, and S. Zhang, "Structured nonconvex and nonsmooth optimization: Algorithms and iteration complexity analysis." <https://arxiv.org/abs/1605.02408>, 2016.
- [33] Y. E. Nesterov, "Gradient methods for minimizing composite functions," *Mathematical Programming*, vol. 140, no. 1, pp. 125–161, 2013.
- [34] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [35] P. Tseng, "On accelerated proximal gradient methods for convex-concave optimization," *Manuscript*, 2008.
- [36] P. L. Combettes and J.-C. Pesquet, "A douglas-rachford splitting approach to nonsmooth convex variational signal recovery," *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 4, pp. 564–574, 2007.
- [37] T. Goldstein and S. Osher, "The split Bregman method for l1-regularized problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 2, pp. 323–343, 2009.
- [38] J. Yang, Y. Zhang, and W. Yin, "A fast alternating direction method for tvl1-l2 signal reconstruction from partial fourier data," *IEEE Journal of Selected Topics in Signal Processing Special Issue on Compressed Sensing*, vol. 4, no. 2, pp. 288–297, 2010.
- [39] J. Yang and Y. Zhang, "Alternating direction algorithms for l1-problems in compressive sensing," *SIAM Journal on Scientific Computing*, vol. 33, no. 1, pp. 250–278, 2011.
- [40] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite-element approximations," *Comp. Math. Appl.*, vol. 2, pp. 17–40, 1976.
- [41] D. Gabay, "Applications of the method of multipliers to variational inequalities," in *Augmented Lagrangian Methods: Applications to the Solution of Boundary Value Problems*, M. Fortin and R. Glowinski, Eds. Amsterdam: North-Holland, 1983.
- [42] M. Fortin and R. Glowinski, *Augmented Lagrangian methods: applications to the numerical solution of boundary-value problems*. North-Holland Pub. Co., 1983.
- [43] P. L. Lions and B. Mercier, "Splitting algorithms for the sum of two nonlinear operators," *SIAM Journal on Numerical Analysis*, vol. 16, pp. 964–979, 1979.
- [44] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, pp. 293–318, 1992.
- [45] B. He and X. Yuan, "On the $\mathcal{O}(1/n)$ convergence rate of douglas-rachford alternating direction method," *SIAM Journal on Numerical Analysis*, vol. 50, pp. 700–709, 2012.
- [46] R. D. C. Monteiro and B. F. Svaiter, "Iteration-complexity of block-decomposition algorithms and the alternating direction method of multipliers," *SIAM Journal on Optimization*, vol. 23, pp. 475–507, 2013.
- [47] T. Lin, S. Ma, and S. Zhang, "On the sublinear convergence rate of multi-block ADMM," *Journal of the Operations Research Society of China*, vol. 3, no. 3, pp. 251–274, 2015.
- [48] C. Chen, B. He, Y. Ye, and X. Yuan, "The direct extension of ADMM for multi-block convex minimization problems is not necessarily convergent," *Mathematical Programming*, vol. 155, pp. 57–79, 2016.

- [49] N. S. Aybat, D. Goldfarb, and G. Iyengar, "Fast first-order methods for stable principal component pursuit," *arXiv preprint arXiv:1105.2126*, 2011.
- [50] R. M. Larsen, "Lanczos bidiagonalization with partial reorthogonalization," *DAIMI Report Series*, vol. 27, no. 537, 1998.
- [51] A. Y. Aravkin, J. V. Burke, and M. P. Friedlander, "Variational properties of value functions," *SIAM Journal on optimization*, vol. 23, no. 3, pp. 1689–1717, 2013.
- [52] M. Frank and P. Wolfe, "An algorithm for quadratic programming," *Naval Res. Logis. Quart.*, vol. 3, pp. 95–110, 1956.
- [53] M. Jaggi, "Revisiting frank-wolfe: Projection-free sparse convex optimization," in *ICML*, 2013.
- [54] —, "Sparse convex optimization methods for machine learning," Ph.D. dissertation, ETH Zurich, Oct. 2011.
- [55] X. Li, M. K. Ng, and X. Yuan, "Median filtering-based methods for static background extraction from surveillance video," *Numerical Linear Algebra with Applications*, vol. 22, pp. 845–865, 2015.
- [56] L. Li, P. Wang, Q. Hu, and S. Cai, "Efficient background modeling based on sparse representation and outlier iterative removal," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 26, no. 2, pp. 278–289, 2014.
- [57] S. Ma, D. Johnson, C. Ashby, D. Xiong, C. L. Cramer, J. H. Moore, S. Zhang, and X. Huang, "SPARCoC: a new framework for molecular pattern discovery and cancer gene identification," *PLoS ONE*, vol. 10, no. 3, p. e0117135, 2015.
- [58] L. Yang, T. K. Pong, and X. Chen, "Alternating direction method of multipliers for a class of nonconvex and nonsmooth problems with applications to background/foreground extraction," *SIAM J. Imaging Sciences*, vol. 10, pp. 74–110, 2017.
- [59] K. Kurdyka, "On gradients of functions definable in o-minimal structures," *Annales de l'institut Fourier*, vol. 146, pp. 769–783, 1998.
- [60] S. Lojasiewicz, *Une propriété topologique des sous-ensembles analytiques réels, Les Équations aux Dérivées Partielles*. Paris: Éditions du centre National de la Recherche Scientifique, 1963.
- [61] A. S. Lewis and J. Malick, "Alternating projections on manifolds," *Mathematics of Operations Research*, vol. 33, no. 1, pp. 216–234, 2008.
- [62] A. E. Waters, A. C. Sankaranarayanan, and R. Baraniuk, "Spars: Recovering lowrank and sparse matrices from compressive measurements," in *NIPS*, 2011.
- [63] S. Ghadimi, G. Lan, and H. Zhang, "Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization," *Math. Program.*, vol. 155, no. 1, pp. 267–305, 2016.
- [64] J. Bolte, S. Sabach, and M. Teboulle, "Proximal alternating linearized minimization for nonconvex and nonsmooth problems," *Mathematical Programming*, vol. 146, no. 1-2, pp. 459–494, 2014.
- [65] P. P. Markopoulos, G. N. Karystinos, and D. A. Pados, "Optimal algorithms for L1-subspace signal processing," *IEEE Transactions on Signal Processing*, vol. 62, pp. 5046–5058, 2014.
- [66] A. P. Ruszczyński, *Nonlinear optimization*. Princeton university press, 2006, vol. 13.
- [67] G. Cohen and D. L. Zhu, "Decomposition coordination methods in large scale optimization problems. the nondifferentiable case and the use of augmented lagrangians," *Advances in large scale systems*, vol. 1, pp. 203–266, 1984.
- [68] R. H. Keshavan and S. Oh, "Optspace: A gradient descent algorithm on the grassman manifold for matrix completion," *Arxiv preprint arXiv:0910.5260v2*, 2009.
- [69] L. Balzano, R. Nowak, and B. Recht, "Online identification and tracking of subspaces from highly incomplete information," in *Allerton Conference*, 2010.
- [70] N. Boumal and P.-A. Absil, "RTRMC: A Riemannian trust-region method for low-rank matrix completion," in *NIPS*, 2011.
- [71] B. Vandereycken, "Low-rank matrix completion by Riemannian optimization," *SIAM J. Optim.*, vol. 23, no. 2, pp. 1214–1236, 2013.
- [72] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung, "Guarantees of riemannian optimization for low rank matrix completion," *preprint https://arxiv.org/abs/1603.06610*, 2016.
- [73] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2008.
- [74] A. Podosinnikova, S. Setzer, and M. Hein, "Robust PCA: Optimization of the robust reconstruction error over the stiefel manifold," *GCPR*, 2014.
- [75] J. Zhang, S. Ma, and S. Zhang, "Multi-block optimization over Riemannian manifolds: an iteration complexity analysis," *https://arxiv.org/abs/1710.02236*, 2017.
- [76] D. Drusvyatskiy and H. Wolkowicz, "The many faces of degeneracy in conic optimization," *Foundations and Trends in Optimization*, vol. 3, no. 2, pp. 77–170, 2016.
- [77] S. Huang and H. Wolkowicz, "Low-rank matrix completion using nuclear norm with facial reduction," *Journal of Global Optimization*, 2017.
- [78] S. Ma, F. Wang, L. Wei, and H. Wolkowicz, "Robust principal component analysis using facial reduction," *http://www.optimization-online.org/DB_HTML/2018/03/6535.html*, 2018.