

Manifold Proximal Point Algorithms for Dual Principal Component Pursuit and Orthogonal Dictionary Learning

Shixiang Chen, Zengde Deng, Shiqian Ma, Anthony Man-Cho So

Abstract—We consider the problem of minimizing the ℓ_1 norm of a linear map over the sphere, which arises in various machine learning applications such as orthogonal dictionary learning (ODL) and robust subspace recovery (RSR). The problem is numerically challenging due to its nonsmooth objective and nonconvex constraint, and its algorithmic aspects have not been well explored. In this paper, we show how the manifold structure of the sphere can be exploited to design fast algorithms with provable guarantees for tackling this problem. Specifically, our contribution is fourfold. First, we present a manifold proximal point algorithm (ManPPA) for the problem and show that it converges at a global sublinear rate. Furthermore, we show that ManPPA can achieve a local quadratic convergence rate when applied to sharp instances of the problem. Second, we develop a semismooth Newton-based inexact augmented Lagrangian method for computing the search direction in each iteration of ManPPA and show that it has an asymptotic superlinear convergence rate. Third, we propose a stochastic variant of ManPPA called StManPPA, which is well suited for large-scale computation, and establish its sublinear convergence rate. Both ManPPA and StManPPA have provably faster convergence rates than existing subgradient-type methods. Fourth, using ManPPA as a building block, we propose a new heuristic method for solving a matrix analog of the problem, in which the sphere is replaced by the Stiefel manifold. The results from our extensive numerical experiments on the ODL and RSR problems demonstrate the efficiency and efficacy of our proposed methods.

I. INTRODUCTION

The problem of finding a subspace that captures the features of a given dataset and possesses certain properties is at the heart of many machine learning applications. One commonly encountered formulation of the problem, which is motivated largely by sparsity or robustness considerations, is given by

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) := \|\mathbf{Y}^\top \mathbf{x}\|_1 \quad \text{s.t.} \quad \|\mathbf{x}\|_2 = 1, \quad (1)$$

A short version of this paper appeared in the 53rd Annual Asilomar Conference on Signals, Systems, and Computers, Nov. 3–6, 2019.

S. Chen is with the Department of Industrial and Systems Engineering at Texas A&M University.

Z. Deng is with Cainiao Network. Part of the research was done when the author was a PhD student at the Department of Systems Engineering and Engineering Management at The Chinese University of Hong Kong.

S. Ma is with the Department of Mathematics at University of California, Davis. Research supported in part by NSF grants DMS-1953210 and CCF-2007797, and UC Davis CeDAR Innovative Data Science Seed Funding Program.

A. M.-C. So is with the Department of Systems Engineering and Engineering Management at The Chinese University of Hong Kong. Research supported in part by the Hong Kong RGC GRF Project CUHK 14203218 and in part by the CUHK Research Sustainability of Major RGC Funding Schemes Project 3133236.

where $\mathbf{Y} \in \mathbb{R}^{n \times p}$ is a given matrix and $\|\cdot\|_r$ denotes the ℓ_r norm of a vector. To better understand how problem (1) arises in applications, let us consider two representative examples.

- **Orthogonal Dictionary Learning (ODL).** The goal of ODL is to find an orthonormal basis that can compactly represent a given set of p ($p \gg n$) data points $\mathbf{y}_1, \dots, \mathbf{y}_p \in \mathbb{R}^n$. Such a problem arises in many signal and image processing applications; see, e.g., [1], [2] and the references therein. By letting $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_p] \in \mathbb{R}^{n \times p}$, the problem can be understood as finding an orthogonal matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ and a sparse matrix $\mathbf{A} \in \mathbb{R}^{n \times p}$ such that $\mathbf{Y} \approx \mathbf{X}\mathbf{A}$. Noting that this means $\mathbf{X}^\top \mathbf{Y} \approx \mathbf{A}$ should be sparse, one approach is to find a collection of sparse vectors from the row space of \mathbf{Y} and apply some post-processing procedure to the collection to form the orthogonal matrix \mathbf{X} . This has been pursued in various works; see, e.g., [3]–[6]. In particular, the work [6] considers the formulation (1) and shows that under a standard generative model of the data, one can recover \mathbf{X} from certain local minimizers of problem (1).
- **Robust Subspace Recovery (RSR).** RSR is a fundamental problem in machine learning and data mining [7]. It is concerned with fitting a linear subspace to a dataset corrupted by outliers. Specifically, given a dataset $\mathbf{Y} = [\mathbf{X}, \mathbf{O}]\Gamma \in \mathbb{R}^{n \times (p_1+p_2)}$, where the columns of $\mathbf{X} \in \mathbb{R}^{n \times p_1}$ are the inlier points spanning a d -dimensional subspace \mathcal{S} of \mathbb{R}^n ($d < p_1$), the columns of $\mathbf{O} \in \mathbb{R}^{n \times p_2}$ are outlier points without a linear structure, and $\Gamma \in \mathbb{R}^{(p_1+p_2) \times (p_1+p_2)}$ is an unknown permutation, the goal is to recover the inlier subspace \mathcal{S} , or equivalently, to cluster the points into inliers and outliers. One recently proposed approach for solving this problem is the so-called dual principal component pursuit (DPCP) [8], [9]. A key task in DPCP is to find a hyperplane that contains all the inliers. Such a task can be tackled by solving problem (1). In fact, it has been shown in [8], [9] that under certain conditions on the inliers and outliers, any global minimizer of problem (1) yields a vector that is orthogonal to the inlier subspace \mathcal{S} .

Despite its attractive theoretical properties in various applications, problem (1) is numerically challenging to solve due to its nonsmooth objective and nonconvex constraint. Nevertheless, the manifold structure of the constraint set $\mathcal{M} := \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\|_2 = 1\}$ suggests that problem (1) could be amenable to manifold optimization techniques [10]. One approach is to apply smoothing to the nonsmooth objective

in (1) and use existing algorithms for Riemannian smooth optimization to solve the resulting problem. For instance, when tackling the ODL problem, Sun et al. [5], [11] and Gilboa et al. [12] proposed to replace the absolute value function $t \mapsto |t|$ by the smooth surrogate $t \mapsto h_\mu(t) = \mu \log(\cosh(t/\mu))$ with $\mu > 0$ being a smoothing parameter, while Qu et al. [13] proposed to replace the ℓ_1 norm with the ℓ_4 norm. They then solve the resulting smoothed problems by either the Riemannian trust-region method [5], [11] or the Riemannian gradient descent method [12], [13]. Although it can be shown that these methods will yield the desired orthonormal basis under a standard generative model of the data, the smoothing approach can introduce significant analytic and computational difficulties [6]. Another approach, which avoids smoothing the objective, is to solve (1) directly using Riemannian nonsmooth optimization techniques. For instance, in the recent work [6], Bai et al. proposed to solve (1) using the Riemannian subgradient method (RSGM), which generates the iterates via

$$\mathbf{x}^{k+1} = \frac{\mathbf{x}^k - \eta_k \mathbf{v}^k}{\|\mathbf{x}^k - \eta_k \mathbf{v}^k\|_2}, \quad \mathbf{v}^k \in \partial_R f(\mathbf{x}^k). \quad (2)$$

Here, $\eta_k > 0$ is the step size; $\partial_R f(\cdot)$ denotes the Riemannian subdifferential of f and is given by

$$\partial_R f(\mathbf{x}) = (\mathbf{I}_n - \mathbf{x}\mathbf{x}^\top) \partial f(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{M},$$

where \mathbf{I}_n is the $n \times n$ identity matrix and $\partial f(\cdot)$ is the usual subdifferential of the convex function f [14, Section 5]. Bai et al. [6] showed that for the ODL problem, RSGM with a suitable initialization will converge at a sublinear rate to a basis vector with high probability under a standard generative model of the data. Moreover, by running RSGM $O(n \log n)$ times, each time with an independent random initialization, one can recover the entire orthonormal basis with high probability. Around the same time, Zhu et al. [9] proposed a projected subgradient method (PSGM) for solving (1). The method generates the iterates via

$$\mathbf{x}^{k+1} = \frac{\mathbf{x}^k - \eta_k \mathbf{v}^k}{\|\mathbf{x}^k - \eta_k \mathbf{v}^k\|_2}, \quad \mathbf{v}^k \in \partial f(\mathbf{x}^k). \quad (3)$$

The updates (2) and (3) differ in the choice of the direction \mathbf{v}^k —the former uses a *Riemannian* subgradient of f at \mathbf{x}^k , while the latter uses a usual *Euclidean* subgradient. For the DPCP formulation of the RSR problem, Zhu et al. [9] showed that under certain assumptions on the data, PSGM with suitable initialization and piecewise geometrically diminishing step sizes will converge at a linear rate to a vector that is orthogonal to the inlier subspace \mathcal{S} . The step sizes take the form $\eta_k = \eta^{\lfloor (k-K_0)/K \rfloor + 1}$, where $\eta \in (0, 1)$ and $K_0, K \geq 1$ satisfy certain conditions. In practice, however, the parameters η, K_0, K are difficult to determine. Therefore, Zhu et al. [9] also proposed a PSGM with modified backtracking line search (PSGM-MBLS), which works well in practice but has no convergence guarantee.

A. Motivations for this Work

Although the results in [6], [9] demonstrate, both theoretically and computationally, the efficacy of RSGM and PSGM for solving instances of (1) that arise from the ODL and RSR

problems, respectively, two fundamental questions remain. First, while PSGM can be shown to achieve a *linear* convergence rate on the DPCP formulation of the RSR problem [9], only a *sublinear* convergence rate has been established for RSGM on the ODL problem [6]. Given the similarity of the updates (2) and (3), it is natural to ask whether the slower convergence rate of RSGM is an artifact of the analysis or due to the inherent structure of the ODL problem. Second, the convergence analyses in [6], [9] focus only on the ODL and RSR problems. In particular, they do not shed light on the performance of RSGM or PSGM when tackling general instances of problem (1). It would be of interest to fill this gap by identifying or developing practically fast methods that have more general convergence guarantees, especially since different applications may give rise to instances of problem (1) with different structures. In a recent attempt to address these questions, Li et al. [15] showed, among other things, that RSGM will converge at the sublinear rate of $\mathcal{O}(k^{-1/4})$ (here, k is the iteration counter) when applied to a general instance of problem (1) and at a linear rate when applied to a so-called *sharp* instance of problem (1). Informally, an optimization problem is said to possess the sharpness property if the objective function grows linearly with the distance to a set of local minima [16]. Such a property plays a crucial role in establishing fast convergence guarantees for a host of iterative methods; see, e.g., [15]–[17] and also [18]–[20] for related results. Since the ODL problem and the DPCP formulation of the RSR problem are known to possess the sharpness property under certain assumptions on the data [6], [9], the results in [15] imply that RSGM will converge linearly on these problems.

B. Our Contributions

In this paper, we depart from the subgradient-type approaches (such as RSGM (2) and PSGM (3)) and present another method called the manifold proximal point algorithm (ManPPA) to tackle problem (1). At each iterate \mathbf{x}^k , ManPPA computes a search direction by minimizing the sum of f and a proximal term defined in terms of the *Euclidean* distance over the tangent space to \mathcal{M} at \mathbf{x}^k . This should be contrasted with other existing PPAs on manifolds (see, e.g., [21], [22]), in which the proximal term is defined in terms of the *Riemannian* distance. Such a difference is important. Indeed, although the search direction defined in ManPPA does not admit a closed-form formula, it can be computed in a highly efficient manner by exploiting the structure of problem (1); see Section II-B. However, the search direction defined in the existing PPAs on manifolds can be as difficult to compute as a solution to the original problem. Consequently, the applicability of those methods is rather limited.

We now summarize our contributions as follows:

- 1) We show that ManPPA has a global sublinear convergence rate of $\mathcal{O}(k^{-1/2})$ when applied to a general instance of problem (1). Moreover, we show that if the instance has the sharpness property, then the local convergence rate of ManPPA is at least quadratic. Although the sublinear rate result follows from the results in [23], the quadratic rate result is new. Moreover, both rates are superior to those of RSGM established in [15]. Key to

the proof of the quadratic rate result is a new *Riemannian subgradient inequality* (see Appendix C, Proposition 5), which extends the classic *subgradient inequality* in the Euclidean space to the sphere \mathcal{M} . Such an inequality allows us to analyze ManPPA in a similar way as its Euclidean counterpart. It can also be of independent interest.

- 2) To compute the search direction in each iteration of ManPPA, we develop a semismooth Newton (SSN)-based inexact augmented Lagrangian method (ALM). Numerically, the proposed method can accurately compute the search direction in a highly efficient manner, which is crucial to the fast convergence of ManPPA. Theoretically, we show, for the first time, that the proposed SSN-based inexact ALM has an asymptotic superlinear convergence rate when finding the search direction.
- 3) We propose a stochastic version of ManPPA called StManPPA to tackle problem (1). StManPPA is well suited for the setting where the number of the data points p is extremely large, as each iteration involves only a simple closed-form update. We also analyze the convergence behavior of StManPPA. In particular, we show that it converges at the sublinear rate of $\mathcal{O}(k^{-1/4})$ when applied to a general instance of problem (1), which matches the convergence rate of RSGM established in [15]. Again, the aforementioned Riemannian subgradient inequality plays an important role in establishing this result, as it connects the analysis of StManPPA to those of various Euclidean stochastic methods.
- 4) Using ManPPA as a building block, we develop a new method for solving the following matrix analog of problem (1):

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times q}} \|\mathbf{Y}^\top \mathbf{X}\|_1 \quad \text{s.t.} \quad \mathbf{X}^\top \mathbf{X} = \mathbf{I}_q. \quad (4)$$

Our interest in problem (4) stems from the observation that it provides alternative formulations of the ODL and RSR problems. Indeed, for the ODL problem, one can recover the entire orthonormal basis all at once by solving problem (4) with $q = n$. For the RSR problem, if one knows the dimension d of the inlier subspace \mathcal{S} , then one can recover it by solving problem (4) with $q = n - d$. We show that a good feasible solution to problem (4) can be found in a column-by-column manner by suitably modifying ManPPA. Although the proposed method is only a heuristic, our extensive numerical experiments show that it yields solutions of comparable quality to but is significantly faster than existing methods on the ODL and RSR problems.

C. Organization and Notation

The rest of the paper is organized as follows. In Section II, we present ManPPA for solving problem (1) and describe a highly efficient method for solving the subproblem that arises in each iteration of ManPPA. We also analyze the convergence behavior of ManPPA. In Section III, we propose StManPPA, a stochastic version of ManPPA that is well suited for large-scale computation, and analyze its convergence behavior. In Section

IV we discuss an extension of ManPPA for solving the matrix analog (4) of problem (1). In Section V, we apply ManPPA to solve the ODL problem and the DPCP formulation of the RSR problem and compare its performance with some existing methods. We draw our conclusions in Section VI.

Besides the notation introduced earlier, we use L to denote the Lipschitz constant of f ; i.e., $|f(\mathbf{x}) - f(\mathbf{y})| \leq L\|\mathbf{x} - \mathbf{y}\|_2$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ (note that $L \leq \sqrt{n} \max_{\mathbf{u} \in \mathbb{R}^n} \|\mathbf{Y}^\top \mathbf{u}\|_1 / \|\mathbf{u}\|_1$). Given a closed set $\mathcal{C} \subseteq \mathbb{R}^n$, we use $\text{Proj}_{\mathcal{C}}(\mathbf{x})$ to denote the projection of \mathbf{x} onto \mathcal{C} and $\text{dist}(\mathbf{x}, \mathcal{C}) := \inf_{\mathbf{y} \in \mathcal{C}} \|\mathbf{y} - \mathbf{x}\|_2$ to denote the distance between \mathbf{x} and \mathcal{C} . Given a proper lower semicontinuous function $h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$, its proximal mapping is given by $\text{prox}_h(\mathbf{x}) = \arg\min_{\mathbf{w} \in \mathbb{R}^n} h(\mathbf{w}) + \frac{1}{2}\|\mathbf{w} - \mathbf{x}\|_2^2$. Given two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we use $\langle \mathbf{x}, \mathbf{y} \rangle$ or $\mathbf{x}^\top \mathbf{y}$ to denote their usual inner product. Other notation is standard.

II. A MANIFOLD PROXIMAL POINT ALGORITHM

Since problem (1) is nonconvex, our goal is to compute a *stationary point* of (1), which is a point $\bar{\mathbf{x}} \in \mathcal{M}$ that satisfies the first-order optimality condition

$$\mathbf{0} \in \partial_R f(\bar{\mathbf{x}}) = (\mathbf{I}_n - \bar{\mathbf{x}}\bar{\mathbf{x}}^\top)\partial f(\bar{\mathbf{x}})$$

(see [14]). In the recent work [23], Chen et al. considered the more general problem of minimizing the sum of a smooth function and a nonsmooth convex function over the Stiefel manifold and developed a manifold proximal gradient method (ManPG) for finding a stationary point of it. When specialized to solve problem (1), the method generates the iterates via

$$\mathbf{x}^{k+1} = \text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \alpha_k \mathbf{d}^k) = \frac{\mathbf{x}^k + \alpha_k \mathbf{d}^k}{\|\mathbf{x}^k + \alpha_k \mathbf{d}^k\|_2}, \quad (5)$$

where the search direction \mathbf{d}^k is given by

$$\begin{aligned} \mathbf{d}^k &= \arg\min_{\mathbf{d} \in \mathbb{R}^n} \|\mathbf{Y}^\top(\mathbf{x}^k + \mathbf{d})\|_1 + \frac{1}{2t}\|\mathbf{d}\|_2^2 \\ &\text{s.t.} \quad \mathbf{d}^\top \mathbf{x}^k = 0 \end{aligned} \quad (6)$$

and $\alpha_k > 0$, $t > 0$ are the step sizes. As the reader may readily recognize, without the constraint $\mathbf{d}^\top \mathbf{x}^k = 0$, the subproblem (6) is simply computing the proximal mapping of f at \mathbf{x}^k and coincides with the update of the classic proximal point algorithm (PPA) [24]. The constraint $\mathbf{d}^\top \mathbf{x}^k = 0$ in (6), which states that the search direction \mathbf{d} should lie on the tangent space to \mathcal{M} at \mathbf{x}^k , is introduced to account for the manifold constraint in problem (1) and ensures that the next iterate \mathbf{x}^{k+1} achieves sufficient decrease in objective value. Motivated by the above discussion, we call the method obtained by specializing ManPG to the setting of problem (1) *ManPPA* and present its details below:

Algorithm 1 ManPPA for Solving Problem (1)

- 1: Input: $\mathbf{x}^0 \in \mathcal{M}$, $\beta \in (0, 1)$, $t > 0$.
- 2: **for** $k = 0, 1, 2, \dots$ **do**
- 3: Solve the subproblem (6) to obtain \mathbf{d}^k .
- 4: Let j_k be the smallest nonnegative integer such that
- 5: $\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \beta^{j_k} \mathbf{d}^k)$ with $\|\mathbf{x}^k + \beta^{j_k} \mathbf{d}^k\|_2^2$.
- 6: **end for**

Naturally, ManPPA inherits the properties of ManPG established in [23]. However, due to the structure of problem (1), many of the developments in [23] have to be refined when designing ManPPA. In particular, the SSN method used by ManPG for finding the search direction in each iteration requires the computation of the proximal mapping of the nonsmooth part of the objective function. However, due to the presence of the matrix \mathbf{Y} , the objective function f of problem (1) does not have an easily computable proximal mapping. As such, the SSN method proposed in [23] cannot efficiently solve the subproblem (6). To circumvent this difficulty, we propose to use an inexact ALM, which can efficiently compute an accurate solution to (6); see Section II-B.

Now, let us state the following result, which shows that the line search step in line 4 of Algorithm 1 is well defined. It simplifies [23, Lemma 5.2] and yields sharper constants. The proof can be found in Appendix B.

Proposition 1. *Let $\{(\mathbf{x}^k, \mathbf{d}^k)\}_k$ be the sequence generated by Algorithm 1. Define $\bar{\alpha} = \min\{1, 1/(tL)\}$. For any $\alpha \in (0, \bar{\alpha}]$, we have*

$$f(\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \alpha \mathbf{d}^k)) \leq f(\mathbf{x}^k) - \frac{\alpha}{2t} \|\mathbf{d}^k\|_2^2. \quad (7)$$

As a result, we have $\alpha_k = \beta^{j_k} > \beta\bar{\alpha}$ for any $k \geq 0$ in Algorithm 1, which implies that the line search step terminates after at most $\lceil \log_{\beta} \bar{\alpha} \rceil + 1$ iterations. In particular, if $t \leq 1/L$, then we have $\bar{\alpha} = 1$, which implies that we can take $j_k = 0$ in line 4 of Algorithm 1; i.e., no line search is needed.

A. Convergence Analysis of ManPPA

In this subsection, we study the convergence behavior of ManPPA. Recall from [23, Lemma 5.3] that if $\mathbf{d}^k = \mathbf{0}$ in (6), then $\mathbf{x}^k \in \mathcal{M}$ is a stationary point of problem (1). This motivates us to call $\mathbf{x}^k \in \mathcal{M}$ an ϵ -stationary point of problem (1) with $\epsilon \geq 0$ if the solution \mathbf{d}^k to (6) satisfies $\|\mathbf{d}^k\|_2 \leq \epsilon$. By specializing the convergence results in [23, Theorem 5.5] for ManPG to ManPPA, we obtain the following theorem:

Theorem 1. *Any limit point of the sequence $\{\mathbf{x}^k\}_k$ generated by Algorithm 1 is a stationary point of problem (1). Moreover, Algorithm 1 with $t = 1/L$ returns an ϵ -stationary point \mathbf{x}^k in at most $\lceil 2(f(\mathbf{x}^0) - f^*)/(L\epsilon^2) \rceil$ iterations, where f^* is the optimal value of problem (1).*

Theorem 1 shows that ManPPA has an iteration complexity of $\mathcal{O}(\epsilon^{-2})$, which is superior to the $\mathcal{O}(\epsilon^{-4})$ bound established for RSGM in [15].

Now, let us analyze the convergence rate of ManPPA in the setting where problem (1) possesses the sharpness property. Such a setting is highly relevant in applications, as both the ODL problem and DPCP formulation of the RSR problem give rise to sharp instances of problem (1) under certain assumptions on the data; see [6, Proposition C.8] and [15, Proposition 4]. To proceed, we first introduce the notion of sharpness.

Definition 1 (Sharpness; see, e.g., [16]). *We say that $\mathcal{X} \subseteq \mathcal{M}$ is a set of weak sharp minima for the function f with parameters*

(α, δ) (where $\alpha, \delta > 0$) if for any $\mathbf{x} \in \mathcal{B}(\delta) := \{\mathbf{x} \in \mathcal{M} \mid \text{dist}(\mathbf{x}, \mathcal{X}) \leq \delta\}$, we have

$$f(\mathbf{x}) - f(\bar{\mathbf{x}}) \geq \alpha \cdot \text{dist}(\mathbf{x}, \mathcal{X}), \quad \forall \bar{\mathbf{x}} \in \mathcal{X}. \quad (8)$$

From the definition, we see that if \mathcal{X} is a set of weak sharp minima of f , then it is the set of minimizers of f over $\mathcal{B}(\delta)$. Moreover, the function value grows linearly with the distance to \mathcal{X} . In the presence of such a regularity property, ManPPA can be shown to converge at a much faster rate. The following result, which has not appeared in the literature before and is thus new, constitutes the first main contribution of this paper.

Theorem 2. *Suppose that $\mathcal{X} \subseteq \mathcal{M}$ is a set of weak sharp minima for the function f with parameters (α, δ) . Let $\{\mathbf{x}^k\}_k$ be the sequence generated by Algorithm 1 with $\text{dist}(\mathbf{x}^0, \mathcal{X}) < \bar{\delta} := \min\{\delta, \frac{\alpha}{L}\}$ and $t \leq \min\left\{\frac{\bar{\delta}}{2\alpha - L\delta}, \frac{2\bar{\delta}\alpha - L\bar{\delta}^2}{L^2}\right\}$. Then, we have*

$$\text{dist}(\mathbf{x}^k, \mathcal{X}) \leq \bar{\delta}, \quad \forall k \geq 0, \quad (9)$$

$$\text{dist}(\mathbf{x}^{k+1}, \mathcal{X}) \leq \mathcal{O}(\text{dist}^2(\mathbf{x}^k, \mathcal{X})), \quad \forall k \geq 0. \quad (10)$$

Theorem 2 establishes the quadratic convergence rate of ManPPA when applied to a sharp instance of problem (1). Again, this is superior to the linear convergence rate of RSGM established in [15] for this setting. The proof of Theorem 2 can be found in Appendix C. Note that since \mathcal{M} is nonconvex, one cannot directly apply standard convergence analysis techniques for PPA (see, e.g., [24]) to obtain Theorem 2. The key to overcoming this difficulty is the new Riemannian subgradient inequality we establish in Proposition 5 (see Appendix C), which provides a path for extending the convergence analysis of PPA to that of ManPPA.

It should be pointed out that the results in Theorems 1 and 2 do not assume any generative model of the data matrix \mathbf{Y} . By contrast, the results developed in, e.g., [6] for the ODL problem and [9] for the DPCP formulation of the RSR problem do assume certain generative models of the data. Although the latter results may yield qualitatively sharper convergence guarantees for instances of (1) that arise from the ODL problem or the DPCP formulation of the RSR problem, the former apply to arbitrary instances of (1).

B. Solving the Subproblem (6)

Observe that each iteration of ManPPA requires solving the subproblem (6) to obtain the search direction. Thus, the efficiency of ManPPA depends not only on its convergence rate (which has already been studied in Theorems 1 and 2) but also on how fast the subproblem (6) can be solved. To address the latter, we note that (6) is a linearly constrained strongly convex quadratic minimization problem. This motivates us to adopt the SSN-based ALM originally developed in [25] for LASSO-type problems to solve it. As we shall see, such an approach yields a highly efficient method for solving the subproblem (6).

To set the stage for our development, let us drop the index k from (6) for simplicity and set $\mathbf{c} = \mathbf{Y}^\top \mathbf{x}$. Then, the subproblem (6) can be equivalently written as

$$\min_{\substack{\mathbf{d} \in \mathbb{R}^n, \\ \mathbf{u} \in \mathbb{R}^p}} \quad \frac{1}{2} \|\mathbf{d}\|_2^2 + t \|\mathbf{u}\|_1 \quad \text{s.t.} \quad \mathbf{Y}^\top \mathbf{d} + \mathbf{c} = \mathbf{u}, \quad \mathbf{d}^\top \mathbf{x} = 0. \quad (11)$$

At this point, one may be tempted to use ADMM to solve problem (11). However, from a practical point of view, ADMM is often unable to return a high-accuracy solution in an efficient manner. Since (11) is a subproblem in ManPPA, a low-accuracy solution will adversely affect the convergence rate of ManPPA. In fact, this has been observed in our numerical experiments. Therefore, we propose to use an inexact ALM, which can solve problem (11) efficiently and accurately. This makes it possible for ManPPA to achieve fast convergence. To describe the algorithm, let us first write down the augmented Lagrangian function corresponding to (11):

$$\begin{aligned} \mathcal{L}_\sigma(\mathbf{d}, \mathbf{u}; y, \mathbf{z}) \\ := & \frac{1}{2}\|\mathbf{d}\|_2^2 + t\|\mathbf{u}\|_1 + y \cdot \mathbf{d}^\top \mathbf{x} + \langle \mathbf{z}, \mathbf{Y}^\top \mathbf{d} + \mathbf{c} - \mathbf{u} \rangle \\ & + \frac{\sigma}{2}(\mathbf{d}^\top \mathbf{x})^2 + \frac{\sigma}{2}\|\mathbf{Y}^\top \mathbf{d} + \mathbf{c} - \mathbf{u}\|_2^2. \end{aligned} \quad (12)$$

Here, $y \in \mathbb{R}$ and $\mathbf{z} \in \mathbb{R}^p$ are Lagrange multipliers (dual variables) associated with the constraints in (11) and $\sigma > 0$ is a penalty parameter. Then, a generic iteration of the inexact ALM is given by

$$(\mathbf{d}^{j+1}, \mathbf{u}^{j+1}) \approx \underset{\substack{\mathbf{d} \in \mathbb{R}^n, \\ \mathbf{u} \in \mathbb{R}^p}}{\operatorname{argmin}} \Psi_j(\mathbf{d}, \mathbf{u}) := \mathcal{L}_{\sigma_j}(\mathbf{d}, \mathbf{u}; y^j, \mathbf{z}^j), \quad (13a)$$

$$y^{j+1} = y^j + \sigma_j(\mathbf{d}^{j+1})^\top \mathbf{x}, \quad (13b)$$

$$\mathbf{z}^{j+1} = \mathbf{z}^j + \sigma_j(\mathbf{Y}^\top \mathbf{d}^{j+1} + \mathbf{c} - \mathbf{u}^{j+1}), \quad (13c)$$

where the penalty parameters $\{\sigma_j\}_j$ are chosen such that $0 < \sigma_j \nearrow \sigma_\infty \leq +\infty$ [25], [26]. Since the subproblem (13a) can only be solved inexactly in general, we adopt the following stopping criteria, which are standard in the literature (see [25], [26]):

$$\Psi_j(\mathbf{d}^{j+1}, \mathbf{u}^{j+1}) - \Psi_j^* \leq \frac{\varepsilon_j^2}{2\sigma_j}, \quad \sum_{j=0}^{\infty} \varepsilon_j < \infty, \quad (14a)$$

$$\begin{aligned} \Psi_j(\mathbf{d}^{j+1}, \mathbf{u}^{j+1}) - \Psi_j^* \\ \leq \frac{\delta_j^2}{2\sigma_j} \|(y^{j+1}, \mathbf{z}^{j+1}) - (y^j, \mathbf{z}^j)\|_2^2, \quad \sum_{j=0}^{\infty} \delta_j < \infty, \end{aligned} \quad (14b)$$

$$\begin{aligned} \operatorname{dist}(\mathbf{0}, \partial\Psi_j(\mathbf{d}^{j+1}, \mathbf{u}^{j+1})) \\ \leq \frac{\delta'_j}{\sigma_j} \|(y^{j+1}, \mathbf{z}^{j+1}) - (y^j, \mathbf{z}^j)\|_2, \quad \delta'_j \searrow 0. \end{aligned} \quad (14c)$$

Here, Ψ_j^* is the optimal value of (13a). Equipped with conditions (14a)–(14c), we can show that starting from any initial point $(\mathbf{d}^0, \mathbf{u}^0; y^0, \mathbf{z}^0)$, the inexact ALM (13) will converge at an asymptotic superlinear rate to an optimal solution to problem (11). This result, which constitutes the second main contribution of this paper, is obtained from a new perturbation analysis of the solution set of the subproblem (6) and its dual. The details can be found in the companion technical report of this paper [27].

Now, it remains to discuss how to solve the subproblem (13a) in an efficient manner. Again, let us drop the index j in (13a)

for simplicity. By simple manipulation, we have

$$\begin{aligned} \Psi(\mathbf{d}, \mathbf{u}) = & \frac{1}{2}\|\mathbf{d}\|_2^2 + \frac{\sigma}{2} \left(\mathbf{d}^\top \mathbf{x} + \frac{y}{\sigma} \right)^2 - \frac{y^2}{2\sigma} - \frac{\|\mathbf{z}\|_2^2}{2\sigma} \\ & + t\|\mathbf{u}\|_1 + \frac{\sigma}{2} \left\| \mathbf{Y}^\top \mathbf{d} + \mathbf{c} + \frac{\mathbf{z}}{\sigma} - \mathbf{u} \right\|_2^2. \end{aligned}$$

Consider the function $\mathbf{d} \mapsto \psi(\mathbf{d}) := \inf_{\mathbf{u} \in \mathbb{R}^p} \Psi(\mathbf{d}, \mathbf{u})$. Upon letting $\mathbf{w} = \mathbf{Y}^\top \mathbf{d} + \mathbf{c} + \frac{\mathbf{z}}{\sigma} \in \mathbb{R}^p$ and using the definition of the proximal mapping of $\mathbf{u} \mapsto h(\mathbf{u}) := t\|\mathbf{u}\|_1$, we have

$$\begin{aligned} \psi(\mathbf{d}) = & \frac{1}{2}\|\mathbf{d}\|_2^2 + \frac{\sigma}{2} \left(\mathbf{d}^\top \mathbf{x} + \frac{y}{\sigma} \right)^2 - \frac{y^2}{2\sigma} - \frac{\|\mathbf{z}\|_2^2}{2\sigma} \\ & + h(\operatorname{prox}_{h/\sigma}(\mathbf{w})) + \frac{\sigma}{2}\|\mathbf{w} - \operatorname{prox}_{h/\sigma}(\mathbf{w})\|_2^2. \end{aligned}$$

It follows that $(\bar{\mathbf{d}}, \bar{\mathbf{u}}) = \underset{\mathbf{d} \in \mathbb{R}^n, \mathbf{u} \in \mathbb{R}^p}{\operatorname{argmin}} \Psi(\mathbf{d}, \mathbf{u})$ if and only if

$$\bar{\mathbf{d}} = \underset{\mathbf{d} \in \mathbb{R}^n}{\operatorname{argmin}} \psi(\mathbf{d}), \quad \bar{\mathbf{u}} = \operatorname{prox}_{h/\sigma} \left(\mathbf{Y}^\top \bar{\mathbf{d}} + \mathbf{c} + \frac{\mathbf{z}}{\sigma} \right).$$

Using [28, Theorem 2.26] and the Moreau decomposition $\mathbf{w} = \operatorname{prox}_{h/\sigma}(\mathbf{w}) + (1/\sigma)\operatorname{prox}_{\sigma h^*}(\sigma\mathbf{w})$, where h^* is the conjugate function of h , it can be deduced that ψ is strongly convex and continuously differentiable with

$$\nabla\psi(\mathbf{d}) = \mathbf{d} + \sigma \left(\mathbf{d}^\top \mathbf{x} + \frac{y}{\sigma} \right) \mathbf{x} + \mathbf{Y} \operatorname{prox}_{\sigma h^*}(\sigma\mathbf{w}).$$

Thus, we can find $\bar{\mathbf{d}}$ by solving the nonsmooth equation

$$\nabla\psi(\mathbf{d}) = \mathbf{0}. \quad (15)$$

Towards that end, we apply an SSN method, which finds the solution by successive linearization of the map $\nabla\psi$. To implement the method, we first need to compute the generalized Jacobian of $\nabla\psi$ [29, Definition 2.6.1], denoted by $\partial(\nabla\psi)$. By the chain rule [29, Corollary of Theorem 2.6.6] and the Moreau decomposition, each element $\mathbf{V} \in \partial(\nabla\psi)$ takes the form

$$\mathbf{V} = \mathbf{I}_n + \sigma \mathbf{Y} (\mathbf{I}_p - \mathbf{Q}) \mathbf{Y}^\top + \sigma \mathbf{x} \mathbf{x}^\top, \quad (16)$$

where $\mathbf{Q} \in \partial \operatorname{prox}_{h/\sigma}(\mathbf{w})$. Using the definition of h , it can be shown that the diagonal matrix $\mathbf{Q} = \operatorname{Diag}(\mathbf{q})$ with

$$q_i = \begin{cases} 0 & \text{if } |w_i| \leq t/\sigma, \\ 1 & \text{otherwise,} \end{cases} \quad i = 1, \dots, p$$

is an element of $\partial \operatorname{prox}_{h/\sigma}(\mathbf{w})$ [25, Section 3.3] and hence can be used to define an element $\mathbf{V} \in \partial(\nabla\psi)$ via (16). Note that the matrix \mathbf{V} so defined is positive definite. As such, the following generic iteration of the SSN method for solving (15) is well defined:

$$\mathbf{v} = -\mathbf{V}^{-1} \nabla\psi(\mathbf{d}^j), \quad (17a)$$

$$\mathbf{d}^{j+1} = \mathbf{d}^j + \rho_j \mathbf{v}. \quad (17b)$$

Here, $\rho_j > 0$ is the step size. Moreover, since $\mathbf{I}_p - \mathbf{Q}$ is a diagonal matrix whose entries are either 0 or 1, the matrix \mathbf{V} can be assembled in a very efficient manner; again, see [25]. In the companion technical report [27], we present a detailed implementation of the SSN method (17) and show that it will converge at a superlinear rate to the unique solution $\bar{\mathbf{d}}$ to (15).

III. STOCHASTIC MANIFOLD PROXIMAL POINT ALGORITHM

In this section, we propose, for the first time, a stochastic ManPPA (StManPPA) for solving problem (1), which is well suited for the setting where p (typically representing the number of data points) is much larger than n (typically representing the ambient dimension of the data points). To begin, observe that problem (1) has the finite-sum structure

$$\min_{\mathbf{x} \in \mathbb{R}^n} \sum_{j=1}^p |\mathbf{y}_j^\top \mathbf{x}| \quad \text{s.t.} \quad \|\mathbf{x}\|_2 = 1,$$

where $\mathbf{y}_j \in \mathbb{R}^n$ is the j -th column of \mathbf{Y} . When p is extremely large, computing the matrix-vector product $\mathbf{Y}^\top \mathbf{x}$ can be expensive. To circumvent this difficulty, in each iteration of StManPPA, a column of \mathbf{Y} , say \mathbf{y}_j , is randomly chosen and the search direction \mathbf{d}^k is given by

$$\begin{aligned} \mathbf{d}^k &= \underset{\substack{\mathbf{d} \in \mathbb{R}^n \\ \mathbf{d}^\top \mathbf{x}^k = 0}}{\operatorname{argmin}} |\mathbf{y}_j^\top (\mathbf{x}^k + \mathbf{d})| + \frac{1}{2t} \|\mathbf{d}\|_2^2 \end{aligned} \quad (18)$$

The key advantage of StManPPA is that the subproblem (18) admits a closed-form solution that is very easy to compute.

Proposition 2. *Let $\mu = t(\mathbf{y}_j^\top \mathbf{x}^k)$. Then, the solution to (18) is given by*

$$\mathbf{d}^k = \begin{cases} \mu \mathbf{x}^k - t \mathbf{y}_j & \text{if } (1 + \mu)\mu/t - t \|\mathbf{y}_j\|_2^2 > 0, \\ -\mu \mathbf{x}^k + t \mathbf{y}_j & \text{if } (1 - \mu)\mu/t + t \|\mathbf{y}_j\|_2^2 < 0, \\ \frac{\mu^2 \mathbf{x}^k - t \mu \mathbf{y}_j}{t^2 \|\mathbf{y}_j\|_2^2 - \mu^2} & \text{otherwise.} \end{cases} \quad (19)$$

Proof. The first-order optimality conditions of (18) are

$$\mathbf{0} \in \frac{1}{t} \mathbf{d} + \partial |\mathbf{y}_j^\top (\mathbf{x}^k + \mathbf{d})| \mathbf{y}_j - \lambda \mathbf{x}^k, \quad (20a)$$

$$0 = \mathbf{d}^\top \mathbf{x}^k. \quad (20b)$$

Suppose that \mathbf{d} is a solution to (20). If $\mathbf{y}_j^\top (\mathbf{x}^k + \mathbf{d}) > 0$, then $\partial |\mathbf{y}_j^\top (\mathbf{x}^k + \mathbf{d})| = 1$, and (20a) implies that $\mathbf{0} = \mathbf{d}/t + \mathbf{y}_j - \lambda \mathbf{x}^k$. This, together with (20b) and the fact that $\|\mathbf{x}^k\|_2 = 1$, gives $\lambda = \mathbf{y}_j^\top \mathbf{x}^k$. It follows that $\mathbf{d} = t(\mathbf{y}_j^\top \mathbf{x}^k) \mathbf{x}^k - t \mathbf{y}_j = \mu \mathbf{x}^k - t \mathbf{y}_j$ and hence $\mathbf{y}_j^\top (\mathbf{x}^k + \mathbf{d}) > 0$ is equivalent to $(1 + \mu)\mu/t - t \|\mathbf{y}_j\|_2^2 > 0$. This establishes the first case in (19). The other two cases in (19), which correspond to $\mathbf{y}_j^\top (\mathbf{x}^k + \mathbf{d}) < 0$ and $\mathbf{y}_j^\top (\mathbf{x}^k + \mathbf{d}) = 0$, can be derived using a similar argument. \square

We now present the details of StManPPA in Algorithm 2. It is worth noting that our proposed StManPPA is different from the ones developed in the recent work [30]. Indeed, in each iteration, the former only needs to solve a subproblem that involves a single component of the objective function, while the latter need to compute the proximal mapping of the entire objective function.

Algorithm 2 StManPPA for Solving Problem (1)

- 1: Input: $\mathbf{x}^0 \in \mathcal{M}$, $t_0, t_1, \dots, t_T > 0$.
- 2: **for** $k = 0, 1, \dots, T$ **do**
- 3: Select $j_k \in \{1, \dots, p\}$ uniformly at random and solve the subproblem (18) with $j = j_k$, $t = t_k$ to obtain \mathbf{d}^k .
- 4: Set $\mathbf{x}^{k+1} = \operatorname{Proj}_{\mathcal{M}}(\mathbf{x}^k + \mathbf{d}^k)$.
- 5: **end for**
- 6: Output: $\bar{\mathbf{x}} = \mathbf{x}^k$ with probability $t_k / \sum_{k=0}^T t_k$.

A. Convergence Analysis of StManPPA

In this section, we present our convergence results for StManPPA. Let us begin with some preparations. Define $f_j : \mathbb{R}^n \rightarrow \mathbb{R}$ to be the function $f_j(\mathbf{x}) = |\mathbf{y}_j^\top \mathbf{x}|$ and let $L_j > 0$ denote the Lipschitz constant of f_j , where $j = 1, \dots, p$. Set $\bar{L} := \max_{j \in \{1, \dots, p\}} L_j$. Furthermore, define the *Moreau envelope* and *proximal mapping* on \mathcal{M} by

$$e_f(\mathbf{z}) = \min_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2, \quad (21a)$$

$$\operatorname{mprox}_f(\mathbf{z}) \in \operatorname{argmin}_{\mathbf{x} \in \mathcal{M}} f(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2, \quad (21b)$$

respectively. The proximal mapping mprox is well defined since the constraint set \mathcal{M} is compact. As it turns out, the proximal mapping mprox can be used to define an alternative notion of stationarity for problem (1). Indeed, for any $\lambda > 0$ and $\mathbf{x} \in \mathcal{M}$, if we denote $\hat{\mathbf{x}} = \operatorname{mprox}_{\lambda f}(\mathbf{x})$, then the optimality condition of (21) yields

$$\mathbf{0} \in \partial_R f(\hat{\mathbf{x}}) + \frac{1}{\lambda} (\mathbf{I}_n - \hat{\mathbf{x}} \hat{\mathbf{x}}^\top)(\hat{\mathbf{x}} - \mathbf{x}).$$

Since $\mathbf{I}_n - \hat{\mathbf{x}} \hat{\mathbf{x}}^\top$ is a projection operator and hence nonexpansive, we obtain

$$\operatorname{dist}(\mathbf{0}, \partial_R f(\hat{\mathbf{x}})) \leq \frac{1}{\lambda} \|\mathbf{x} - \hat{\mathbf{x}}\|_2.$$

In particular, if $\frac{1}{\lambda} \|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq \epsilon$, then (i) $\hat{\mathbf{x}}$ is ϵ -stationary in the sense that $\operatorname{dist}(\mathbf{0}, \partial_R f(\hat{\mathbf{x}})) \leq \epsilon$ and (ii) \mathbf{x} is close to the ϵ -stationary point $\hat{\mathbf{x}}$. This motivates us to use

$$\mathcal{M} \ni \mathbf{x} \mapsto \Theta_\lambda(\mathbf{x}) := \frac{1}{\lambda} \|\mathbf{x} - \operatorname{mprox}_{\lambda f}(\mathbf{x})\|_2$$

as a stationarity measure for problem (1). We call $\mathbf{x} \in \mathcal{M}$ an *ϵ -nearly stationary point* of problem (1) if $\Theta_\lambda(\mathbf{x}) \leq \epsilon$. It is worth noting that such a notion has also been used in [15] to study the stochastic RSGM.

We are now ready to establish the following convergence rate result for StManPPA, which constitutes the third main contribution of this paper.

Theorem 3. *For any $\lambda \in (0, 1/(p\bar{L}))$, the point $\bar{\mathbf{x}}$ output by Algorithm 2 satisfies*

$$\mathbb{E} [\Theta_\lambda(\bar{\mathbf{x}})^2] \leq \frac{2\lambda e_\lambda(\mathbf{x}^0) + \bar{L}^2 \sum_{k=0}^T t_k^2}{\lambda((1/p) - \lambda\bar{L}) \sum_{k=0}^T t_k},$$

where the expectation is taken over all random choices made by the algorithm. In particular, if the step sizes $\{t_k\}_k$ satisfy $\sum_{k=0}^\infty t_k = \infty$ and $\sum_{k=0}^\infty t_k^2 < \infty$, then $\mathbb{E} [\Theta_\lambda(\bar{\mathbf{x}})^2] \rightarrow 0$. Moreover, if we take $t_k = \frac{1}{\sqrt{T+1}}$ for $k = 0, 1, \dots, T$, then the

number of iterations needed by StManPPA to obtain a point $\bar{\mathbf{x}} \in \mathcal{M}$ satisfying $\mathbb{E}[\Theta_\lambda(\bar{\mathbf{x}})] \leq \epsilon$ is $\mathcal{O}(\epsilon^{-4})$.

The proof of Theorem 3 can be found in Appendix D. Again, it makes crucial use of our newly established Riemannian subgradient inequality (see Appendix C, Proposition 5), which allows StManPPA to be analyzed in a similar way as various Euclidean stochastic methods. We remark that the iteration complexity bound $\mathcal{O}(\epsilon^{-4})$ of StManPPA established in Theorem 3 is comparable to that of RSGM established in [15].

IV. EXTENSION TO STIEFEL MANIFOLD CONSTRAINT

In this section, we consider the matrix analog (4) of problem (1), which also arises in many applications such as certain ‘‘one-shot’’ formulations of the ODL and RSR problems (see Section I-B). Currently, there are two existing approaches for solving problem (4), namely a sequential linear programming (SLP) approach and an iteratively reweighted least squares (IRLS) approach [8]. In the SLP approach, the columns of \mathbf{X} are extracted one at a time. Suppose that we have already obtained the first ℓ columns of \mathbf{X} ($\ell = 0, 1, \dots, q-1$) and arrange them in the matrix $\mathbf{X}_\ell \in \mathbb{R}^{n \times \ell}$ (with $\mathbf{X}_0 = \mathbf{0}$). Then, the $(\ell+1)$ -st column of \mathbf{X} is obtained by solving

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{Y}^\top \mathbf{x}\|_1 \quad \text{s.t.} \quad \|\mathbf{x}\|_2 = 1, \quad \mathbf{X}_\ell^\top \mathbf{x} = \mathbf{0}.$$

This is achieved by the alternating linearization and projection (ALP) method, which generates the iterates via

$$\begin{aligned} \mathbf{z}^k &= \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} \|\mathbf{Y}^\top \mathbf{z}\|_1 \quad \text{s.t.} \quad \mathbf{z}^\top \mathbf{x}^{k-1} = 1, \quad \mathbf{X}_\ell^\top \mathbf{z} = \mathbf{0}, \\ \mathbf{x}^k &= \mathbf{z}^k / \|\mathbf{z}^k\|_2. \end{aligned}$$

Note that the \mathbf{z} -subproblem is a linear program, which can be efficiently solved by off-the-shelf solvers.

In the IRLS approach, the following variant of problem (4), which favors row-wise sparsity of $\mathbf{Y}^\top \mathbf{X}$ and has been studied by Lerman et al. in [31], is considered:

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times q}} \|\mathbf{Y}^\top \mathbf{X}\|_{1,2} \quad \text{s.t.} \quad \mathbf{X}^\top \mathbf{X} = \mathbf{I}_q. \quad (22)$$

Here, $\|\mathbf{Y}^\top \mathbf{X}\|_{1,2}$ denotes the sum of the Euclidean norms of the rows of $\mathbf{Y}^\top \mathbf{X}$. The IRLS method for solving (22) generates the iterates via

$$\mathbf{X}^k = \underset{\mathbf{X} \in \mathbb{R}^{n \times q}}{\operatorname{argmin}} \sum_{j=1}^p w_{j,k} \|\mathbf{X}^\top \mathbf{y}_j\|_2^2 \quad \text{s.t.} \quad \mathbf{X}^\top \mathbf{X} = \mathbf{I}_q,$$

where $w_{j,k} = 1 / \max\{\delta, \|\mathbf{X}^{k-1\top} \mathbf{y}_j\|_2\}$ and $\delta > 0$ is a perturbation parameter to prevent the denominator from being 0. The solution \mathbf{X}^k can be obtained via an SVD and is thus easy to implement. However, there has been no convergence guarantee for the IRLS method so far.

Recently, Wang et al. [32] proposed a proximal alternating maximization method for solving a maximization version of (4), which arises in the so-called ℓ_1 -PCA problem (see [7]). However, the method cannot be easily adapted to solve problem (4).

The similarity between problems (1) and (4) suggests that the latter can also be solved by ManPPA, which is indeed the case. However, the SSN method for solving the resulting nonsmooth

equation (i.e., the matrix analog of (15)) can be slow, as the dimension of the linear system (17) is high. Here, we propose an alternative method called *sequential ManPPA*, which solves (4) in a column-by-column manner and constitutes the fourth main contribution of this paper. The method is based on the observation that the objective function in (4) is separable in the columns of $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$. To find \mathbf{x}_1 , we simply solve

$$\min_{\mathbf{x}_1 \in \mathbb{R}^n} \|\mathbf{Y}^\top \mathbf{x}_1\|_1 \quad \text{s.t.} \quad \|\mathbf{x}_1\|_2 = 1$$

using ManPPA as it is an instance of problem (1). Suppose that we have found the first ℓ columns of \mathbf{X} ($\ell = 0, 1, \dots, q-1$) and arrange them in the matrix $\mathbf{Q}_\ell \in \mathbb{R}^{n \times \ell}$ (with $\mathbf{Q}_0 = \mathbf{0}$). Then, we find the $(\ell+1)$ -st column $\mathbf{x}_{\ell+1}$ by solving

$$\min_{\mathbf{x}_{\ell+1} \in \mathbb{R}^n} \|\mathbf{Y}^\top \mathbf{x}_{\ell+1}\|_1 \quad \text{s.t.} \quad \|\mathbf{x}_{\ell+1}\|_2 = 1, \quad \mathbf{Q}_\ell^\top \mathbf{x}_{\ell+1} = \mathbf{0}. \quad (23)$$

As it turns out, problem (23) is equivalent to the deflation process discussed in [11] for sequentially recovering the columns of an orthogonal dictionary. Specifically, let \mathbf{V}_ℓ be an orthonormal basis of the orthogonal complement of \mathbf{Q}_ℓ . We can then find $\mathbf{x}_{\ell+1}$ by solving

$$\min_{\mathbf{q} \in \mathbb{R}^{n-\ell}} \|\mathbf{Y}^\top \mathbf{V}_\ell \mathbf{q}\|_1 \quad \text{s.t.} \quad \|\mathbf{q}\|_2 = 1. \quad (24)$$

Note that (24) has the same form as (1) and thus can be solved by RSGM or PSGM. By contrast, problem (23) has an extra linear constraint and cannot be solved by RSGM or PSGM directly. Nevertheless, our ManPPA can solve both (23) and (24). Let us now briefly discuss how to use ManPPA to solve the former. To simplify notation, let us drop the index ℓ and denote $\mathbf{x} = \mathbf{x}_{\ell+1}$, $\mathbf{Q} = \mathbf{Q}_\ell$. In the k -th iteration of ManPPA, we update the iterate by (5), where the search direction \mathbf{d}^k is computed by

$$\begin{aligned} \min_{\substack{\mathbf{d} \in \mathbb{R}^n, \\ \mathbf{u} \in \mathbb{R}^p}} \quad & \frac{1}{2} \|\mathbf{d}\|_2^2 + t \|\mathbf{u}\|_1 \\ \text{s.t.} \quad & \mathbf{Y}^\top (\mathbf{x}^k + \mathbf{d}) = \mathbf{u}, \quad \mathbf{d}^\top [\mathbf{Q}, \mathbf{x}^k] = \mathbf{0}. \end{aligned}$$

This subproblem can be solved using the SSN-based inexact ALM framework in Section II. We omit the details for succinctness. The sequential ManPPA is guaranteed to find a feasible solution to problem (4). Moreover, as we shall see in Section V, it is computationally much more efficient than the ALP and IRLS methods on the ODL problem and the DPCP formulation of the RSR problem. However, it remains open whether the solution found by sequential ManPPA is a stationary point of problem (4). We leave this question for future research.

V. NUMERICAL EXPERIMENTS

In this section, we compare our proposed ManPPA, StManPPA, and sequential ManPPA with the existing methods PSGM-MBLS, ALP, and IRLS. We do not include RSGM with diminishing step sizes in our comparison, as the numerical results in [9] show that they are slower than PSGM-MBLS and cannot achieve high accuracy. In all the tests, we used the step size $t = 0.1$ for ManPPA and set the maximum number of iterations of ManPPA, inexact ALM, and SSN to 100, 30, and 20,

respectively. We stopped ManPPA if the relative change in the objective values satisfies $|f(\mathbf{x}^k) - f(\mathbf{x}^{k-1})|/f(\mathbf{x}^{k-1}) \leq 10^{-9}$. In the k -th iteration of ManPPA, we stopped the inexact ALM if the primal and dual feasibility of problem (11) satisfy

$$\max \left\{ \sqrt{\|\mathbf{Y}^\top \mathbf{d}^{j+1} + \mathbf{c} - \mathbf{u}^{j+1}\|_2^2 + ((\mathbf{d}^{j+1})^\top \mathbf{x})^2}, \|\nabla \psi_j(\mathbf{d}^{j+1})\|_2 \right\} \leq \epsilon_k = 0.1^k.$$

For SSN, we used the same termination criteria as the ones given in [25] and solved the linear equation (17a) by Cholesky decomposition. We refer the reader to the companion technical report [27] for the setting of the parameters in the inexact ALM and SSN.

For StManPPA, we used the piecewise geometrically diminishing step sizes $t_k = \beta^{\lfloor k/p \rfloor} t_0$ for $k = 0, 1, \dots, pT$ with $t_0 = 0.6$ and $T = 500$. Such step sizes are motivated by those used in PSGM [9]. We use StManPPA- β to specify the parameter β used in the algorithm. We stopped the algorithm if the relative change in the objective values satisfies $|f(\mathbf{x}^k) - f(\mathbf{x}^{k-1})|/f(\mathbf{x}^{k-1}) \leq 10^{-12}$. For PSGM-MBLS, ALP, and IRLS, we used their default settings of the parameters. We stopped ALP if the change in the objective values satisfies $|f(\mathbf{x}^k) - f(\mathbf{x}^{k-1})| \leq 10^{-6}$, while we stopped IRLS if the change in the objective values satisfies $|f(\mathbf{x}^k) - f(\mathbf{x}^{k-1})| \leq 10^{-11}$. With these stopping criteria, the solutions returned by these algorithms usually achieve the same accuracy.

A. DPCP Formulations of the RSR Problem

In this section, we consider the DPCP formulations of the RSR problem. For the recovery of a vector that is orthogonal to the inlier subspace (the *vector* case), we compared the performance of ManPPA, StManPPA, PSGM-MBLS, ALP, and IRLS on problem (1). For the recovery of the entire inlier subspace of known dimension d (the *matrix* case), we compared the performance of sequential ManPPA, PSGM-MBLS, ALP, and IRLS on problem (4) with $q = n - d$.

1) Synthetic Data: We first test the algorithms on synthetic data. The data matrix takes the form $\mathbf{Y} = [\mathbf{X}, \mathbf{O}] \in \mathbb{R}^{n \times (p_1 + p_2)}$. We generated the inlier data \mathbf{X} as $\mathbf{X} = \mathbf{Q}\mathbf{C}$, where $\mathbf{Q} \in \mathbb{R}^{n \times d}$ is an orthonormal matrix and $\mathbf{C} \in \mathbb{R}^{d \times p_1}$ is a coefficient matrix. The matrix \mathbf{Q} was generated by orthonormalizing an $n \times d$ standard Gaussian random matrix via QR decomposition, while the matrix \mathbf{C} was generated as a $d \times p_1$ standard Gaussian random matrix. The outlier data $\mathbf{O} \in \mathbb{R}^{n \times p_2}$ was generated as a standard Gaussian random matrix. Finally, the columns of the matrix \mathbf{Y} were normalized. The d -dimensional subspace spanned by \mathbf{X} is denoted by \mathcal{S} and its orthogonal complement is denoted by \mathcal{S}^\perp .

Vector case. We set the initial point \mathbf{x}^0 of all algorithms to be the eigenvector of $\mathbf{Y}\mathbf{Y}^\top$ corresponding to the eigenvalue with minimum magnitude. We compared the performance of the algorithms on problem (1) with different dimension n , number of inliers p_1 , and number of outliers p_2 in Figure 1. The first row of Figure 1 reports the principal angle

¹The principal angle is the distance between \mathbf{x}^k and \mathcal{S}^\perp . Any optimal solution \mathbf{x}^* to problem (1) is orthogonal to the inlier subspace \mathcal{S} [9, Theorem 1]. Using the Lipschitz continuity of f , we know that θ also measures the function value gap $f(\mathbf{x}) - f(\mathbf{x}^*)$.

θ between \mathbf{x}^k and \mathcal{S}^\perp versus the iteration number. The second row reports θ versus CPU time. Note that $\mathbf{x}^k = \sin(\theta)\mathbf{n} + \cos(\theta)\mathbf{s}$, where $\mathbf{n} = \text{Proj}_{\mathcal{S}}(\mathbf{x}^k)/\|\text{Proj}_{\mathcal{S}}(\mathbf{x}^k)\|_2$ and $\mathbf{s} = \text{Proj}_{\mathcal{S}^\perp}(\mathbf{x}^k)/\|\text{Proj}_{\mathcal{S}^\perp}(\mathbf{x}^k)\|_2$. From Figure 1, we see that PSGM-MBLS is the fastest, while ManPPA is slightly slower. However, they are both much faster than other compared methods. Moreover, the principal angle θ of ManPPA decreases much faster than PSGM-MBLS in terms of iteration number. This can be attributed to the quadratic convergence rate of ManPPA (Theorem 2).




Fig. 1. Numerical results for the DPCP formulation (1). (a): Principal angle versus iteration number. (b): Principal angle versus CPU time.

In Figure 2 we report the quadratic fitting curves of the different algorithms. As shown in [9], the DPCP formulation can tolerate $O((\# \text{inliers})^2)$ outliers; i.e., $p_2 = \mathcal{O}(p_1^2)$. For different $p_2 \in \{40, 80, 120, \dots, 600\}$, we find the smallest $p_1 \in \{60, 70, 80, \dots, 260\}$ such that $\theta < 10^{-1}$. Here, the principal angle θ is the mean value of 10 trials; i.e., we find pairs (p_1, p_2) such that for a fixed p_1 , p_2 is the largest number of outliers that can be tolerated. We then use a quadratic function to fit these pairs (p_1, p_2) . A higher curve indicates that more outliers can be tolerated and hence the algorithm is more robust. From Figure 2, we see that the curve corresponding to PSGM-MBLS is the lowest one and thus the least robust, while ManPPA and StManPPA are more robust. In Figure 3 we report the CPU time versus p_1 and p_2 . For each algorithm, the shadow area represents the standard deviation (std) of 10 random trials, while the line within the shadow is the mean of those trials. From the left two subfigures of Figure 3, we see that the stds of IRLS and ALP are quite significant. In

particular, they are usually more than ten times larger than the stds of other compared algorithms. To better illustrate the stds of the other four algorithms, we plot their CPU times in the right two subfigures of Figure 3. From these two figures, we see that ManPPA has a larger std than those of the other three algorithms. Overall, we see that PSGM-MBLS is the fastest and ManPPA is second, and they are both much faster than the other compared algorithms. Figures 2 and 3 suggest that ManPPA is slightly slower than PSGM-MBLS but is more robust. Moreover, the choice of the parameter β for StManPPA is crucial and challenging, as StManPPA-0.9 is faster but less robust than StManPPA-0.8. We leave the determination of the best parameter β for StManPPA as a future work.




Fig. 2. Quadratic fitting curves ($n = 30$).




Fig. 3. Upper: CPU time versus the number of inliers p_1 ($n = 30$, $p_2 = 320$). Lower: CPU time versus the number of outliers p_2 ($n = 30$, $p_1 = 200$). The shadow area corresponds to the std and the line within the shadow is the mean of 10 random trials.

Matrix case. We solved problem (4) using sequential ManPPA and compared its performance with PSGM-MBLS (applied to (24)), ALP, and IRLS. We report the results for $q = 2$ and $q = 4$ in Figures 4 and 5, respectively. Since ALP has a very high std, for better illustration of the CPU time comparison, we exclude it from the right two subfigures of Figures 4 and 5. The results suggest that sequential ManPPA is not as robust as IRLS but is the second most efficient one among the four compared algorithms. Moreover, we find that




Fig. 4. Comparison on the DPCP formulation (4) with $n = 30$, $q = 2$. First row: Quadratic fitting curves. Second row: CPU time versus the number of inliers p_1 ($p_2 = 320$). Third row: CPU time versus the number of outliers p_2 ($p_1 = 200$). The shadow area corresponds to the std and the line within the shadow is the mean of 10 random trials.

although PSGM-MBLS is very efficient, its fitting curve is not very good. This is due to the fact that PSGM-MBLS is very sensitive to the choice of step size.

2) Real 3D Point Cloud Road Data: Next, we compared ManPPA with PSGM-MBLS on the road detection challenge of the KITTI dataset [33]. This dataset contains image data together with the corresponding 3D points collected by a rotating 3D laser scanner. Similar to [9], we only used the 360° 3D point clouds to determine which points lie on the road plane (inliers) and which do not (outliers). By using homogeneous coordinates, this can be cast as a robust hyperplane learning problem (1) in \mathbb{R}^4 . As reported in [9], PSGM-MBLS is the fastest algorithm when compared with other state-of-the-art methods. Thus, we only compared the performance of ManPPA and PSGM-MBLS on problem (1). Table I reports the area under the Receiver Operator Curve (ROC) and the CPU time. We see that all ROC values of ManPPA are better than those of PSGM-MBLS, with some sacrifice on the CPU time.

B. ODL Problem

We generated instances of the ODL problem by first randomly generating an orthogonal matrix $\hat{\mathbf{X}} \in \mathbb{R}^{n \times n}$ and a Bernoulli-Gaussian matrix $\hat{\mathbf{A}} \in \mathbb{R}^{n \times p}$ with parameter γ (see, e.g., [3]), then setting $\mathbf{Y} = \hat{\mathbf{X}}\hat{\mathbf{A}}$.

Vector case. We first compared the performance of ManPPA and StManPPA with ALP, IRLS, and PSGM-MBLS on problem (1) with $n = 30$, $p = \lceil 10n^{1.5} \rceil$. Figures 6 and 7

TABLE I

AREA UNDER ROC AND CPU TIME FOR ANNOTATED 3D POINT CLOUDS WITH INDEX 0, 21 IN KITTY-CITY-48 AND 1, 45, 120, 137, 153 IN KITTYCITY-5. THE NUMBER IN PARENTHESIS IS THE PERCENTAGE OF OUTLIERS.

KITTY-CITY-48		KITTY-CITY-5					
0 (56%)	21 (57%)	1 (37%)	45 (38%)	120 (53%)	137 (48%)	153(67%)	
Area under ROC							
ManPPA	0.99437	0.99077	0.99810	0.99898	0.87629	0.99969	0.75481
PSGM-MBLS	0.99420	0.99062	0.99802	0.99891	0.86782	0.99968	0.74933
CPU time							
ManPPA	0.129	0.174	0.091	0.066	0.106	0.108	0.066
PSGM-MBLS	0.028	0.015	0.034	0.029	0.029	0.014	0.017




Fig. 5. Comparison on the DPCP formulation (4) with $n = 30$, $q = 4$. First row: Quadratic fitting curves. Second row: CPU time versus the number of inliers p_1 ($p_2 = 320$). Third row: CPU time versus the number of outliers p_2 ($p_1 = 200$). The shadow area corresponds to the std and the line within the shadow is the mean of 10 random trials.

report the iteration numbers and CPU times of the compared algorithms. The quantity θ is the angle between \mathbf{x}^k returned by the algorithm and its nearest column in $\hat{\mathbf{X}}$. From Figures 6 and 7, we see that PSGM-MBLS is the fastest algorithm in terms of CPU time, while ManPPA is slightly slower. However, they are both much faster than the other compared algorithms. Moreover, we see that ManPPA is much faster than PSGM-MBLS in terms of iteration number. This again can be attributed to the quadratic convergence rate of ManPPA (Theorem 2).

In Figures 8 and 9 we report the linear fitting curves for $\log(n)$ and $\log(p)$ and CPU times of ManPPA, StManPPA-0.8, StManPPA-0.9, IRLS, ALP, and PSGM-MBLS. Note that for the ODL problem, it has been found empirically in [6] that the sample size p and dimension n should satisfy




Fig. 6. Numerical results for the ODL problem (1): $n = 30$, $p = \lceil 10n^{1.5} \rceil$, $\gamma = 0.1$.

$p = O(n^2)$ to guarantee recovery. The linear fitting curves were found in the following manner. For a given dimension $n \in \{5, 10, 15, \dots, 50\}$, we find the smallest sample number $p \in 2n + \{10, 20, 30, \dots, 800\}$ such that $\theta < 10^{-1}$. Here the principal angle θ is the mean value of 10 trials. We then use a linear function to fit the points $\{(\log(n), \log(p))\}_{n,p}$. From Figures 8 and 9, we find that PSGM-MBLS is the fastest but its fitting curve is high, which suggests that it is not robust. This is because PSGM-MBLS is very sensitive to the choice of step size. ManPPA appears to be the second fastest but is very robust based on the fitting curve.

Matrix case. To find the entire orthogonal basis, we use sequential ManPPA to solve problem (4) with $q = n$. We compared sequential ManPPA with PSGM-MBLS (applied to (24)) and SLP based on ALP, and report the results in Figures 10 and 11. We see that sequential ManPPA is not as robust as ALP, but it is much faster than ALP. Note that there is nothing




Fig. 7. Numerical results for the ODL problem (1): $n = 30$, $p = \lceil 10n^{1.5} \rceil$, $\gamma = 0.3$.

for IRLS to do, as it tackles the objective function $\|\mathbf{Y}^\top \mathbf{X}\|_{1,2}$, which is a constant when $\mathbf{X}^\top \mathbf{X} = \mathbf{I}_n$. We also find that the fitting curve of PSGM-MBLS is very high, which indicates that it fails to recover the dictionary in many cases. Again, this is due to the high sensitivity of PSGM-MBLS to the choice of step size.

VI. CONCLUSIONS

In this paper, we presented ManPPA and its stochastic variant StManPPA for solving problem (1). By exploiting the manifold structure of the constraint set \mathcal{M} , these methods not only are practically efficient but also possess convergence guarantees that are provably superior to those of existing subgradient-type methods. Using ManPPA as a building block, we also proposed a new sequential approach to solving the matrix analog (4) of problem (1). We conducted extensive numerical experiments to compare the performance of our proposed algorithms with existing ones on the ODL problem and DPCP formulation of the RSR problem. The results demonstrated the efficiency and efficacy of our proposed methods.

APPENDIX

A. Useful Properties of $\text{Proj}_{\mathcal{M}}$

In this section, we collect some useful properties of the projector $\text{Proj}_{\mathcal{M}}$.

Proposition 3. For any $\mathbf{x} \in \mathcal{M}$ and $\mathbf{d} \in \mathbb{R}^n$ satisfying $\mathbf{d}^\top \mathbf{x} = 0$ (i.e., \mathbf{d} is a tangent vector at \mathbf{x}), we have

$$\|\text{Proj}_{\mathcal{M}}(\mathbf{x} + \mathbf{d}) - (\mathbf{x} + \mathbf{d})\|_2 \leq \frac{1}{2} \|\mathbf{d}\|_2^2. \quad (25)$$




Fig. 8. Comparison on the ODL problem (1) with $\gamma = 0.1$. First row: Fitting curves. Second row: CPU time versus the number of samples p , ($n = 30$). Third row: CPU time versus the number of dimension n , ($p = 300$). The shadow area corresponds to the std and the line within the shadow is the mean of 10 random trials.

Moreover, if $\|\mathbf{d}\|_2 \leq D$ for some $D \in (0, +\infty)$, then

$$\|\text{Proj}_{\mathcal{M}}(\mathbf{x} + \mathbf{d}) - \mathbf{x}\|_2 \geq \frac{1}{(1 + D^2)^{3/4}} \|\mathbf{d}\|_2. \quad (26)$$

Proof. It is straightforward to verify that

$$\left\| \frac{\mathbf{x} + \mathbf{d}}{\|\mathbf{x} + \mathbf{d}\|_2} - (\mathbf{x} + \mathbf{d}) \right\|_2 = \sqrt{1 + \|\mathbf{d}\|_2^2} - 1.$$

We then have (25) by using the fact that $\sqrt{1 + x^2} - 1 \leq \frac{1}{2} x^2$ for all $x \in \mathbb{R}$. Similarly, since

$$\left\| \frac{\mathbf{x} + \mathbf{d}}{\|\mathbf{x} + \mathbf{d}\|_2} - \mathbf{x} \right\|_2^2 = 2 \left(1 - \frac{1}{\sqrt{1 + \|\mathbf{d}\|_2^2}} \right)$$

and

$$\frac{1}{\sqrt{1 + x^2}} \leq 1 - \frac{1}{2(1 + D^2)^{3/2}} x^2, \quad \forall x \in [0, D],$$

we get (26). \square

Proposition 4. For any $\mathbf{x}, \mathbf{z} \in \mathcal{M}$ and $\mathbf{d} \in \mathbb{R}^n$ satisfying $\mathbf{d}^\top \mathbf{x} = 0$, we have

$$\|\text{Proj}_{\mathcal{M}}(\mathbf{x} + \mathbf{d}) - \mathbf{z}\|_2 \leq \|\mathbf{x} + \mathbf{d} - \mathbf{z}\|_2.$$




Fig. 9. Comparison on the ODL problem (1) with $\gamma = 0.3$. First row: Fitting curves. Second row: CPU time versus the number of samples p , ($n = 30$). Third row: CPU time versus the number of dimension n , ($p = 300$). The shadow area corresponds to the std and the line within the shadow is the mean of 10 random trials.




Fig. 10. Comparison on the ODL problem (4) with $\gamma = 0.1$. First row: Linear fitting curves. Second row: CPU time versus the number of samples p , ($n = 30$). Third row: CPU time versus the number of dimension n , ($p = 300$). The shadow area corresponds to the std and the line within the shadow is the mean of 10 random trials.

Proof. We compute

$$\begin{aligned}
 & \left\| \frac{\mathbf{x} + \mathbf{d}}{\|\mathbf{x} + \mathbf{d}\|_2} - \mathbf{z} \right\|_2 = 2 - 2 \frac{(\mathbf{x} + \mathbf{d})^\top \mathbf{z}}{\|\mathbf{x} + \mathbf{d}\|_2} \\
 &= 2 + 2(\mathbf{x} + \mathbf{d})^\top \mathbf{z} \left(1 - \frac{1}{\|\mathbf{x} + \mathbf{d}\|_2} \right) - 2(\mathbf{x} + \mathbf{d})^\top \mathbf{z} \\
 &\leq 2 + 2(\|\mathbf{x} + \mathbf{d}\|_2 - 1) - 2(\mathbf{x} + \mathbf{d})^\top \mathbf{z} \\
 &\leq \|\mathbf{x}\|_2^2 + \|\mathbf{z}\|_2^2 + \|\mathbf{d}\|_2^2 - 2(\mathbf{x} + \mathbf{d})^\top \mathbf{z} \\
 &= \|\mathbf{x} + \mathbf{d} - \mathbf{z}\|_2,
 \end{aligned} \tag{27}$$

where (27) follows from the fact that $\|\mathbf{x} + \mathbf{d}\|_2 - 1 = \sqrt{1 + \|\mathbf{d}\|_2^2} - 1 \leq \frac{1}{2} \|\mathbf{d}\|_2^2$. \square

B. Proof of Proposition 1

Since f is Lipschitz with constant L and $\mathbf{d}^\top \mathbf{x}^k = 0$, we have

$$\begin{aligned}
 & |f(\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \alpha \mathbf{d}^k)) - f(\mathbf{x}^k + \alpha \mathbf{d}^k)| \\
 &\leq L \left\| \frac{\mathbf{x}^k + \alpha \mathbf{d}^k}{\|\mathbf{x}^k + \alpha \mathbf{d}^k\|_2} - (\mathbf{x}^k + \alpha \mathbf{d}^k) \right\|_2 \leq \frac{\alpha^2 L}{2} \|\mathbf{d}^k\|_2^2
 \end{aligned}$$

by Proposition 3. Hence, for any $\alpha \in (0, \bar{\alpha}]$, we have

$$\begin{aligned}
 & f(\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \alpha \mathbf{d}^k)) \leq f(\mathbf{x}^k + \alpha \mathbf{d}^k) + \frac{\alpha^2 L}{2} \|\mathbf{d}^k\|_2^2 \\
 &\leq (1 - \alpha) f(\mathbf{x}^k) + \alpha f(\mathbf{x}^k + \mathbf{d}^k) + \frac{\alpha^2 L}{2} \|\mathbf{d}^k\|_2^2 \tag{28a}
 \end{aligned}$$

$$\leq f(\mathbf{x}^k) - \frac{\alpha}{t} \|\mathbf{d}^k\|_2^2 + \frac{\alpha^2 L}{2} \|\mathbf{d}^k\|_2^2 \tag{28b}$$

$$\leq f(\mathbf{x}^k) - \frac{\alpha}{2t} \|\mathbf{d}^k\|_2^2, \tag{28c}$$

where (28a) follows from the convexity of f , (28b) holds because the strong convexity of the objective function in subproblem (6), together with the optimality of $\mathbf{d} = \mathbf{d}^k$ and feasibility of $\mathbf{d} = \mathbf{0}$ for (6), implies that $f(\mathbf{x}^k + \mathbf{d}^k) + \frac{1}{t} \|\mathbf{d}^k\|_2^2 \leq f(\mathbf{x}^k)$, and (28c) is due to $\alpha \leq 1/(tL)$. If $t \leq 1/L$, then $\bar{\alpha} = 1$. This completes the proof.

C. Proof of Theorem 2

We begin with two preparatory results. The first states that the restriction of the objective function f in (1) on the nonconvex constraint set \mathcal{M} satisfies a Riemannian subgradient inequality, which means that with respect to its Riemannian subgradient, the function f behaves almost like a convex function on \mathcal{M} .

Proposition 5. Let $\mathbf{x} \in \mathcal{M}$ and $\mathbf{d} \in \mathbb{R}^n$ be such that $\mathbf{d}^\top \mathbf{x} = 0$. Define $\mathbf{x}^+ = \mathbf{x} + \mathbf{d}$. Then, for any $\mathbf{z} \in \mathcal{M}$ and $\mathbf{s} \in \partial f(\mathbf{x}^+)$, we have

$$f(\mathbf{z}) - f(\mathbf{x}^+) \geq \langle (\mathbf{I}_n - \mathbf{x} \mathbf{x}^\top) \mathbf{s}, \mathbf{z} - \mathbf{x}^+ \rangle - \frac{L}{2} \|\mathbf{z} - \mathbf{x}\|_2^2.$$




Fig. 11. Comparison on the ODL problem (4) with $\gamma = 0.3$. First row: Linear fitting curves. Second row: CPU time versus the number of samples p , ($n = 30$). Third row: CPU time versus the number of dimension n , ($p = 300$). The shadow area corresponds to the std and the line within the shadow is the mean of 10 random trials.

Proof. Since f is convex on \mathbb{R}^n , we have

$$\begin{aligned} f(\mathbf{z}) - f(\mathbf{x}^+) &\geq \langle \mathbf{s}, \mathbf{z} - \mathbf{x}^+ \rangle \\ &= \langle (\mathbf{I}_n - \mathbf{x}\mathbf{x}^\top)\mathbf{s}, \mathbf{z} - \mathbf{x}^+ \rangle + \langle \mathbf{x}\mathbf{x}^\top\mathbf{s}, \mathbf{z} - \mathbf{x}^+ \rangle \end{aligned}$$

Now, observe that

$$\begin{aligned} \langle \mathbf{x}\mathbf{x}^\top\mathbf{s}, \mathbf{z} - \mathbf{x}^+ \rangle &= \langle \mathbf{s}, \mathbf{x}\mathbf{x}^\top(\mathbf{z} - (\mathbf{x} + \mathbf{d})) \rangle \\ &= \langle \mathbf{s}, \mathbf{x}(\mathbf{x}^\top\mathbf{z} - 1) \rangle \geq -\frac{1}{2}\|\mathbf{s}\|_2\|\mathbf{z} - \mathbf{x}\|_2^2, \end{aligned}$$

where the second equality is due to $\mathbf{x}^\top\mathbf{x} = 1$ and $\mathbf{d}^\top\mathbf{x} = 0$, and the inequality follows from the fact that $|\mathbf{x}^\top\mathbf{z} - 1| = \frac{1}{2}\|\mathbf{z} - \mathbf{x}\|_2^2$. Since f is Lipschitz with constant L , we have $\|\mathbf{s}\|_2 \leq L$ by [28, Theorem 9.13]. \square

The second establishes a key recursion for the iterates generated by ManPPA.

Proposition 6. Let $\{\mathbf{x}^k\}_k$ be the sequence generated by Algorithm 1 with $t \leq 1/L$. For any $\bar{\mathbf{x}} \in \mathcal{M}$, we have

$$\begin{aligned} \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}\|_2^2 &\leq (1 + tL)\|\mathbf{x}^k - \bar{\mathbf{x}}\|_2^2 \\ &\quad - 2t(f(\mathbf{x}^k) - f(\bar{\mathbf{x}})) + t^2L^2. \end{aligned}$$

Proof. Since $t \leq 1/L$, we have $\mathbf{x}^{k+1} = \text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \mathbf{d}^k)$ by from Proposition 1. From the optimality condition of the subproblem (6), there exists an $\mathbf{s}^k \in \partial f(\mathbf{x}^k + \mathbf{d}^k)$ such that

$$\mathbf{d}^k = -t(\mathbf{I}_n - \mathbf{x}^k\mathbf{x}^{k\top})\mathbf{s}^k. \quad (29)$$

Denoting $\mathbf{x}^{k+} = \mathbf{x}^k + \mathbf{d}^k$, we have

$$\begin{aligned} \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}\|_2^2 &= \|\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \mathbf{d}^k) - \bar{\mathbf{x}}\|_2^2 \\ &\leq \|\mathbf{x}^k + \mathbf{d}^k - \bar{\mathbf{x}}\|_2^2 \end{aligned} \quad (30a)$$

$$\begin{aligned} &= \|\mathbf{x}^k - \bar{\mathbf{x}}\|_2^2 + 2\langle \mathbf{d}^k, \mathbf{x}^{k+} - \bar{\mathbf{x}} \rangle - \|\mathbf{d}^k\|_2^2 \\ &= \|\mathbf{x}^k - \bar{\mathbf{x}}\|_2^2 - 2t\langle (\mathbf{I}_n - \mathbf{x}^k\mathbf{x}^{k\top})\mathbf{s}^k, \mathbf{x}^{k+} - \bar{\mathbf{x}} \rangle - \|\mathbf{d}^k\|_2^2 \end{aligned} \quad (30b)$$

$$\leq (1 + tL)\|\mathbf{x}^k - \bar{\mathbf{x}}\|_2^2 + 2t(f(\bar{\mathbf{x}}) - f(\mathbf{x}^{k+})) - \|\mathbf{d}^k\|_2^2 \quad (30c)$$

$$\begin{aligned} &\leq (1 + tL)\|\mathbf{x}^k - \bar{\mathbf{x}}\|_2^2 + 2t(f(\bar{\mathbf{x}}) - f(\mathbf{x}^k)) \\ &\quad + 2tL\|\mathbf{d}^k\|_2^2 - \|\mathbf{d}^k\|_2^2 \end{aligned} \quad (30d)$$

$$\leq (1 + tL)\|\mathbf{x}^k - \bar{\mathbf{x}}\|_2^2 + 2t(f(\bar{\mathbf{x}}) - f(\mathbf{x}^k)) + t^2L^2, \quad (30e)$$

where (30a) follows from Proposition 4, (30b) follows from (29), (30c) follows from Proposition 5, (30d) follows from the Lipschitz continuity of f , and (30e) follows from the fact that $2tL\|\mathbf{d}^k\|_2^2 - \|\mathbf{d}^k\|_2^2 = -(\|\mathbf{d}^k\|_2 - tL)^2 + t^2L^2 \leq t^2L^2$. \square

We are now ready to prove Theorem 2. We first prove (9) by induction. Let $\mathbf{x}^* \in \mathcal{X}$ be such that $\text{dist}(\mathbf{x}^k, \mathcal{X}) = \|\mathbf{x}^k - \mathbf{x}^*\|_2$. By invoking Proposition 6 with $\bar{\mathbf{x}} = \mathbf{x}^*$, we have

$$\begin{aligned} \text{dist}^2(\mathbf{x}^{k+1}, \mathcal{X}) &\leq \|\mathbf{x}^{k+1} - \mathbf{x}^*\|_2^2 \\ &\leq (1 + tL)\|\mathbf{x}^k - \mathbf{x}^*\|_2^2 - 2t(f(\mathbf{x}^k) - f(\mathbf{x}^*)) + t^2L^2 \\ &\leq (1 + tL)\text{dist}^2(\mathbf{x}^k, \mathcal{X}) - 2\alpha t \cdot \text{dist}(\mathbf{x}^k, \mathcal{X}) + t^2L^2, \end{aligned}$$

where the last inequality follows from (8). Consider the function $[0, \bar{\delta}] \ni s \mapsto \phi(s) = (1 + tL)s^2 - 2t\alpha s + t^2L^2$. Observe that ϕ attains its maximum at $s = \bar{\delta}$ if $\bar{\delta} \geq \frac{2t\alpha}{1+tL}$. Given that $t \leq \min\left\{\frac{\bar{\delta}}{2\alpha - L\bar{\delta}}, \frac{2\bar{\delta}\alpha - L\bar{\delta}^2}{L^2}\right\}$, we indeed have $\bar{\delta} \geq \frac{2t\alpha}{1+tL}$ and hence $\phi(s) \leq \phi(\bar{\delta}) \leq \bar{\delta}^2$ for all $s \in [0, \bar{\delta}]$. In particular, we have $\text{dist}(\mathbf{x}^{k+1}, \mathcal{X}) \leq \bar{\delta}$ whenever $\text{dist}(\mathbf{x}^k, \mathcal{X}) \leq \bar{\delta}$. This establishes (9).

Next, we prove (10). Again, let $\mathbf{x}^* \in \mathcal{X}$ be such that $\text{dist}(\mathbf{x}^k, \mathcal{X}) = \|\mathbf{x}^k - \mathbf{x}^*\|_2$. Since $\alpha \leq L$ by (8), we have $\bar{\delta} \leq 1$. This implies that $\mathbf{x}^{k\top}\mathbf{x}^* = \frac{2 - \|\mathbf{x}^k - \mathbf{x}^*\|_2^2}{2} \geq \frac{1}{2}$. Hence, the vector $\bar{\mathbf{d}}^k = \frac{\mathbf{x}^*}{\mathbf{x}^{k\top}\mathbf{x}^*} - \mathbf{x}^k$ is well defined and satisfies $\bar{\mathbf{d}}^{k\top}\mathbf{x}^k = 0$ (i.e., $\bar{\mathbf{d}}^k$ is a tangent vector at \mathbf{x}^k), $\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \bar{\mathbf{d}}^k) = \mathbf{x}^*$, and $\|\bar{\mathbf{d}}^k\|_2 \leq \sqrt{3}$. By the strong convexity of the objective function in subproblem (6) and noting the optimality of \mathbf{d}^k and feasibility of $\bar{\mathbf{d}}^k$ for (6), we have

$$\begin{aligned} f(\mathbf{x}^k + \mathbf{d}^k) + \frac{1}{2t}\|\mathbf{d}^k\|_2^2 + \frac{1}{2t}\|\mathbf{d}^k - \bar{\mathbf{d}}^k\|_2^2 \\ \leq f(\mathbf{x}^k + \bar{\mathbf{d}}^k) + \frac{1}{2t}\|\bar{\mathbf{d}}^k\|_2^2. \end{aligned} \quad (31)$$

Furthermore, by the Lipschitz continuity of f and Proposition 3, we get

$$\begin{aligned} f(\mathbf{x}^k + \bar{\mathbf{d}}^k) &\leq f(\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \bar{\mathbf{d}}^k)) + \frac{L}{2}\|\bar{\mathbf{d}}^k\|_2^2 \\ &= f(\mathbf{x}^*) + \frac{L}{2}\|\bar{\mathbf{d}}^k\|_2^2, \end{aligned} \quad (32a)$$

$$\begin{aligned} f(\mathbf{x}^{k+1}) &= f(\text{Proj}_{\mathcal{M}}(\mathbf{x}^k + \mathbf{d}^k)) \\ &\leq f(\mathbf{x}^k + \mathbf{d}^k) + \frac{L}{2}\|\mathbf{d}^k\|_2^2. \end{aligned} \quad (32b)$$

Combining (31), (32a), and (32b), we have

$$\begin{aligned} & f(\mathbf{x}^{k+1}) + \left(\frac{1}{2t} - \frac{L}{2} \right) \|\mathbf{d}^k\|_2^2 + \frac{1}{2t} \|\mathbf{d}^k - \bar{\mathbf{d}}^k\|_2^2 \\ & \leq f(\mathbf{x}^*) + \left(\frac{L}{2} + \frac{1}{2t} \right) \|\bar{\mathbf{d}}^k\|_2^2. \end{aligned} \quad (33)$$

Since $t \leq \frac{\bar{\delta}}{2\alpha - L\bar{\delta}} \leq 1/L$, we have $\frac{1}{2t} - \frac{L}{2} \geq 0$. Moreover, since $\text{dist}(\mathbf{x}^{k+1}, \mathcal{X}) \leq \bar{\delta} \leq \delta$, we have $f(\mathbf{x}^{k+1}) - f(\mathbf{x}^*) \geq \alpha \cdot \text{dist}(\mathbf{x}^{k+1}, \mathcal{X})$ by (8). It then follows from (33) and Proposition 3 that

$$\begin{aligned} \alpha \cdot \text{dist}(\mathbf{x}^{k+1}, \mathcal{X}) & \leq \left(\frac{L}{2} + \frac{1}{2t} \right) \|\bar{\mathbf{d}}^k\|_2^2 \\ & \leq 8 \left(\frac{L}{2} + \frac{1}{2t} \right) \|\mathbf{x}^k - \mathbf{x}^*\|_2^2 = 4 \left(L + \frac{1}{t} \right) \text{dist}^2(\mathbf{x}^k, \mathcal{X}). \end{aligned}$$

This completes the proof.

D. Proof of Theorem 3

Let $\hat{\mathbf{x}}^k = \text{mprox}_{\lambda f}(\mathbf{x}^k) \in \mathcal{M}$. By definition of e_λ in (21a) and Proposition 4, we have

$$\begin{aligned} e_\lambda(\mathbf{x}^{k+1}) & \leq f(\hat{\mathbf{x}}^k) + \frac{1}{2\lambda} \|\hat{\mathbf{x}}^k - \mathbf{x}^{k+1}\|_2^2 \\ & \leq f(\hat{\mathbf{x}}^k) + \frac{1}{2\lambda} \|\hat{\mathbf{x}}^k - (\mathbf{x}^k + \mathbf{d}^k)\|_2^2. \end{aligned} \quad (34)$$

From the optimality condition of (18), we get

$$\mathbf{d}^k \in -t_k(\mathbf{I}_n - \mathbf{x}^k \mathbf{x}^{k\top}) \partial f_{j_k}(\mathbf{x}^k + \mathbf{d}^k).$$

Hence, we compute

$$\begin{aligned} & \|\hat{\mathbf{x}}^k - (\mathbf{x}^k + \mathbf{d}^k)\|_2^2 \\ & = \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 - \|\mathbf{d}^k\|_2^2 - 2\langle \hat{\mathbf{x}}^k - \mathbf{x}^k - \mathbf{d}^k, \mathbf{d}^k \rangle \\ & \leq \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 - \|\mathbf{d}^k\|_2^2 \\ & \quad + 2t_k \left(f_{j_k}(\hat{\mathbf{x}}^k) - f_{j_k}(\mathbf{x}^k + \mathbf{d}^k) + \frac{L_{j_k}}{2} \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 \right) \end{aligned} \quad (35a)$$

$$\begin{aligned} & \leq \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 - \|\mathbf{d}^k\|_2^2 + 2t_k (f_{j_k}(\hat{\mathbf{x}}^k) - f_{j_k}(\mathbf{x}^k)) \\ & \quad + 2t_k L_{j_k} \left(\frac{1}{2} \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 + \|\mathbf{d}^k\|_2 \right), \end{aligned} \quad (35b)$$

where (35a) follows from Proposition 5 and (35b) is due to the Lipschitz continuity of f_{j_k} . Upon taking the expectation on both sides of (35) with respect to j_k conditioned on \mathbf{x}^k , we obtain

$$\begin{aligned} & \mathbb{E} [\|\hat{\mathbf{x}}^k - (\mathbf{x}^k + \mathbf{d}^k)\|_2^2 | \mathbf{x}^k] \\ & \leq (1 + t_k \bar{L}) \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 + \frac{2t_k}{p} (f(\hat{\mathbf{x}}^k) - f(\mathbf{x}^k)) \\ & \quad + 2t_k \bar{L} \cdot \mathbb{E} [\|\mathbf{d}^k\|_2 | \mathbf{x}^k] - \mathbb{E} [\|\mathbf{d}^k\|_2^2 | \mathbf{x}^k] \\ & \leq (1 + t_k \bar{L}) \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 + \frac{2t_k}{p} (f(\hat{\mathbf{x}}^k) - f(\mathbf{x}^k)) \\ & \quad + t_k^2 \bar{L}^2 \end{aligned} \quad (36a)$$

$$\begin{aligned} & = \left(1 + t_k \bar{L} - \frac{t_k}{p\lambda} \right) \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 \\ & \quad + \frac{2t_k}{p} (e_\lambda(\mathbf{x}^k) - f(\mathbf{x}^k)) + t_k^2 \bar{L}^2 \\ & \leq \left(1 + t_k \bar{L} - \frac{t_k}{p\lambda} \right) \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 + t_k^2 \bar{L}^2, \end{aligned} \quad (36b)$$

where (36a) follows from the fact that $\mathbb{E} [\|\mathbf{d}^k\|_2 | \mathbf{x}^k] \leq \sqrt{\mathbb{E} [\|\mathbf{d}^k\|_2^2 | \mathbf{x}^k]}$ and $a\sqrt{x} - x \leq a^2/4$ for any $a, x \geq 0$; (36b) follows from the definition of e_λ . Putting (34) and (36b) together gives

$$\begin{aligned} & e_\lambda(\mathbf{x}^{k+1}) \\ & \leq f(\hat{\mathbf{x}}^k) + \frac{1}{2\lambda} \left(1 + t_k \bar{L} - \frac{t_k}{p\lambda} \right) \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 + \frac{t_k^2 \bar{L}^2}{2\lambda} \\ & = e_\lambda(\mathbf{x}^k) + \frac{(\bar{L} - 1/(p\lambda))t_k}{2\lambda} \|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2 + \frac{t_k^2 \bar{L}^2}{2\lambda}. \end{aligned}$$

Taking expectation on both sides with respect to \mathbf{x}^k yields

$$\begin{aligned} & \frac{(1/(p\lambda) - \bar{L})t_k}{2\lambda} \mathbb{E} [\|\hat{\mathbf{x}}^k - \mathbf{x}^k\|_2^2] \\ & \leq \mathbb{E} [e_\lambda(\mathbf{x}^k)] - \mathbb{E} [e_\lambda(\mathbf{x}^{k+1})] + \frac{t_k^2 \bar{L}^2}{2\lambda}. \end{aligned}$$

Upon summing the above inequality over $k = 0, 1, \dots, T$ and noting that $\lambda < 1/(p\bar{L})$ and $e_\lambda(\mathbf{z}) \geq 0$ for any $\mathbf{z} \in \mathbb{R}^n$, we obtain

$$\begin{aligned} & \sum_{k=0}^T t_k \mathbb{E} \left[\frac{1}{\lambda^2} \|\mathbf{x}^k - \text{mprox}_{\lambda f}(\mathbf{x}^k)\|_2^2 \right] \\ & \leq \frac{2}{1/p - \lambda \bar{L}} e_\lambda(\mathbf{x}^0) + \frac{\bar{L}^2}{\lambda(1/p - \lambda \bar{L})} \sum_{k=0}^T t_k^2. \end{aligned}$$

Upon dividing both sides of the above inequality by $\sum_{k=0}^T t_k$ and noting that the left-hand side becomes $\mathbb{E} [\Theta_\lambda(\bar{\mathbf{x}})^2]$, the proof is complete.

REFERENCES

- [1] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Rev.*, vol. 51, no. 1, pp. 34–81, 2009.
- [2] R. Rubinstein, A. M. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," *Proc. IEEE*, vol. 98, no. 6, pp. 1045–1057, 2010.
- [3] D. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," in *COLT*, 2012.
- [4] Q. Qu, J. Sun, and J. Wright, "Finding a sparse vector in subspace: linear sparsity using alternating directions," *IEEE Trans. Inf. Theory*, vol. 62, no. 10, pp. 5855–5880, 2016.
- [5] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery over the sphere I: Overview and the geometric picture," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 853–884, 2017.
- [6] Y. Bai, Q. Jiang, and J. Sun, "Subgradient descent learns orthogonal dictionaries," in *ICLR*, 2019.
- [7] G. Lerman and T. Maunu, "An overview of robust subspace recovery," *Proc. IEEE*, vol. 106, no. 8, pp. 1380–1410, 2018.
- [8] M. C. Tsakiris and R. Vidal, "Dual principal component pursuit," *J. Mach. Learn. Res.*, vol. 19, no. 18, pp. 1–50, 2018.
- [9] Z. Zhu, Y. Wang, D. Robinson, D. Naiman, R. Vidal, and M. Tsakiris, "Dual principal component pursuit: Improved analysis and efficient algorithms," in *NeurIPS*, 2018, pp. 2171–2181.
- [10] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.
- [11] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery over the sphere II: Recovery by Riemannian trust-region method," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 885–914, 2017.
- [12] D. Gilboa, S. Buchanan, and J. Wright, "Efficient dictionary learning with gradient descent," in *ICML*, 2019, pp. 2252–2259.
- [13] Q. Qu, Y. Zhai, X. Li, Y. Zhang, and Z. Zhu, "Geometric analysis of nonconvex optimization landscapes for overcomplete learning," in *ICLR*, 2020.
- [14] W. H. Yang, L.-H. Zhang, and R. Song, "Optimality conditions for the nonlinear programming problems on Riemannian manifolds," *Pacific J. Optim.*, vol. 10, no. 2, pp. 415–434, 2014.

- [15] X. Li, S. Chen, Z. Deng, Q. Qu, Z. Zhu, and A. M.-C. So, "Weakly convex optimization over Stiefel manifold using Riemannian subgradient-type methods," *SIAM J. Optim.*, vol. 31, no. 3, pp. 1605–1634, 2021.
- [16] J. V. Burke and M. C. Ferris, "Weak sharp minima in mathematical programming," *SIAM J. Control Optim.*, vol. 31, no. 5, pp. 1340–1359, 1993.
- [17] X. Li, Z. Zhu, A. M.-C. So, and R. Vidal, "Nonconvex robust low-rank matrix recovery," *SIAM J. Optim.*, vol. 30, no. 1, pp. 660–686, 2020.
- [18] H. Liu, M.-C. Yue, and A. M.-C. So, "On the Estimation Performance and Convergence Rate of the Generalized Power Method for Phase Synchronization," *SIAM J. Optim.*, vol. 27, no. 4, pp. 2426–2446, 2017.
- [19] Z. Zhou and A. M.-C. So, "A Unified Approach to Error Bounds for Structured Convex Optimization Problems," *Math. Programm., Ser. A*, vol. 165, no. 2, pp. 689–728, 2017.
- [20] H. Liu, A. M.-C. So, and W. Wu, "Quadratic optimization with orthogonality constraint: Explicit Łojasiewicz exponent and linear convergence of retraction-based line-search and stochastic variance-reduced gradient methods," *Math. Programm., Ser. A*, vol. 178, no. 1–2, pp. 215–262, 2019.
- [21] O. P. Ferreira and P. R. Oliveira, "Proximal point algorithm on Riemannian manifolds," *Optim.*, vol. 51, no. 2, pp. 257–270, 2002.
- [22] G. C. Bento, O. P. Ferreira, and J. G. Melo, "Iteration-complexity of gradient, subgradient and proximal point methods on Riemannian manifolds," *J. Optim. Theory Appl.*, vol. 173, pp. 548–562, 2017.
- [23] S. Chen, S. Ma, A. M.-C. So, and T. Zhang, "Proximal gradient method for nonsmooth optimization over the Stiefel manifold," *SIAM J. Optim.*, vol. 30, no. 1, pp. 210–239, 2020.
- [24] R. T. Rockafellar, "Monotone operators and the proximal point algorithm," *SIAM J. Control Optim.*, vol. 14, pp. 877–898, 1976.
- [25] X. Li, D. Sun, and K.-C. Toh, "A highly efficient semi-smooth Newton augmented Lagrangian method for solving Lasso problems," *SIAM J. Optim.*, vol. 28, pp. 433–458, 2018.
- [26] R. T. Rockafellar, "Augmented Lagrangians and applications of the proximal point algorithm in convex programming," *Math. Oper. Res.*, vol. 1, no. 2, pp. 97–116, 1976.
- [27] S. Chen, Z. Deng, S. Ma, and A. M.-C. So, "Manifold proximal point algorithms for dual principal component pursuit and orthogonal dictionary learning," *arXiv preprint https://arxiv.org/abs/2005.02356*, 2020.
- [28] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*. Springer-Verlag, 2004.
- [29] F. H. Clarke, *Optimization and Nonsmooth Analysis*. Society for Industrial and Applied Mathematics, 1990.
- [30] B. Wang, S. Ma, and L. Xue, "Riemannian stochastic proximal gradient methods for nonsmooth optimization over the Stiefel manifold," *arXiv preprint arXiv:2005.01209*, 2020.
- [31] G. Lerman, M. McCoy, J. A. Tropp, and T. Zhang, "Robust computation of linear models by convex relaxation," *Found. Comput. Math.*, vol. 15, no. 2, pp. 363–410, 2015.
- [32] P. Wang, H. Liu, and A. M.-C. So, "Globally convergent accelerated proximal alternating maximization method for L1-principal component analysis," in *ICASSP*, 2019, pp. 8147–8151.
- [33] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, 2013.