

On Asymptotic Linear Convergence of Projected Gradient Descent for Constrained Least Squares

Trung Vu, *Graduate Student Member, IEEE*, Raviv Raich, *Senior Member, IEEE*

Abstract—Many recent problems in signal processing and machine learning such as compressed sensing, image restoration, matrix/tensor recovery, and non-negative matrix factorization can be cast as constrained optimization. Projected gradient descent is a simple yet efficient method for solving such constrained optimization problems. Local convergence analysis furthers our understanding of its asymptotic behavior near the solution, offering sharper bounds on the convergence rate compared to global convergence analysis. However, local guarantees often appear scattered in problem-specific areas of machine learning and signal processing. This manuscript presents a unified framework for the local convergence analysis of projected gradient descent in the context of constrained least squares. The proposed analysis offers insights into pivotal local convergence properties such as the conditions for linear convergence, the region of convergence, the exact asymptotic rate of convergence, and the bound on the number of iterations needed to reach a certain level of accuracy. To demonstrate the applicability of the proposed approach, we present a recipe for the convergence analysis of projected gradient descent and demonstrate it via a beginning-to-end application of the recipe on four fundamental problems, namely, linear equality-constrained least squares, sparse recovery, least squares with the unit norm constraint, and matrix completion.

Index Terms—Projected gradient descent, constrained least squares, local linear convergence, asymptotic convergence rate.

I. INTRODUCTION

CONSTRAINED least squares can be formulated as the following optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \quad \text{s.t. } \mathbf{x} \in \mathcal{C}, \quad (1)$$

where $\mathcal{C} \subseteq \mathbb{R}^n$ is a non-empty closed set, $\mathbf{A} \in \mathbb{R}^{m \times n}$, and $\mathbf{b} \in \mathbb{R}^m$ is the observation from which we wish to recover the solution \mathbf{x}^* efficiently. With the surge in the amount of data over the past decades, modern learning problems have become increasingly complex and optimization in the presence of constraints is frequently used to capture accurately their inherent structure. Examples in the area of machine learning and signal processing include, but are not limited to, compressed sensing [1]–[3], image restoration [4]–[6], seismic inversion [7]–[9], and phase-only beamforming [10], [11]. Since the set of real $n_1 \times n_2$ matrices is isomorphic to $\mathbb{R}^{n_1 n_2}$, application of (1) is also found in problems such as low-rank matrix recovery [12]–[14] and non-negative matrix factorization [15]–[17].

Trung Vu and Raviv Raich are with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331, USA. (e-mails: vutru@oregonstate.edu and raich@oregonstate.edu).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the author. The material includes a tutorial-style treatment of convergence analysis and the proofs of Example 1 and Subsection IV-B. This material is 283 KB in size.

Projected gradient descent (PGD) is one of the most popular methods for solving constrained optimization, thanks to its simplicity and efficiency. In theory, convergence properties of this method are natural extensions of the classical results for unconstrained optimization [18]–[21]. When the constraint set \mathcal{C} is convex, PGD is also known as the projected Landweber iteration [22] and is shown to converge sublinearly to the global solution of (1). Moreover, when the least-squares objective is strongly convex, the algorithm enjoys fast linear convergence. For non-convex settings, with the recent introduction of restricted (strong) convexity, global convergence has been guaranteed for certain structural constraints such as sparsity constraint [23], low-rank constraint [24], and L2-norm constraint [25]. For a more comprehensive review of convergence analysis for PGD in the literature, we refer the reader to Section I of the Supplementary Material.

From a different perspective, problem (1) can be viewed as a manifold optimization problem in which the intrinsic structure of manifolds can be exploited. Dating back to the 1970s, Luenberger [26] studied a variant of gradient projection method using the concept of geodesic descent. Under the assumption that \mathcal{C} is a differentiable manifold in Euclidean space, the author provided sufficient conditions for global convergence and established a sharp bound on the asymptotic convergence rate near a strict local minimum. Later on, this result was extended to a broader class of Riemannian manifolds and has been widely known as the Riemannian steepest descent method [26]–[29]. The asymptotic convergence rate of Riemannian steepest descent (with exact line search) is given by the Kantorovich ratio $(\beta - \alpha)^2 / (\beta + \alpha)^2$, where α and β are the smallest and largest eigenvalues of the second derivative of the Lagrangian restricted to the subspace tangent to the constraint manifold at the solution. Remarkably, such local convergence bounds are tighter than those obtained from the aforementioned global convergence analysis in the optimization literature since the former exploits the local structure of the problem. The global convergence bounds, on the other hand, take into account the worst-case behavior of the algorithm that might occur far away from the solution of interest. In certain situations, global convergence analysis suggests sublinear convergence while local convergence analysis offers linear convergence thanks to the benign structure near the solution [30]. One key element in the asymptotic convergence analysis of Riemannian steepest descent is Kantorovich inequality [31]. However, this technique depends on the optimal choice of step size in the exact line search scheme and is not straightforwardly generalized to other variants of gradient projection. To the best of our knowledge, there has been

no direct extension of the analysis for Riemannian steepest descent method to plain PGD with a fixed step size.

Our Contribution. In this paper, we develop a unified framework for a local convergence analysis of the PGD algorithm. We leverage our earlier preliminary work, in which we developed a convergence rate only analysis for the specific problems of low-rank matrix completion [32] and minimization of a quadratic with spherical constraints [33]. For the former, we developed two acceleration approaches that leverage on the rate analysis [34], [35]. The key approach used in these works is to represent each algorithm as a fixed point iteration and to approximate the fixed point operator as locally linear. This idea extends to other algorithms (i.e., non PGD) that can be represented using a fixed point iteration (e.g., see our work on analyzing GD for symmetric matrix completion [36]). For each problem, problem-specific properties have been utilized to facilitate the analysis. Here, our goal is to develop a *unified* framework for convergence rate analysis of PGD for constrained least-squares. Our framework relies on three key steps: (i) the introduction of Lipschitz-continuous differentiability to provide tight error bounds on the linear approximation of the projection operator near the solution, (ii) the establishment of an asymptotically-linear recursion on the error iterations, and (iii) the derivation of the linear rate and the region of convergence (ROC) of the error sequence by leveraging our work on the convergence of nonlinear difference equations [37]. Our approach shifts the burden of the analysis to the characterization of the projection operator (for an example of such characterization of the projection onto the rank- r manifold, see [38]-Theorem 1). In the context of PGD for the general constrained least squares, the proposed framework is the first to offer a closed-form expression of the exact asymptotic rate of local linear convergence, the ROC, and a bound on the number of iterations needed to reach a certain level of accuracy.¹ To illustrate the utility of the approach, we apply our framework to four well-known problems in machine learning and signal processing, namely, linear equality-constrained least squares, sparse recovery, least squares with spherical constraint, and matrix completion. We show that the obtained asymptotic rate of convergence matches existing results in the literature. For problems in which the exact convergence rate of PGD has not been studied, we verify the asymptotic rate obtained by our analysis against the rate of convergence obtained in numerical experiments. We believe that this framework can be used as a general recipe to develop quick yet sharp local convergence results for PGD in other applications in the field as well as to complement conservative analysis of global convergence.

Organization. The rest of this paper is organized as follows. Section II provides a brief background of PGD for constrained least squares, including properties of the orthogonal projection, stationary points of the problem, and the PGD algorithm along with its fixed points. Next, we present our unified

¹We note that the classic work of Polyak [39] can be considered as a replacement for our analysis in the third step. While such result is more general in the context of nonlinear different equations, we do not find a straightforward extension to obtain the ROC and the guarantees on the number of required iterations in our context of convergence analysis.

framework for the local convergence analysis of PGD in Section III, followed by the proof of the main theorem. Then, Section IV demonstrates the application of the proposed recipe to four well-known problems in machine learning and signal processing. Finally, we summarize our results and discuss some of the possible extensions in Section V.

II. PRELIMINARIES

This section presents key concepts and background results that will be used as the basic premise of our subsequent convergence analysis.

A. Notation

Throughout the paper, we use the notation $\|\cdot\|$ to denote the Euclidean norm for vectors. For matrices, $\|\cdot\|_F$ and $\|\cdot\|_2$ denote the Frobenius norm and the spectral norm, respectively. Boldfaced symbols are reserved for vectors and matrices. Additionally, the $t \times t$ identity matrix is denoted by \mathbf{I}_t and the i th vector in the natural basis of \mathbb{R}^n is denoted by \mathbf{e}_i . We use \otimes to denote the Kronecker product between two matrices. The vectorization of a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, denoted by $\text{vec}(\mathbf{X})$, is the concatenation of the columns of a matrix one on top of another in their original order, i.e., for $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$, $\text{vec}(\mathbf{X}) = [\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top]^\top$. Given a vector $\mathbf{x} \in \mathbb{R}^n$, $\text{diag}(\mathbf{x})$ denotes the a square diagonal matrix such that $[\text{diag}(\mathbf{x})]_{ii} = x_i$. For a scalar $r > 0$, denote the open ball of center \mathbf{x} and radius r by $\mathcal{B}(\mathbf{x}, r) = \{\mathbf{y} \mid \|\mathbf{y} - \mathbf{x}\| < r\}$. Correspondingly, the closed ball of center \mathbf{x} and radius r is denoted by $\mathcal{B}[\mathbf{x}, r] = \{\mathbf{y} \mid \|\mathbf{y} - \mathbf{x}\| \leq r\}$. The lexicographical order between two vectors \mathbf{x} and \mathbf{y} of the same length is defined by $\mathbf{x} < \mathbf{y}$ if $x_i < y_i$ for the first i (i goes from 1) where x_i and y_i differ. The lexicographical order between two matrices \mathbf{X} and \mathbf{Y} of the same size is define by the lexicographical order between $\text{vec}(\mathbf{X})$ and $\text{vec}(\mathbf{Y})$.

Given $\mathbf{A} \in \mathbb{R}^{m \times n}$, the i th largest eigenvalue and the i th largest singular value of \mathbf{A} are denoted by $\lambda_i(\mathbf{A})$ and $\sigma_i(\mathbf{A})$, respectively. The spectral radius of \mathbf{A} is defined as $\rho(\mathbf{A}) = \max_i |\lambda_i(\mathbf{A})|$ and is less than or equal to the spectral norm, i.e., $\rho(\mathbf{A}) \leq \|\mathbf{A}\|_2$. Gelfand's formula [40] states that $\rho(\mathbf{A}) = \lim_{k \rightarrow \infty} \|\mathbf{A}^k\|_2^{1/k}$. If \mathbf{A} is square and invertible, the condition number of \mathbf{A} is defined as $\kappa(\mathbf{A}) = \sigma_1(\mathbf{A})/\sigma_n(\mathbf{A})$.

B. Nonlinear Orthogonal Projections

Given a non-empty set $\mathcal{C} \subseteq \mathbb{R}^n$, let us define the distance from a point $\mathbf{x} \in \mathbb{R}^n$ to \mathcal{C} as

$$d(\mathbf{x}, \mathcal{C}) = \inf_{\mathbf{y} \in \mathcal{C}} \{\|\mathbf{y} - \mathbf{x}\|\}. \quad (2)$$

The set of all projections of \mathbf{x} onto \mathcal{C} is defined by

$$\Pi_{\mathcal{C}}(\mathbf{x}) = \{\mathbf{y} \in \mathcal{C} \mid \|\mathbf{y} - \mathbf{x}\| = d(\mathbf{x}, \mathcal{C})\}. \quad (3)$$

It is well-known [41] that if \mathcal{C} is closed, then for any $\mathbf{x} \in \mathbb{R}^n$, $\Pi_{\mathcal{C}}(\mathbf{x})$ is non-empty². An orthogonal projection onto \mathcal{C} is defined as $\mathcal{P}_{\mathcal{C}} : \mathbb{R}^n \rightarrow \mathcal{C}$ such that $\mathcal{P}_{\mathcal{C}}(\mathbf{x})$ is chosen as an element of $\Pi_{\mathcal{C}}(\mathbf{x})$ based on a prescribed scheme (e.g., based

²In addition, if \mathcal{C} is convex, then $\Pi_{\mathcal{C}}(\mathbf{x})$ is singleton.

on lexicographic order). There exists a non-empty subset of \mathbb{R}^n such that $\Pi_{\mathcal{C}}$ is uniquely defined, given by

$$\text{singleton } \Pi_{\mathcal{C}} = \{\mathbf{x} \in \mathbb{R}^n \mid \Pi_{\mathcal{C}}(\mathbf{x}) \text{ is singleton}\}. \quad (4)$$

We can now consider the differentiability of $\mathcal{P}_{\mathcal{C}}$ over singleton $\Pi_{\mathcal{C}}$ as follows.

Definition 1 (Point-wise differentiability). *The projection $\mathcal{P}_{\mathcal{C}}$ is said to be **differentiable** at $\mathbf{x} \in \text{singleton } \Pi_{\mathcal{C}}$ if there exists $\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}) \in \mathbb{R}^{n \times n}$ such that*

$$\limsup_{\delta \rightarrow 0} \sup_{\mathbf{y} \in \Pi_{\mathcal{C}}(\mathbf{x} + \delta)} \frac{\|\mathbf{y} - \mathcal{P}_{\mathcal{C}}(\mathbf{x}) - \nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x})\delta\|}{\|\delta\|} = 0.$$

The operator $\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x})$ is said to be the derivative of $\mathcal{P}_{\mathcal{C}}$ at \mathbf{x} .

Definition 2 (Point-wise Lipschitz-continuous differentiability). *The projection $\mathcal{P}_{\mathcal{C}}$ is said to be Lipschitz-continuously differentiable at $\mathbf{x} \in \text{singleton } \Pi_{\mathcal{C}}$ if $\mathcal{P}_{\mathcal{C}}$ is differentiable at \mathbf{x} and there exist $0 < c_1(\mathbf{x}) \leq \infty$ and $0 \leq c_2(\mathbf{x}) < \infty$ such that for any $\delta \in \mathcal{B}(\mathbf{0}, c_1(\mathbf{x}))$, we have*

$$\sup_{\mathbf{y} \in \Pi_{\mathcal{C}}(\mathbf{x} + \delta)} \|\mathbf{y} - \mathcal{P}_{\mathcal{C}}(\mathbf{x}) - \nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x})\delta\| \leq c_2(\mathbf{x})\|\delta\|^2. \quad (5)$$

It is noted that the supremum in (5) implies

$$\|\mathcal{P}_{\mathcal{C}}(\mathbf{x} + \delta) - \mathcal{P}_{\mathcal{C}}(\mathbf{x}) - \nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x})\delta\| \leq c_2(\mathbf{x})\|\delta\|^2$$

holds for any choice of $\mathcal{P}_{\mathcal{C}}(\mathbf{x} + \delta)$ in $\Pi_{\mathcal{C}}(\mathbf{x} + \delta)$. Note that while $\mathcal{P}_{\mathcal{C}}(\mathbf{x})$ is uniquely defined for $\mathbf{x} \in \text{singleton } \Pi_{\mathcal{C}}$, $\mathcal{P}_{\mathcal{C}}(\mathbf{x} + \delta)$ is not since $\mathbf{x} + \delta$ may not be in $\text{singleton } \Pi_{\mathcal{C}}$.

Example 1. *Let $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| = 1\}$ be the unit sphere of dimension $n - 1$. For any $\mathbf{x} \neq \mathbf{0}$, the projection onto \mathcal{C} is uniquely given by $\mathcal{P}_{\mathcal{C}}(\mathbf{x}) = \mathbf{x}/\|\mathbf{x}\|$. For $\mathbf{x} = \mathbf{0}$, we have $\Pi_{\mathcal{C}}(\mathbf{0}) = \mathcal{C}$ and $\mathcal{P}_{\mathcal{C}}(\mathbf{0})$ can be chosen as any point on the unit sphere. In Section II of the Supplementary Material, we prove that $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at any $\mathbf{x} \in \text{singleton } \Pi_{\mathcal{C}} = \mathbb{R}^n \setminus \{\mathbf{0}\}$. In particular, for any $\mathbf{x} \neq \mathbf{0}$ and $\delta \in \mathbb{R}^n$, we have*

$$\sup_{\mathbf{y} \in \Pi_{\mathcal{C}}(\mathbf{x} + \delta)} \mathbf{y} - \frac{\mathbf{x}}{\|\mathbf{x}\|} - \mathbf{I}_n - \frac{\mathbf{x}\mathbf{x}^{\top}}{\|\mathbf{x}\|^2} \frac{\delta}{\|\mathbf{x}\|} \leq \frac{2}{\|\mathbf{x}\|^2} \|\delta\|^2. \quad (6)$$

For $\delta \neq -\mathbf{x}$, $\Pi_{\mathcal{C}}(\mathbf{x} + \delta) = \{(\mathbf{x} + \delta)/\|\mathbf{x} + \delta\|\}$ is singleton and the supremum is evaluated at only one point $\mathbf{y} = (\mathbf{x} + \delta)/\|\mathbf{x} + \delta\|$. For $\delta = -\mathbf{x}$, $\Pi_{\mathcal{C}}(\mathbf{x} + \delta) = \Pi_{\mathcal{C}}(\mathbf{0}) = \mathcal{C}$ is not singleton and the supremum is taken over the entire sphere independent of \mathbf{x} . In either case regardless the value of δ , comparing (6) with (5), we recognize the projection onto the unit sphere is Lipschitz-continuously differentiable at $\mathbf{x} \in \text{singleton } \Pi_{\mathcal{C}}$ with

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}) = \frac{1}{\|\mathbf{x}\|} \mathbf{I}_n - \frac{\mathbf{x}\mathbf{x}^{\top}}{\|\mathbf{x}\|^2},$$

$$c_1(\mathbf{x}) = \infty, \quad c_2(\mathbf{x}) = \frac{2}{\|\mathbf{x}\|^2}.$$

In 1984, Foote [42] showed that if \mathcal{C} is a C^k ($k \geq 2$) submanifold of \mathbb{R}^n , then \mathcal{C} has a neighborhood \mathcal{E} such that $\mathcal{E} \subseteq \text{singleton } \Pi_{\mathcal{C}}$ and the projection $\mathcal{P}_{\mathcal{C}}$ restricted to \mathcal{E} is a C^{k-1} mapping. Later on, Dudek and Holly [43] proved

Algorithm 1 Projected Gradient Descent (PGD)

Input: $\mathbf{A}, \mathbf{b}, \mathcal{C}, \eta, \mathbf{x}^{(0)}$

Output: $\{\mathbf{x}^{(k)}\}_{k=0}^{\infty}$

- 1: **for** $k = 0, 1, \dots$ **do**
 - 2: $\mathbf{z}_{\eta}^{(k)} = \mathbf{x}^{(k)} - \eta \mathbf{A}^{\top} \mathbf{A} \mathbf{x}^{(k)} - \mathbf{b}$
 - 3: $\mathbf{x}^{(k+1)} = \mathcal{P}_{\mathcal{C}} \mathbf{z}_{\eta}^{(k)}$
-

the derivative $\nabla \mathcal{P}_{\mathcal{C}}$ is a linear map to the tangent bundle of \mathcal{C} and more importantly, for any $\mathbf{x}^* \in \mathcal{C}$, $\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*)$ is the (linear) orthogonal projection onto the tangent space to \mathcal{C} at \mathbf{x}^* . Recently, a local version of this result has been proposed by Lewis and Malick [44]:

Proposition 1. (Rephrased from Lemma 4 in [44]) *Assume \mathcal{C} is a C^k ($k \geq 2$) manifold around a point $\mathbf{x}^* \in \mathcal{C}$. Denote the tangent space to \mathcal{C} at \mathbf{x}^* by $T_{\mathbf{x}^*}(\mathcal{C})$. Then, the set of projections $\Pi_{\mathcal{C}}$ is (locally) singleton around \mathbf{x}^* . Moreover, $\mathcal{P}_{\mathcal{C}}$ is a C^{k-1} mapping around \mathbf{x}^* and*

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathcal{P}_{T_{\mathbf{x}^*}(\mathcal{C})}, \quad (7)$$

where $\mathcal{P}_{T_{\mathbf{x}^*}(\mathcal{C})}$ is the orthogonal projection onto $T_{\mathbf{x}^*}(\mathcal{C})$.

Further works on the uniqueness and regularity of $\mathcal{P}_{\mathcal{C}}$ can also be found in [45]–[48]. We note that the assumption \mathcal{C} is a C^2 manifold around \mathbf{x}^* requires the existence of a neighborhood of \mathbf{x}^* in which $\mathcal{P}_{\mathcal{C}}$ is uniformly differentiable. Our result in this manuscript, while strongly motivated by the aforementioned results, only requires \mathcal{C} to be differentiable at two points (see Theorem 1).

C. Stationary Points of (1)

We defined the (Lipschitz-continuous) differentiability of the projection $\mathcal{P}_{\mathcal{C}}$ at a point in \mathcal{C} . We are now in position to define stationary points of (1) as those where the gradient of the objective function on the constraint set vanishes [49]:

Definition 3. *$\mathbf{x}^* \in \mathcal{C}$ is a **stationary point** of (1) if $\mathcal{P}_{\mathcal{C}}$ is differentiable at \mathbf{x}^* and*

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) \mathbf{A}^{\top} \mathbf{A} \mathbf{x}^* - \mathbf{b} = \mathbf{0}. \quad (8)$$

Assume in addition that $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at \mathbf{x}^* with constants $c_1(\mathbf{x}^*)$ and $c_2(\mathbf{x}^*)$. Then \mathbf{x}^* is called a **Lipschitz stationary point** of (1) with constants $c_1(\mathbf{x}^*)$ and $c_2(\mathbf{x}^*)$.

Similar to unconstrained optimization, stationary points in Definition 3 can be local minimizers, local maximizers, or saddle points of the constrained problem (1).

D. Projected Gradient Descent

Algorithm 1 describes the projected gradient descent algorithm for solving (1). Starting at some $\mathbf{x}^{(0)} \in \mathcal{C}$, the algorithm iteratively updates the current value by (i) taking a step in the opposite direction of the gradient and (ii) projecting the result back onto \mathcal{C} , i.e.,

$$\mathbf{x}^{(k+1)} = \mathcal{P}_{\mathcal{C}} \mathbf{x}^{(k)} - \eta \mathbf{A}^{\top} \mathbf{A} \mathbf{x}^{(k)} - \mathbf{b}, \quad (9)$$

where \mathbf{Q} is an orthogonal matrix with $\kappa(\mathbf{Q}) = 1$.

Next, we study a special case of Theorem 1 in which

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*) = \nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathcal{P}_{T_{\mathbf{x}}(\mathcal{C})} = \mathbf{U}_{\mathbf{x}} \mathbf{U}_{\mathbf{x}}^{\top}, \quad (17)$$

where $\mathbf{U}_{\mathbf{x}} \in \mathbb{R}^{n \times d}$ ($d \leq n$) is the matrix whose columns provide an orthonormal basis for the tangent space to \mathcal{C} at \mathbf{x}^* . A typical example in which (17) holds is when (i) \mathcal{C} is a C^2 d -dimensional submanifold around \mathbf{x}^* ; and (ii) $\mathbf{z}_{\eta}^* = \mathbf{x}^*$. The first condition (i) stems from Proposition 1 in order to guarantee $\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathcal{P}_{T_{\mathbf{x}}(\mathcal{C})}$. The second condition (ii) is equivalent to $\mathbf{A}^{\top}(\mathbf{A}\mathbf{x} - \mathbf{b}) = \mathbf{0}$, which means \mathbf{x}^* is also a stationary point of the unconstrained problem. Conveniently, this coincidence eliminates the task of characterizing the projection $\mathcal{P}_{\mathcal{C}}$ and its derivative $\nabla \mathcal{P}_{\mathcal{C}}$ at a point outside \mathcal{C} , which can be a challenging task in many problems.

Corollary 1. Consider the same setting as in Theorem 1 with the additional assumption that (17) holds. If $(\mathbf{A}\mathbf{U}_{\mathbf{x}})^{\top}\mathbf{A}\mathbf{U}_{\mathbf{x}}$ has full rank and

$$0 < \eta < \frac{2}{\|\mathbf{A}\mathbf{U}_{\mathbf{x}}\|_2^2}, \quad (18)$$

then Algorithm 1 with fixed step size η converges locally to \mathbf{x}^* at an asymptotic linear rate

$$\rho(\mathbf{H}) = \max\{|1 - \eta\lambda_1|, |1 - \eta\lambda_d|\}, \quad (19)$$

where λ_1 and λ_d are the largest and smallest eigenvalues of $(\mathbf{A}\mathbf{U}_{\mathbf{x}})^{\top}\mathbf{A}\mathbf{U}_{\mathbf{x}}$, respectively. The region of linear convergence is given by

$$\|\mathbf{x}^{(0)} - \mathbf{x}^*\| < \min \left\{ c_1(\mathbf{x}^*), \frac{c_1(\mathbf{z}_{\eta}^*)}{u_{\eta}}, \frac{1 - \rho(\mathbf{H})}{u_{\eta}c_2(\mathbf{x}^*) + u_{\eta}^2c_2(\mathbf{z}_{\eta}^*)} \right\}, \quad (20)$$

where u_{η} is given by (14).

The proof of Corollary 1 is given in Appendix B.

Remark 2. Recall that u_{η} defined in (14) is also the asymptotic linear rate of gradient descent for the unconstrained least squares [50], i.e.,

$$u_{\eta} = \max\{|1 - \eta\lambda_1(\mathbf{A}^{\top}\mathbf{A})|, |1 - \eta\lambda_n(\mathbf{A}^{\top}\mathbf{A})|\}.$$

Since $\mathbf{U}_{\mathbf{x}}$ is a semi-orthogonal matrix, the eigenvalues of $\mathbf{U}_{\mathbf{x}}^{\top}\mathbf{A}^{\top}\mathbf{A}\mathbf{U}_{\mathbf{x}}$ interlace with those of $\mathbf{A}^{\top}\mathbf{A}$ [51], which in turns implies $\lambda_n(\mathbf{A}^{\top}\mathbf{A}) \leq \lambda_d \leq \lambda_1 \leq \lambda_1(\mathbf{A}^{\top}\mathbf{A})$. Thus, one can show that for $\eta < 2/\|\mathbf{A}\|_2^2$,

$$\rho(\mathbf{H}) \leq u_{\eta} \leq 1, \quad (21)$$

with the equality $u_{\eta} = 1$ holding if and only if $\mathbf{A}^{\top}\mathbf{A}$ is singular. Interestingly, (21) implies the presence of the constraint in this case helps accelerate the convergence of gradient descent to \mathbf{x}^* .

B. Proof of Theorem 1

This section presents the proof of Theorem 1. Our key ideas are: (1) using the Lipschitz-continuous differentiability of $\mathcal{P}_{\mathcal{C}}$ at \mathbf{x}^* and at \mathbf{z}_{η}^* to establish a recursive relation on the error vector $\delta^{(k)} = \mathbf{x}^{(k)} - \mathbf{x}^*$, (2) performing a change of

basis $\tilde{\delta}^{(k)} = \mathbf{Q}^{-1}\delta^{(k)}$ to establish an asymptotically-linear quadratic system dynamic that upper-bounds the norm of the transformed error vector, (3) applying the result on the convergence of an asymptotically-linear quadratic difference equation in [37] to obtain the number of iterations required for $\|\tilde{\delta}^{(k)}\| \leq \tilde{\epsilon}\|\tilde{\delta}^{(0)}\|$, and (4) converting the convergence result on the transformed error $\|\tilde{\delta}^{(k)}\|$ to the convergence result on the original error $\|\delta^{(k)}\|$. In the following, we provide the complete proof, with some details deferred to the appendix.

Step 1: Let us define the error vector of Algorithm 1 as $\delta^{(k)} = \mathbf{x}^{(k)} - \mathbf{x}^*$, for $k \in \mathbb{N}$. Using this definition of the error vector, we can replace $\mathbf{x}^{(k)} = \mathbf{x}^* + \delta^{(k)}$ and $\mathbf{x}^{(k+1)} = \mathbf{x}^* + \delta^{(k+1)}$ into (9) to obtain an equivalent update on the error vector

$$\delta^{(k+1)} = \mathcal{P}_{\mathcal{C}}(\mathbf{x}^* + \delta^{(k)}) - \eta\mathbf{A}^{\top}(\mathbf{A}(\mathbf{x}^* + \delta^{(k)}) - \mathbf{b}) - \mathbf{x}^*. \quad (22)$$

Based on the definition of \mathbf{z}_{η}^* in (11) and the fact that \mathbf{x}^* is a fixed point of the algorithm (see (10)), i.e., $\mathbf{x}^* = \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*)$, we can rewrite (22) as

$$\delta^{(k+1)} = \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^* + (\mathbf{I} - \eta\mathbf{A}^{\top}\mathbf{A})\delta^{(k)}) - \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*). \quad (23)$$

We are now in position to analyze the error update as a fixed-point iteration: $\delta^{(k+1)} = \mathbf{f}(\delta^{(k)})$, where $\mathbf{f}(\delta) = \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^* + (\mathbf{I} - \eta\mathbf{A}^{\top}\mathbf{A})\delta) - \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*)$. The following lemma provides a recursive equation on the error vector that is in the form of an asymptotically-linear quadratic system dynamic:

Lemma 2. Recall $\mathbf{H} = \nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*)(\mathbf{I}_n - \eta\mathbf{A}^{\top}\mathbf{A})\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*)$. If the error vector at the k -th iteration satisfies

$$\|\delta^{(k)}\| < \min \left\{ c_1(\mathbf{x}^*), \frac{c_1(\mathbf{z}_{\eta}^*)}{u_{\eta}} \right\}, \quad (24)$$

then the error vector at the $k+1$ -th iteration satisfies

$$\delta^{(k+1)} = \mathbf{H}\delta^{(k)} + \mathbf{q}_2(\delta^{(k)}), \quad (25)$$

where $\mathbf{q}_2: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the residual such that

$$\|\mathbf{q}_2(\delta^{(k)})\| \leq \|\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*)\|_2 u_{\eta} c_2(\mathbf{x}^*) + c_2(\mathbf{z}_{\eta}^*) u_{\eta}^2 \|\delta^{(k)}\|^2. \quad (26)$$

The proof of Lemma 2 is given in Appendix C. Given the nonlinear difference equation of form (25), we proceed with characterizing the convergence of the error sequence $\{\delta^{(k)}\}_{k=0}^{\infty}$.

Remark 3. Dating back to 1964, Polyak [39] studied the convergence of nonlinear difference equations of form

$$\mathbf{a}^{(k+1)} = \mathbf{T}(\mathbf{a}^{(k)}) + \mathbf{q}(\mathbf{a}^{(k)}), \quad \text{for } k \in \mathbb{N}, \quad (27)$$

where $\mathbf{a}^{(0)} \in \mathbb{R}^n$, $\mathbf{T}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a linear operator, and $\mathbf{q}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ satisfies $\lim_{t \rightarrow 0} \sup_{\|\mathbf{a}\| \leq t} \|\mathbf{q}(\mathbf{a})\| / \|\mathbf{a}\| = 0$. The author showed that if the operator \mathbf{T} satisfies $\|\mathbf{T}^k\|_2 \leq c(\zeta)(\rho + \zeta)^k$, for some $\rho < 1$ and arbitrarily small $\zeta > 0$, then $\{\mathbf{a}^{(k)}\}_{k=0}^{\infty}$ approaches zero with sufficiently small $\|\mathbf{a}^{(0)}\|$:

$$\|\mathbf{a}^{(k)}\| \leq C(\zeta)\|\mathbf{a}^{(0)}\|(\rho + \zeta)^k. \quad (28)$$

Here $c(\zeta)$ and $C(\zeta)$ are unknown constants that could grow to infinity as $\zeta \rightarrow 0$. Applying this result to (25) with $\mathbf{a}^{(k)} = \delta^{(k)}$

and $\mathbf{T} = \mathbf{H}$, one can show that the error vector of Algorithm 1 converges to $\mathbf{0}$ with the asymptotic linear rate $\rho(\mathbf{H})$, provided that $\rho(\mathbf{H}) < 1$ and $\|\delta^{(0)}\|$ is sufficiently small. However, we note that the proof of (28) in [39] is adapted from a more general result on the stability of differential equations in [52]. This technique can not provide the precise control of the ROC and the number of iterations required to reach a certain accuracy (i.e., how small $\|\mathbf{a}^{(0)}\|$ is as well as how large the factor $C(\zeta)$ is) needed for our convergence analysis of PGD. Alternatively, we utilize our previous result in [37] that eliminates the dependence on ζ in the expression of the linear rate, at the cost of an additional assumption on the diagonalizability of \mathbf{H} .³ Additionally, our approach offers explicit expressions of the ROC and the number of required iterations (as in (13) and (16), respectively).

Step 2: Our approach for analyzing the convergence of the nonlinear difference equation (24) is to leverage the eigen-decomposition $\mathbf{H} = \mathbf{Q} \mathbf{Q}^{-1}$ and consider the transformed error vector as follows.

Lemma 3. Let $\tilde{\delta}^{(k)} = \mathbf{Q}^{-1}\delta^{(k)}$ be the transformed error vector. If (13) holds and the spectral radius of \mathbf{H} is strictly less than 1, i.e., $\rho(\mathbf{H}) < 1$, then, for all $k \in \mathbb{N}$, we have

$$\tilde{\delta}^{(k+1)} = \tilde{\delta}^{(k)} + \mathbf{q}_3(\tilde{\delta}^{(k)}), \quad (29)$$

where the residual $\mathbf{q}_3(\tilde{\delta}^{(k)}) = \mathbf{Q}^{-1}\mathbf{q}_2(\mathbf{Q}\tilde{\delta}^{(k)})$ satisfies $\|\mathbf{q}_3(\tilde{\delta}^{(k)})\| \leq (q/\|\mathbf{Q}^{-1}\|_2)\|\tilde{\delta}^{(k)}\|^2$ for q given in (15).

The proof of Lemma 3 is given in Appendix D. Taking the norms of both sides of (29) and applying the triangle inequality, we obtain

$$\|\tilde{\delta}^{(k+1)}\| \leq \rho(\mathbf{H})\|\tilde{\delta}^{(k)}\| + \frac{q}{\|\mathbf{Q}^{-1}\|_2}\|\tilde{\delta}^{(k)}\|^2. \quad (30)$$

This inequality, holding for all $k \in \mathbb{N}$, is the key to the convergence of the transformed error sequence in the next step.

Step 3: If we replace the inequality symbol in (30) by the equality symbol, then we obtain an asymptotically-linear quadratic difference equation whose convergence is studied in [37]. Indeed, the following lemma states that the norm of the transformed error vector is governed by this asymptotically-linear quadratic system dynamic:

Lemma 4. Assume the same setting as Lemma 3. Then, for any desired accuracy $0 < \tilde{\epsilon} < 1$, we have $\|\tilde{\delta}^{(k)}\| \leq \tilde{\epsilon}\|\tilde{\delta}^{(0)}\|$ for all

$$k \geq \frac{\log(1/\tilde{\epsilon})}{\log(1/\rho(\mathbf{H}))} + c_3(\rho(\mathbf{H}), \tau), \quad (31)$$

where $\tau = q\|\tilde{\delta}^{(0)}\|/\|\mathbf{Q}^{-1}\|_2/(1 - \rho(\mathbf{H})) \in (0, 1)$ and

$$c_3(\rho, \tau) = \frac{E_1 \log \frac{1}{\rho + \tau(1 - \rho)} - E_1 \log \frac{1}{\rho}}{\rho \log(1/\rho)} + \frac{1}{2\rho} \log \frac{\log(1/\rho)}{\log 1/(\rho + \tau(1 - \rho))} + 1, \quad (32)$$

for $E_1(t) = \int_t^\infty \frac{e^{-z}}{z} dz$ being the exponential integral [53].

³In particular, the bound in (16) suggests $k\mathbf{a}^{(k)} \leq Ck\mathbf{a}^{(0)}$, for constant $C = (\mathbf{Q})e^{c_3}$, which is tighter than (28).

The proof of Lemma 4 is given in Appendix E.

Step 4: Finally, we show the convergence of $\|\delta^{(k)}\|$ based on the convergence of $\|\tilde{\delta}^{(k)}\|$. From (31), substituting $\tilde{\epsilon} = \epsilon/\kappa(\mathbf{Q})$ and identifying c_3 as $c_3(\rho(\mathbf{H}), \tau)$, we obtain (16). Thus, it remains to prove that the accuracy on the transformed error vector $\|\tilde{\delta}^{(k)}\| \leq \tilde{\epsilon}\|\tilde{\delta}^{(0)}\|$ is sufficient for the accuracy on the original error vector $\|\delta^{(k)}\| \leq \epsilon\|\delta^{(0)}\|$. Indeed, given

$$\|\tilde{\delta}^{(k)}\| \leq \tilde{\epsilon}\|\tilde{\delta}^{(0)}\| = \frac{\epsilon}{\|\mathbf{Q}\|_2\|\mathbf{Q}^{-1}\|_2}\|\tilde{\delta}^{(0)}\|,$$

we have

$$\begin{aligned} \|\delta^{(k)}\| &= \|\mathbf{Q}\tilde{\delta}^{(k)}\| \leq \|\mathbf{Q}\|_2\|\tilde{\delta}^{(k)}\| \\ &\leq \|\mathbf{Q}\|_2 \frac{\epsilon}{\|\mathbf{Q}\|_2\|\mathbf{Q}^{-1}\|_2}\|\tilde{\delta}^{(0)}\| \\ &= \frac{\epsilon}{\|\mathbf{Q}^{-1}\|_2}\|\tilde{\delta}^{(0)}\| \leq \epsilon\|\delta^{(0)}\|, \end{aligned}$$

where the last inequality stems from $\|\tilde{\delta}^{(0)}\| = \|\mathbf{Q}^{-1}\delta^{(0)}\| \leq \|\mathbf{Q}^{-1}\|_2\|\delta^{(0)}\|$. This completes our proof of Theorem 1.

IV. APPLICATIONS

In this section, we demonstrate the application of our proposed framework to a collection of well-known problems in machine learning and signal processing. The constraint sets in these problems vary from as simple as an affine subspace (A) and a sphere (C) to more complex algebraic varieties such as the s -sparse vector set (B) and the low-rank matrix set (D). We consider both problems with known convergence rate results and problems for which the rate is unavailable. The former allows us to verify the correctness of our analysis against the known rate results, while for the latter numerical experiments are used to verify the rate. Additionally, we illustrate how ROC can be obtained for each problem. Due to space limitation, we restrict the illustration of our framework to the four aforementioned applications. While we believe that additional applications can be considered (see the potential applications of our framework in Section V), such applications may require a more elaborate development. Our goal in this section is to offer a recipe for analyzing the convergence of PGD for different applications using the proposed framework. Table I describes the steps we follow to obtain the asymptotic linear rate and the region of linear convergence in each application. Table II summarizes our local convergence results on the four problems presented in this section. The detailed analysis is given below.

A. Linear Equality-Constrained Least Squares

As a sanity check, we start with a simple example of the so-called linear equality-constrained least squares (LECLS)

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \quad \text{s.t. } \mathbf{C}\mathbf{x} = \mathbf{d}, \quad (33)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, $\mathbf{C} \in \mathbb{R}^{p \times n}$, and $\mathbf{d} \in \mathbb{R}^p$. In addition, we assume that $p < n$ and \mathbf{C} has linearly independent rows. The LECLS problem finds application in a wide range of areas such as linear-phase system identification [54], antenna array processing [55], and adaptive array processing [56].

TABLE I: General recipe for local convergence analysis.

Step 1:	Identify \mathbf{A} , \mathbf{b} , \mathcal{C} , and $\mathcal{P}_{\mathcal{C}}$.
Step 2:	Establish the conditions for $\mathbf{x} \in \mathcal{C}$ to be a Lipschitz stationary point of (1). In particular, (i) $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at every \mathbf{x} with $r_{\mathcal{P}_{\mathcal{C}}}(\mathbf{x})$, $c_1(\mathbf{x})$, and $c_2(\mathbf{x})$; and (ii) the stationarity equation (8) holds.
Step 3:	Establish the conditions for $\eta > 0$ such that \mathbf{x}^* is a fixed point of Algorithm 1 with step size η , i.e., $\mathbf{x}^* = \mathcal{P}_{\mathcal{C}}(\mathbf{z})$, for $\mathbf{z} = \mathbf{x}^* - \eta \mathbf{A}^{\top}(\mathbf{A}\mathbf{x}^* - \mathbf{b})$; and (ii) $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at \mathbf{z} with $r_{\mathcal{P}_{\mathcal{C}}}(\mathbf{z})$, $c_1(\mathbf{z})$, and $c_2(\mathbf{z})$.
Step 4:	Determine the asymptotic linear rate as the spectral radius of \mathbf{H} given by (12). (If $r_{\mathcal{P}_{\mathcal{C}}}(\mathbf{z}) = r_{\mathcal{P}_{\mathcal{C}}}(\mathbf{x})$, (19) can be used instead.)
Step 5:	Establish the conditions for $\rho < 1$, which guarantees local linear convergence. Thereby, combine these conditions with the previous conditions obtained from Steps 2 and 3.
Step 6:	If \mathbf{H} is diagonalizable, determine the region of linear convergence given by (13). (If $r_{\mathcal{P}_{\mathcal{C}}}(\mathbf{z}) = r_{\mathcal{P}_{\mathcal{C}}}(\mathbf{x})$, (20) can be used instead.)

While this problem can be solved efficiently using the method of Lagrange multipliers [57] or the method of weighting [58], we limit our interest to using PGD to solve (33) to demonstrate the applicability of our analysis. In the literature, this algorithm is referred to as the projected Landweber iteration [59]–[62]. While these works provide bounds on the linear convergence of PGD for different variants of linear equality-constrained problems, we have not found any closed-form expression of the asymptotic rate of linear convergence.

Step 1: In this example, \mathbf{A} and \mathbf{b} are given explicitly in (33). The constraint set \mathcal{C} is the closed convex affine subspace

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{C}\mathbf{x} = \mathbf{d}\}.$$

The orthogonal projection onto this subspace is given in a closed-form expression as $\mathcal{P}_{\mathcal{C}}(\mathbf{x}) = \mathbf{x} - \mathbf{C}^{\top}(\mathbf{C}\mathbf{C}^{\top})^{-1}(\mathbf{C}\mathbf{x} - \mathbf{d})$, for all $\mathbf{x} \in \mathbb{R}^n$ [63]. Since \mathbf{C} has full row rank, it admits a compact singular value decomposition (SVD) $\mathbf{C} = \mathbf{U}_{\mathcal{C}} \mathbf{\Sigma}_{\mathcal{C}} \mathbf{V}_{\mathcal{C}}^{\top}$, where $\mathbf{\Sigma}_{\mathcal{C}} \in \mathbb{R}^{p \times p}$ is a diagonal matrix with positive diagonal entries, $\mathbf{U}_{\mathcal{C}} \in \mathbb{R}^{p \times p}$ and $\mathbf{V}_{\mathcal{C}} \in \mathbb{R}^{n \times p}$ satisfy $\mathbf{U}_{\mathcal{C}}^{\top} \mathbf{U}_{\mathcal{C}} = \mathbf{V}_{\mathcal{C}}^{\top} \mathbf{V}_{\mathcal{C}} = \mathbf{I}_p$. Denote $\mathbf{V}_{\mathcal{C}}^{\perp} \in \mathbb{R}^{n \times (n-p)}$ the orthogonal complement of $\mathbf{V}_{\mathcal{C}}$, i.e., $\mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top} = \mathbf{I}_n - \mathbf{V}_{\mathcal{C}} \mathbf{V}_{\mathcal{C}}^{\top}$ and $(\mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{V}_{\mathcal{C}}^{\perp} = \mathbf{I}_{n-p}$. Substituting the SVD of \mathbf{C} back into the aforementioned expression of $\mathcal{P}_{\mathcal{C}}$ yields

$$\mathcal{P}_{\mathcal{C}}(\mathbf{x}) = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{x} + \tilde{\mathbf{d}}, \quad (34)$$

where $\tilde{\mathbf{d}} = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{-1} \mathbf{U}_{\mathcal{C}}^{\top} \mathbf{d} = \mathbf{C}^{\dagger} \mathbf{d}$.

Step 2: From (34), we obtain the difference between the two projections of $\mathbf{x} + \boldsymbol{\delta}$ and \mathbf{x} onto \mathcal{C} , for any $\mathbf{x}, \boldsymbol{\delta} \in \mathbb{R}^n$, as $\mathcal{P}_{\mathcal{C}}(\mathbf{x} + \boldsymbol{\delta}) - \mathcal{P}_{\mathcal{C}}(\mathbf{x}) = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \boldsymbol{\delta}$. Using Definition 1 with the note that $\Pi_{\mathcal{C}}(\mathbf{x} + \boldsymbol{\delta})$ is always singleton, we have $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at every $\mathbf{x} \in \mathbb{R}^n$ with

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}) = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top}, \quad c_1(\mathbf{x}) = \infty, \quad c_2(\mathbf{x}) = 0. \quad (35)$$

Due to the independence from \mathbf{x} , we also have $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at every $\mathbf{x}^* \in \mathcal{C}$ with

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top}, \quad c_1(\mathbf{x}^*) = \infty, \quad c_2(\mathbf{x}^*) = 0.$$

Next, substituting $\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top}$ into the stationarity equation (8) yields $\mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{A}^{\top} \mathbf{A} \mathbf{x}^* - \mathbf{b} = \mathbf{0}$. Since $\mathbf{V}_{\mathcal{C}}^{\perp} \in \mathbb{R}^{n \times (n-p)}$ has full-rank, we can omit the left most $\mathbf{V}_{\mathcal{C}}^{\perp}$ and obtain the condition for $\mathbf{x}^* \in \mathcal{C}$ to be a Lipschitz stationary point of (33) as

$$(\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp})^{\top} (\mathbf{A} \mathbf{x}^* - \mathbf{b}) = \mathbf{0}, \quad (36)$$

which means $\mathbf{A} \mathbf{x}^* - \mathbf{b}$ is in the left null space of $\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp}$.⁴

Step 3: Evaluating the projection in (34) at $\mathbf{z}_{\eta}^* = \mathbf{x}^* - \eta \mathbf{A}^{\top} (\mathbf{A} \mathbf{x}^* - \mathbf{b})$ and using the stationarity condition (36) to eliminate the term $\eta \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{A}^{\top} (\mathbf{A} \mathbf{x}^* - \mathbf{b})$, we have $\mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*) = \mathbf{x}^*$ for any $\eta > 0$. Thus, the condition in this step for \mathbf{x}^* to be a fixed point of Algorithm 1 is $\eta > 0$. In addition, substituting $\mathbf{x} = \mathbf{z}_{\eta}^*$ into (35), we obtain $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at \mathbf{z}_{η}^* with

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*) = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top}, \quad c_1(\mathbf{z}_{\eta}^*) = \infty, \quad c_2(\mathbf{z}_{\eta}^*) = 0.$$

Step 4: Since $\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_{\eta}^*) = \nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathbf{V}_{\mathcal{C}}^{\perp} (\mathbf{V}_{\mathcal{C}}^{\perp})^{\top}$, using (19), we obtain the asymptotic linear rate as

$$\rho = \max\{|1 - \eta \lambda_1|, |1 - \eta \lambda_{n-p}|\}, \quad (37)$$

where λ_1 and λ_{n-p} are the largest and smallest eigenvalues of $(\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp}$, respectively.

Step 5: From (37), we have $\rho < 1$ if and only if $(\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp}$ has full rank and $0 < \eta < 2 / \|\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp}\|_2^2$. It is noted that the latter condition is sufficient for the condition $\eta > 0$ in Step 3.

Step 6: Since $c_1(\mathbf{x}^*) = c_1(\mathbf{z}_{\eta}^*) = \infty$ and $c_2(\mathbf{x}^*) = c_2(\mathbf{z}_{\eta}^*) = 0$, the region of convergence given by (20) is the entire space \mathbb{R}^n , which implies global convergence.

Remark 4. The explicit expression of the convergence rate in (37) offers a simple method to select the optimal step size:

$$\begin{aligned} \eta_{opt} &= \underset{0 < \eta < 2 / \|\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp}\|_2^2}{\operatorname{argmin}} \max\{|1 - \eta \lambda_1|, |1 - \eta \lambda_{n-p}|\} \\ &= \frac{2}{\lambda_1 + \lambda_{n-p}}. \end{aligned} \quad (38)$$

Using $\eta = \eta_{opt}$, we obtain the optimal rate of convergence

$$\rho_{opt} = 1 - \frac{2}{\kappa((\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp}) + 1}. \quad (39)$$

As a comparison, the optimal convergence rate of gradient descent for the unconstrained problem is given by [50]

$$u_{opt} = 1 - \frac{2}{\kappa(\mathbf{A}^{\top} \mathbf{A}) + 1}.$$

Recall from Remark 2 that $\rho_{opt} \leq u_{opt}$ due to the interlacing of eigenvalues of $(\mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp})^{\top} \mathbf{A} \mathbf{V}_{\mathcal{C}}^{\perp}$ and $\mathbf{A}^{\top} \mathbf{A}$.

B. Iterative Hard Thresholding for Sparse Recovery

In compressed sensing, one would like to reconstruct a sparse signal by finding solutions to under-determined linear systems $\mathbf{A} \mathbf{x} = \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$ (for $m < n$).

⁴Here, it is interesting to note that any stationary point of (33) is a global minimizer since (33) is a convex optimization problem.

TABLE II: Summary of local convergence analysis for four problems: linear equality-constrained least squares (Sec. IV-A), sparse recovery (Sec. IV-B), least squares with a unit norm constraint (Sec. IV-C), and matrix completion (Sec. IV-D). In the second row, $\mathbf{v}^* = \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b})$. In the third row, $\mathbf{K} = (\mathbf{A}\mathbf{U}_x)^\top \mathbf{A}\mathbf{U}_x$. We refer the reader to each of the corresponding sections for further details.

Problem formulation	Condition(s) for linear convergence	Asymptotic rate of convergence	Region of convergence
$\min \frac{1}{2} \mathbf{k} \mathbf{A} \mathbf{x} \quad \mathbf{b} \mathbf{k}^2$ s.t. $\mathbf{C} \mathbf{x} = \mathbf{d}$	$\left\{ \begin{array}{l} \mathbf{K} = (\mathbf{A} \mathbf{V}_C^\top)^\top \mathbf{A} \mathbf{V}_C^\top \text{ has full rank} \\ 0 < \lambda_1(\mathbf{K}) < \lambda_2(\mathbf{K}) \end{array} \right.$	$\max \{ \lambda_1(\mathbf{K}), \lambda_2(\mathbf{K}) \}$	$\mathbf{k} \mathbf{x} \quad \mathbf{x} \mathbf{k} < 1$
$\min \frac{1}{2} \mathbf{k} \mathbf{A} \mathbf{x} \quad \mathbf{b} \mathbf{k}^2$ s.t. $\mathbf{k} \mathbf{x} \mathbf{k}_0 \leq s$	$\left\{ \begin{array}{l} \mathbf{K} = (\mathbf{A} \mathbf{S}_x)^\top \mathbf{A} \mathbf{S}_x \text{ has full rank} \\ 0 < \lambda_1(\mathbf{K}) < \min \left\{ \frac{2}{\mathbf{k} \mathbf{A} \mathbf{S}_x^\top \mathbf{k}_2^2}, \frac{\mathbf{j} \mathbf{x}_{[s]} \mathbf{j}}{\mathbf{k} \mathbf{v} \mathbf{k}_1} g \right\} \end{array} \right.$	$\max \{ \lambda_1(\mathbf{K}), \lambda_2(\mathbf{K}) \}$	$\mathbf{k} \mathbf{x} \quad \mathbf{x} \mathbf{k} < \min \left\{ \frac{\mathbf{j} \mathbf{x}_{[s]} \mathbf{j}}{\mathbf{k} \mathbf{v} \mathbf{k}_1}, \frac{\mathbf{j} \mathbf{x}_{[s]} \mathbf{j}}{2 \mathbf{k} \mathbf{I}_n} \frac{\mathbf{k} \mathbf{v} \mathbf{k}_1}{\mathbf{A}^\top \mathbf{A} \mathbf{k}_2} g \right\}$
$\min \frac{1}{2} \mathbf{k} \mathbf{A} \mathbf{x} \quad \mathbf{b} \mathbf{k}^2$ s.t. $\mathbf{k} \mathbf{x} \mathbf{k} = 1$	$\left\{ \begin{array}{l} 0 < \lambda_1(\mathbf{K}) < 1 \text{ if } \lambda_1(\mathbf{K}) < 1 \\ 0 < \lambda_1(\mathbf{K}) < \frac{2}{1 + \lambda_1(\mathbf{K})} \text{ o.t.w.} \end{array} \right.$	$\frac{1}{1 + \lambda_1(\mathbf{K})} \max \{ \lambda_1(\mathbf{K}), \lambda_2(\mathbf{K}) \}$	$\mathbf{k} \mathbf{x} \quad \mathbf{x} \mathbf{k} < \frac{1}{2(1 + \lambda_1(\mathbf{K}))}; \quad t = \frac{\mathbf{k} \mathbf{I}_n}{1 + \lambda_1(\mathbf{K})} \frac{\mathbf{A}^\top \mathbf{A} \mathbf{k}_2}{1}$
$\min \frac{1}{2} \mathbf{k} \mathbf{P}(\mathbf{X}) \quad \mathbf{M} \mathbf{k}_F^2$ s.t. $\text{rank}(\mathbf{X}) = r$	$\left\{ \begin{array}{l} \mathbf{K} = \mathbf{Q}_\Omega^\top \mathbf{S} \mathbf{S}^\top \mathbf{Q}_\Omega \text{ has full rank} \\ 0 < \lambda_1(\mathbf{K}) < \frac{2}{\mathbf{k} \mathbf{Q}_\Omega^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\Omega \mathbf{k}_2} \end{array} \right.$	$\max \{ \lambda_1(\mathbf{K}), \lambda_2(\mathbf{K}) \}$	$\mathbf{k} \mathbf{X} \quad \mathbf{X} \mathbf{k}_F < \frac{(1 - \frac{r}{2}) \mathbf{k}(\mathbf{X})}{8(1 + \frac{r}{2})}$

This problem can be formulated as an L0-norm constrained least squares:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \quad \text{s.t.} \quad \|\mathbf{x}\|_0 \leq s. \quad (40)$$

In the literature, the PGD algorithm for solving (40) is often known as iterative hard thresholding (IHT), with myriad applications in medical imaging [64], MIMO communication [65], [66], antenna arrays [67], and scene recognition [68]. The convergence of a special case of IHT in which $\|\mathbf{A}\|_2 < 1$ and $\eta = 1$ has been well-studied in [2], [3], under the restricted isometry property (RIP) assumption on \mathbf{A} . In the following, we demonstrate the application of our framework to establishing a local convergence analysis of IHT with a range of different step sizes, without requiring the RIP of \mathbf{A} .

Step 1: In this example, \mathbf{A} and \mathbf{b} are given explicitly in (40), and the constraint set \mathcal{C} is the closed non-convex set of s -sparse vectors

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\|_0 \leq s\},$$

with the projection $\mathcal{P}_\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by [2]

$$[\mathcal{P}_\mathcal{C}(\mathbf{x})]_i = \begin{cases} 0 & \text{if } |x_i| < |x_{[s]}| \\ x_i & \text{if } |x_i| \geq |x_{[s]}| \end{cases} \quad \text{for } i = 1, \dots, n, \quad (41)$$

where x_i and $x_{[s]}$ denote the i th coordinate and the s th largest (in magnitude) element of a vector $\mathbf{x} \in \mathbb{R}^n$, respectively. In the case \mathbf{x} has multiple elements with the same magnitude as $x_{[s]}$, e.g., $x_{[s]} = x_{[s+1]} > 0$, we sort these entries based on the (descending) lexicographical order so that (41) is well-defined (see [2]-p. 10).

Step 2: In contrast to the previous example, the projection here is nonlinear and non-unique since the set \mathcal{C} is a real algebraic variety but not smooth in those points in \mathbb{R}^n of sparsity strictly less than s . The smooth part of \mathcal{C} is the subset

$$\mathcal{C}_{=s} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\|_0 = s\}$$

of vectors with exactly s non-zero elements. In Supplementary Material Section III, we show that any $\mathbf{x}^* \in \Phi_{=s}$ and $\mathbf{x} \in \mathcal{B}(\mathbf{x}^*, |x_{[s]}^*|/\sqrt{2})$ share the same index set of s -largest

elements (in magnitude), denoted by $\Omega_s(\mathbf{x}^*)$.⁵ Let the indices in $\Omega_s(\mathbf{x}^*)$ be $i_1 \leq \dots \leq i_s$ and $\mathbf{S}_x = [\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_s}] \in \mathbb{R}^{n \times s}$. Then, we have $(\mathbf{S}_x)^\top \mathbf{S}_x = \mathbf{I}_s$ and

$$\mathcal{P}_\mathcal{C}(\mathbf{x}) = \mathbf{S}_x \mathbf{S}_x^\top \mathbf{x}, \quad \forall \mathbf{x} \in \mathcal{B}(\mathbf{x}^*, |x_{[s]}^*|/\sqrt{2}). \quad (42)$$

By Definition 1, we obtain $\mathcal{P}_\mathcal{C}$ is Lipschitz-continuously differentiable at any $\mathbf{x}^* \in \Phi_{=s}$ with

$$\nabla \mathcal{P}_\mathcal{C}(\mathbf{x}^*) = \mathbf{S}_x \mathbf{S}_x^\top, \quad c_1(\mathbf{x}^*) = \frac{1}{\sqrt{2}} |x_{[s]}^*|, \quad c_2(\mathbf{x}^*) = 0.$$

Similar to the previous example, the stationarity equation for $\mathbf{x}^* \in \mathcal{C}_{=s}$ is given by

$$(\mathbf{A} \mathbf{S}_x)^\top (\mathbf{A} \mathbf{x}^* - \mathbf{b}) = \mathbf{0}. \quad (43)$$

Thus, we obtain the conditions for $\mathbf{x}^* \in \mathcal{C}$ to be a Lipschitz stationary point of (40) are $\mathbf{x}^* \in \mathcal{C}_{=s}$ and the vector $\mathbf{v}^* = \mathbf{A}^\top (\mathbf{A} \mathbf{x}^* - \mathbf{b})$ satisfies $v_i^* = 0$ for all $i \in \Omega_s(\mathbf{x}^*)$.

Step 3: First, following a similar approach to that in [2], we show that the condition in this step for \mathbf{x}^* to be a fixed point of Algorithm 1 is

$$0 < \eta < \frac{|x_{[s]}^*|}{\|\mathbf{v}^*\|_\infty}. \quad (44)$$

Since $v_i^* = 0$ for all $i \in \Omega_s(\mathbf{x}^*)$, we have $\mathbf{z}_\eta^* = \mathbf{x}^* - \eta \mathbf{v}^*$ satisfies $(z_\eta^*)_i = x_i^*$ for all $i \in \Omega_s(\mathbf{x}^*)$. Moreover, for any indices $i \in \Omega_s(\mathbf{x}^*)$ and $j \in \{1, \dots, n\} \setminus \Omega_s(\mathbf{x}^*)$, we have

$$\begin{aligned} |(z_\eta^*)_j| &= |x_j^* - \eta v_j^*| = \eta |v_j^*| \\ &< \frac{|x_{[s]}^*|}{\|\mathbf{v}^*\|_\infty} |v_j^*| \leq |x_{[s]}^*| \leq |x_i^*| = |(z_\eta^*)_i|, \end{aligned}$$

where the second inequality stems from $|v_j^*| \leq \|\mathbf{v}^*\|_\infty$. Therefore, $\Omega_s(\mathbf{x}^*)$ contains the s -largest (in magnitude) elements of \mathbf{z}_η^* , and hence, $\mathbf{x}^* = \mathcal{P}_\mathcal{C}(\mathbf{z}_\eta^*)$.

Second, we consider the Lipschitz-continuously differentiability of $\mathcal{P}_\mathcal{C}$ at \mathbf{z}_η^* . Given η in (44), by the same argument as in Supplementary Material Section III, one can show that every point in $\mathcal{B}(\mathbf{z}_\eta^*, (|(z_\eta^*)_{[s]}| - |(z_\eta^*)_{[s+1]}|)/\sqrt{2})$ shares the

⁵It is interesting to note that $|x_{[s]}^*|/\sqrt{2}$ is the largest possible radius. A counter-example is also constructed in Supplementary Material Section III.

same index set of s -largest elements (in magnitude) with \mathbf{z}_η^* , which is $\Omega_s(\mathbf{x}^*)$. Here, we note that $|(z_\eta^*)_{[s]}| - |(z_\eta^*)_{[s+1]}| = |x_{[s]}^*| - \eta \|\mathbf{v}^*\|_\infty$. Thus, we obtain \mathcal{P}_C is Lipschitz-continuously differentiable at \mathbf{z}_η^* with

$$\begin{aligned} \nabla \mathcal{P}_C(\mathbf{z}_\eta^*) &= \mathbf{S}_x \mathbf{S}_x^\top, \\ c_1(\mathbf{z}_\eta^*) &= \frac{1}{\sqrt{2}} |x_{[s]}^*| - \eta \|\mathbf{v}^*\|_\infty, \quad c_2(\mathbf{z}_\eta^*) = 0. \end{aligned}$$

Step 4: Since $\nabla \mathcal{P}_C(\mathbf{z}_\eta^*) = \nabla \mathcal{P}_C(\mathbf{x}^*) = \mathbf{S}_x \mathbf{S}_x^\top$, using (19), we obtain the asymptotic linear rate as

$$\rho = \max\{|1 - \eta\lambda_1|, |1 - \eta\lambda_s|\}. \quad (45)$$

where λ_1 and λ_s are the largest and smallest eigenvalues of $(\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x$, respectively.

Step 5: From (45), $\rho < 1$ if and only if $(\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x$ has full rank and

$$0 < \eta < \frac{2}{\|\mathbf{A}\mathbf{S}_x\|_2}. \quad (46)$$

Combining (44) and (46) yields the condition on the step size

$$0 < \eta < \min \left\{ \frac{2}{\|\mathbf{A}\mathbf{S}_x\|_2}, \frac{|x_{[s]}^*|}{\|\mathbf{v}^*\|_\infty} \right\}. \quad (47)$$

Here, we note that the condition $(\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x$ has full rank is related to the restricted isometry property (RIP) assumption on \mathbf{A} : $(1 - \delta_s) \|\mathbf{x}\|^2 \leq \|\mathbf{A}\mathbf{x}\|^2 \leq (1 + \delta_s) \|\mathbf{x}\|^2$, for $\delta_s \in (0, 1)$ and any s -sparse vector $\mathbf{x} \in \mathbb{R}^n$ [23]. In the reduced-form, we can rewrite the RIP assumption as

$$0 < (1 - \delta_s) \|\mathbf{y}\|^2 \leq \|\mathbf{A}\mathbf{S}_x \mathbf{y}\|^2 \leq (1 + \delta_s) \|\mathbf{y}\|^2, \quad (48)$$

for any $\mathbf{y} \in \mathbb{R}^s$ and any selection matrix $\mathbf{S} \in \mathbb{R}^{n \times s}$ obtained by randomly choosing s columns from the $n \times n$ identity matrix. Substituting $\mathbf{S} = \mathbf{S}_x$ into (48), we obtain $(\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x$ has full rank.

Step 6: Recall that $c_2(\mathbf{x}^*) = c_2(\mathbf{z}_\eta^*) = 0$. From (20), the region of convergence is given by

$$\|\mathbf{x} - \mathbf{x}^*\| < \min \left\{ \frac{|x_{[s]}^*|}{\sqrt{2}}, \frac{|x_{[s]}^*| - \eta \|\mathbf{v}^*\|_\infty}{\sqrt{2} \|\mathbf{I}_n - \eta \mathbf{A}^\top \mathbf{A}\|_2} \right\}. \quad (49)$$

Remark 5. Similar to (38) and (39), the optimal step size and the optimal convergence rate are given by

$$\begin{aligned} \eta_{opt} &= \frac{2}{\lambda_1((\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x) + \lambda_s((\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x)}, \\ \rho_{opt} &= 1 - \frac{2}{\kappa((\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x) + 1}. \end{aligned} \quad (50)$$

We consider the following numerical experiment to verify the analytical rate in (45). We start by generating \mathbf{A} , \mathbf{x}^* , and \mathbf{b} as follows. First, we sample an 200×300 sensing matrix \mathbf{A} with *i.i.d* Gaussian distributed entries $\mathcal{N}(0, 1/200)$.⁶ Next, we create a 10-sparse solution \mathbf{x}^* by randomly selecting 10 coordinates and assigning non-zero values to them based on *i.i.d* normal distribution $\mathcal{N}(0, 1)$. Finally, we set $\mathbf{b} = \mathbf{A}\mathbf{x}^*$. We apply PGD with different step sizes (listed in Fig. 2) including η_{opt} in (50) and record the value of $\|\mathbf{x}^{(k)} - \mathbf{x}^*\|$ as a function

⁶Note that such random matrix is shown to satisfy the RIP constraint [23].




Fig. 2: (Log-scale) plot of the distance between the current iterate and the local minimizer of the sparse recovery problem as a function of the number of iterations. Each solid line corresponds to PGD with a different fixed step size. Each dashed line represents the respective exponential bound ρ^k up to a constant, where the theoretical rate ρ is given by (45). In the experiment, we select $m = 200$, $n = 300$, and $s = 10$. The optimal step size $\eta_{opt} = 0.97273$ is computed by (50), with the corresponding optimal rate $\rho_{opt} = 0.3613$.

of k . In Fig. 2, the aforementioned curves are presented along with their analytic bounds given by ρ^k (up to a constant). The match in the slope between the analytic rate curve and the empirical rate curve verifies the analytic rate predicts accurately the asymptotic rate obtained empirically.

Remark 6. In Supplementary Material Section III, we further show that any stationary point \mathbf{x}^* must be a local minimum of (40). Moreover, the condition $(\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x$ has full rank in Step 5 implies \mathbf{x}^* is a **strict** local minimum of (40). Finally, it is interesting to note that in [2], the authors assume $\|\mathbf{A}\|_2 < 1$ and select $\eta = 1$. With these assumptions, the rate in (45) simplifies to $\rho = 1 - \lambda_s((\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x) = \mathbf{I}_s - (\mathbf{A}\mathbf{S}_x)^\top \mathbf{A}\mathbf{S}_x$, which is consistent with Eqn. (3.9) in [2].

C. Least Squares with the Unit Norm Constraint

A common constraint that arises in regularization methods for ill-posed problems is the spherical constraint [9], [69], [70]. In particular, we consider the following optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \quad \text{s.t. } \|\mathbf{x}\| = 1, \quad (51)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$.

Step 1: In this example, \mathbf{A} and \mathbf{b} are given explicitly in (51), and the constraint set \mathcal{C} is the closed non-convex sphere

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| = 1\},$$

with the projection $\mathcal{P}_C : \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by

$$\mathcal{P}_C(\mathbf{x}) = \begin{cases} \frac{\mathbf{x}}{\|\mathbf{x}\|} & \text{if } \mathbf{x} \neq \mathbf{0}, \\ \mathbf{e}_1 & \text{if } \mathbf{x} = \mathbf{0}. \end{cases} \quad (52)$$

Step 2: In Example 1, we showed that the projection onto the unit sphere is Lipschitz-continuously differentiable at any $\mathbf{x} \neq \mathbf{0}$. Since $\mathbf{0} \notin \mathcal{C}$, we have $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at every $\mathbf{x}^* \in \mathcal{C}$ with

$$\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathbf{I}_n - \mathbf{x}^*(\mathbf{x}^*)^\top, \quad c_1(\mathbf{x}^*) = \infty, \quad c_2(\mathbf{x}^*) = 2.$$

In addition, substituting $\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{x}^*) = \mathbf{I}_n - \mathbf{x}^*(\mathbf{x}^*)^\top$ into the stationarity equation (8) yields $\mathbf{I}_n - \mathbf{x}^*(\mathbf{x}^*)^\top \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}) = \mathbf{0}$. Equivalently, we have

$$\mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}) = \gamma \mathbf{x}^*, \quad (53)$$

where $\gamma = (\mathbf{x}^*)^\top \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b})$ is the Lagrange multiplier at \mathbf{x}^* (see Lemma 1 in [33]). Thus, we obtain the condition for $\mathbf{x}^* \in \mathcal{C}$ to be a Lipschitz stationary point of (51) is \mathbf{x}^* and $\mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b})$ are collinear.

Step 3: First, the necessary condition for $\mathcal{P}_{\mathcal{C}}$ to be Lipschitz-continuously differentiable at \mathbf{z}_η^* is $\mathbf{z}_\eta^* \in \text{singleton } \Pi_{\mathcal{C}}$, i.e., $\mathbf{z}_\eta^* \neq \mathbf{0}$. From (53), we have $\mathbf{z}_\eta^* = \mathbf{x}^* - \eta \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}) = (1 - \eta\gamma)\mathbf{x}^*$. Hence, $\mathbf{z}_\eta^* \neq \mathbf{0}$ is equivalent to $1 - \eta\gamma \neq 0$. Now the projection $\mathcal{P}_{\mathcal{C}}$ at $\mathbf{z}_\eta^* \neq \mathbf{0}$ is given by

$$\mathcal{P}_{\mathcal{C}}(\mathbf{z}_\eta^*) = \frac{(1 - \eta\gamma)\mathbf{x}^*}{\|(1 - \eta\gamma)\mathbf{x}^*\|} = \frac{1 - \eta\gamma}{|1 - \eta\gamma|} \mathbf{x}^*,$$

which implies $\mathbf{x}^* = \mathcal{P}_{\mathcal{C}}(\mathbf{z}_\eta^*)$ if and only if $1 - \eta\gamma > 0$. Thus, we obtain the condition for $\eta > 0$ such that \mathbf{x}^* is a fixed point of Algorithm 1 and $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at \mathbf{z}_η^* is $1 - \eta\gamma > 0$, which is equivalent to

$$\begin{cases} \eta \in (0, \infty) & \text{if } \gamma \leq 0, \\ \eta \in (0, \frac{1}{\gamma}) & \text{if } \gamma > 0. \end{cases} \quad (54)$$

Second, it follows from (6) that $\mathcal{P}_{\mathcal{C}}$ is Lipschitz-continuously differentiable at \mathbf{z}_η^* with

$$\begin{aligned} \nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_\eta^*) &= \frac{1}{\|\mathbf{z}_\eta^*\|} \mathbf{I}_n - \frac{\mathbf{z}_\eta^*(\mathbf{z}_\eta^*)^\top}{\|\mathbf{z}_\eta^*\|^2} = \frac{\mathbf{I}_n - \mathbf{x}^*(\mathbf{x}^*)^\top}{1 - \eta\gamma}, \\ c_1(\mathbf{z}_\eta^*) &= \infty, \quad c_2(\mathbf{z}_\eta^*) = \frac{2}{\|\mathbf{z}_\eta^*\|^2} = \frac{2}{(1 - \eta\gamma)^2}. \end{aligned}$$

Step 4: Denote $\mathbf{P}_{\mathbf{x}}^\perp = \mathbf{I}_n - \mathbf{x}^*(\mathbf{x}^*)^\top$. From (12), the asymptotic linear rate is given by

$$\begin{aligned} \rho &= \rho \frac{1}{1 - \eta\gamma} \mathbf{P}_{\mathbf{x}}^\perp (\mathbf{I}_n - \eta \mathbf{A}^\top \mathbf{A}) \mathbf{P}_{\mathbf{x}}^\perp \\ &= \frac{1}{1 - \eta\gamma} \|\mathbf{P}_{\mathbf{x}}^\perp (\mathbf{I}_n - \eta \mathbf{A}^\top \mathbf{A}) \mathbf{P}_{\mathbf{x}}^\perp\|_2. \end{aligned}$$

Let $\mathbf{P}_{\mathbf{x}}^\perp = \mathbf{U}_{\mathbf{x}} \mathbf{U}_{\mathbf{x}}^\top$, where $\mathbf{U}_{\mathbf{x}} \in \mathbb{R}^{n \times (n-1)}$ is a semi-orthogonal matrix whose columns provide a basis for the null space of \mathbf{x}^* . Then, following the same derivation as in the proof of Corollary 1, we obtain

$$\rho = \frac{1}{1 - \eta\gamma} \max\{|1 - \eta\lambda_1|, |1 - \eta\lambda_{n-1}|\}, \quad (55)$$

where λ_1 and λ_{n-1} are the largest and smallest eigenvalues of $(\mathbf{A}\mathbf{U}_{\mathbf{x}})^\top \mathbf{A}\mathbf{U}_{\mathbf{x}}$, respectively.

Step 5: Since $|1 - \eta\lambda|/(1 - \eta\gamma) < 1$ is equivalent to $\eta\gamma - 1 < 1 - \eta\lambda < 1 - \eta\gamma$, we have $\rho < 1$ if and only if

$$\gamma < \lambda_{n-1} \quad (56)$$

and

$$\eta \gamma + \lambda_1 < 2. \quad (57)$$

Similar to (54), the inequality in (57) can be rewritten as

$$\begin{cases} \eta \in (0, \infty) & \text{if } \gamma \leq -\lambda_1, \\ \eta \in (0, \frac{2}{\gamma + \lambda_1}) & \text{if } \gamma > -\lambda_1. \end{cases}$$

Finally, we note that conditions (56) and (57) together imply the condition $1 - \eta\gamma > 0$ in Step 3 since $2\eta\gamma < \eta(\gamma + \lambda_{n-1}) \leq \eta(\gamma + \lambda_1) < 2$.

Step 6: To determine the region of linear convergence, we first recall that $c_1(\mathbf{x}^*) = c_1(\mathbf{z}_\eta^*) = \infty$. Second, we have

$$\|\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_\eta^*)\|_2 = \frac{\mathbf{I}_n - \mathbf{x}^*(\mathbf{x}^*)^\top}{1 - \eta\gamma} = \frac{1}{1 - \eta\gamma}.$$

Third, since $\mathbf{H} = \mathbf{P}_{\mathbf{x}}^\perp (\mathbf{I}_n - \eta \mathbf{A}^\top \mathbf{A}) \mathbf{P}_{\mathbf{x}}^\perp / (1 - \eta\gamma)$ is symmetric, one can choose \mathbf{Q} in the eigendecomposition $\mathbf{H} = \mathbf{Q} \mathbf{Q}^{-1}$ to be orthogonal, with $\kappa(\mathbf{Q}) = 1$. Thus, from (13), we obtain the region of linear convergence as

$$\|\mathbf{x} - \mathbf{x}^*\| \leq \frac{1 - \rho}{2(t^2 + t)}, \quad (58)$$

where $t = \|\mathbf{I}_n - \eta \mathbf{A}^\top \mathbf{A}\|_2 / (1 - \eta\gamma)$.

Remark 7. The local linear rate in (55) matches the rate provided by Theorem 1 in [33]. Compared to the setting in [33], here we consider a special case of the quadratic that is convex (and hence, $\lambda_d \geq 0$). By minimizing the rate in (55) over η , we also obtain the same optimal rate of linear convergence given by Lemma 5 in [33]:

$$\rho_{\text{opt}} = \frac{\lambda_1 - \lambda_{n-1}}{\lambda_1 + \lambda_{n-1} - 2\gamma} \text{ with } \eta_{\text{opt}} = \frac{2}{\lambda_1 + \lambda_{n-1}}.$$

Interestingly, condition (56) implies \mathbf{x}^* is a strict local minimum of (51) (see Lemma 2 in [33]). Since $\rho < 1$ is one of the conditions in Theorem 1, our analysis requires \mathbf{x}^* to be a strict local minimum of (51) in order to obtain linear convergence. Finally, our framework provides the region of linear convergence in (58), which is not given in [33].

D. Matrix Completion

1) *Background:* The last application is an application of our framework to the matrix case. In matrix completion [71], given a rank- r matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$ (for $1 \leq r \leq \min\{m, n\}$) with a set of its observed entries indexed by Ω , of cardinality $0 < s < mn$, we wish to recover the unknown entries of \mathbf{M} in the complement set $\bar{\Omega}$ by solving the following optimization:

$$\min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \frac{1}{2} \|\mathcal{P}_{\Omega}(\mathbf{X} - \mathbf{M})\|_F^2 \text{ s.t. } \text{rank}(\mathbf{X}) \leq r, \quad (59)$$

where $\mathcal{P}_{\Omega} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ is the orthogonal projection onto the set of $m \times n$ matrices supported in Ω , i.e.,

$$[\mathcal{P}_{\Omega}(\mathbf{X})]_{ij} = \begin{cases} X_{ij} & \text{if } (i, j) \in \Omega, \\ 0 & \text{if } (i, j) \notin \Omega. \end{cases}$$

It is noted that while \mathbf{M} is unknown, the projection $\mathcal{P}_{\Omega}(\mathbf{M})$ is unambiguously determined by the observed entries in \mathbf{M} .

In the literature, the PGD algorithm for solving (59) is also known as the Singular Value Projection (SVP) algorithm for matrix completion [12], [13], [72], [73], with the update

$$\mathbf{X}^{(k+1)} = \mathcal{P}_{\mathcal{M}_{\leq r}}(\mathbf{X}^{(k)} - \eta \mathcal{P}_{\Omega}(\mathbf{X}^{(k)} - \mathbf{M})).$$

Here, $\mathcal{M}_{\leq r}$ is the set of matrices of rank at most r , i.e.,

$$\mathcal{M}_{\leq r} = \{\mathbf{X} \in \mathbb{R}^{m \times n} \mid \text{rank}(\mathbf{X}) \leq r\}.$$

In addition, the orthogonal projection $\mathcal{P}_{\mathcal{M}_{\leq r}}: \mathbb{R}^{m \times n} \rightarrow \mathcal{M}_{\leq r}$ is defined by Eckart–Young–Mirsky theorem [74] as follows. Let $\text{SVD}(\mathbf{X})$ be the set of all triples (\mathbf{U}, \mathbf{V}) such that $\mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$ and

$$\begin{aligned} &= \text{diag}(\sigma_1(\mathbf{X}), \dots, \sigma_n(\mathbf{X})), \\ &\mathbf{U} \in \mathbb{R}^{m \times n}, \mathbf{V} \in \mathbb{R}^{n \times n} : \mathbf{U}^\top \mathbf{U} = \mathbf{V}^\top \mathbf{V} = \mathbf{I}_n. \end{aligned}$$

Denote $\mathbf{u}_i(\mathbf{X})$ and $\mathbf{v}_i(\mathbf{X})$ the i th columns of \mathbf{U} and \mathbf{V} , respectively. Then, the set of all projections of \mathbf{X} onto $\mathcal{M}_{\leq r}$ is given by

$$\Pi_{\mathcal{M}_{\leq r}}(\mathbf{X}) = \bigcap_{i=1}^r \sigma_i(\mathbf{X}) \mathbf{u}_i(\mathbf{X}) \mathbf{v}_i(\mathbf{X})^\top \circ \left((\mathbf{U}, \mathbf{V}) \in \text{SVD}(\mathbf{X}) \right). \quad (60)$$

The set $\Pi_{\mathcal{M}_{\leq r}}(\mathbf{X})$ is singleton if and only if $\sigma_r(\mathbf{X}) = 0$ or $\sigma_r(\mathbf{X}) > \sigma_{r+1}(\mathbf{X})$. In the case $\Pi_{\mathcal{M}_{\leq r}}(\mathbf{X})$ has multiple elements, we define $\mathcal{P}_{\mathcal{M}_{\leq r}}(\mathbf{X})$ as the greatest element in $\Pi_{\mathcal{M}_{\leq r}}(\mathbf{X})$ based on the lexicographical order. We re-emphasize that our subsequent analysis holds independently of this choice.

In differential geometry, it is well-known that $\mathcal{M}_{\leq r}$ is a closed set of $\mathbb{R}^{m \times n}$ but non-smooth in those points of rank strictly less than r [75]. Similar to sparse recovery, the smooth part of $\mathcal{M}_{\leq r}$ is the set of matrices of fixed rank r :

$$\mathcal{M}_{=r} = \{\mathbf{X} \in \mathbb{R}^{m \times n} \mid \text{rank}(\mathbf{X}) = r\}.$$

At any $\mathbf{X}^* \in \mathcal{M}_{=r}$, it is shown [38] that derivative of $\mathcal{P}_{\mathcal{M}_{\leq r}}$ is a linear mapping from $\mathbb{R}^{m \times n}$ to $\mathbb{R}^{m \times n}$ satisfying

$$\nabla \mathcal{P}_{\mathcal{M}_{\leq r}}(\mathbf{X}^*)(\delta) = -\mathbf{P}_{\mathbf{U}_\perp} \mathbf{P}_{\mathbf{V}_\perp} \delta, \quad (61)$$

where $\mathbf{P}_{\mathbf{U}_\perp}$ and $\mathbf{P}_{\mathbf{V}_\perp}$ are the projections onto the left and right null spaces of \mathbf{X}^* , respectively. More importantly, for any $\delta \in \mathbb{R}^{m \times n}$, Theorem 3 in [38] asserts that

$$\begin{aligned} &\sup_{\mathbf{Y} \in \Pi_{\mathcal{M}_{\leq r}}(\mathbf{X} + \delta)} \|\mathbf{Y} - \mathbf{X}^* - \nabla \mathcal{P}_{\mathcal{M}_{\leq r}}(\mathbf{X}^*)(\delta)\|_F \\ &\leq \frac{4(1 + \sqrt{2})}{\sigma_r(\mathbf{X}^*)} \|\delta\|_F^2. \end{aligned} \quad (62)$$

2) *Vectorized version of matrix completion:* To apply our proposed framework to matrix completion, we consider a vectorized version of (59) as follows. Slightly extending the notation, we denote $\mathcal{C} = \{\text{vec}(\mathbf{X}) \mid \mathbf{X} \in \mathcal{M}_{\leq r}\}$ and $\text{vec}(\Omega) = \{(j-1)m + i \mid (i, j) \in \Omega\}$ with s distinct elements $1 \leq i_1 < \dots < i_s \leq mn$. Let $\mathbf{S}_\Omega = [\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_s}] \in \mathbb{R}^{mn \times s}$ be the selection matrix satisfying

$$\begin{aligned} &\mathbf{S}_\Omega^\top \mathbf{S}_\Omega = \mathbf{I}_s, \\ &\text{vec} \mathcal{P}_\Omega(\mathbf{X}) = \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \text{vec}(\mathbf{X}). \end{aligned}$$

Then, problem (59) can be represented as

$$\min_{\mathbf{x} \in \mathbb{R}^{mn}} \frac{1}{2} \|\mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{x} - \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \text{vec}(\mathbf{M})\|^2 \text{ s.t. } \mathbf{x} \in \mathcal{C}.$$

Step 1: In this vectorized version of matrix completion, we have $\mathbf{A} = \mathbf{S}_\Omega \mathbf{S}_\Omega^\top$, $\mathbf{b} = \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \text{vec}(\mathbf{M})$, and \mathcal{C} is a closed non-convex set. For any vector $\mathbf{x} \in \mathbb{R}^{mn}$, let $\mathbf{X} = \text{vec}^{-1}(\mathbf{x})$ with $\mathcal{P}_{\mathcal{M}_{\leq r}}(\mathbf{X}) = \sum_{i=1}^r \sigma_i(\mathbf{X}) \mathbf{u}_i(\mathbf{X}) \mathbf{v}_i(\mathbf{X})^\top$, for some $(\mathbf{U}, \mathbf{V}) \in \text{SVD}(\mathbf{X})$. The projection $\mathcal{P}_\mathcal{C}$ is given by $\mathcal{P}_\mathcal{C}(\mathbf{x}) = \text{vec}(\sum_{i=1}^r \sigma_i(\mathbf{X}) \mathbf{u}_i(\mathbf{X}) \mathbf{v}_i(\mathbf{X})^\top)$. Using the fact that $\text{vec}(\mathbf{u}\mathbf{v}^\top) = \mathbf{v} \otimes \mathbf{u}$, for any vectors \mathbf{u} and \mathbf{v} of compatible dimensions, $\mathcal{P}_\mathcal{C}$ can then be represented as

$$\mathcal{P}_\mathcal{C}(\mathbf{x}) = \sum_{i=1}^r \sigma_i(\mathbf{X}) \mathbf{v}_i(\mathbf{X}) \otimes \mathbf{u}_i(\mathbf{X}). \quad (63)$$

Step 2: In the following, we show that $\mathcal{P}_\mathcal{C}$ is Lipschitz-continuously differentiable at any point in the set

$$\mathcal{C}_{=r} = \{\text{vec}(\mathbf{X}) \mid \mathbf{X} \in \mathcal{M}_{=r}\}.$$

In particular, for any $\mathbf{x}^* \in \mathcal{C}_{=r}$, we prove that $\mathcal{P}_\mathcal{C}$ is Lipschitz-continuously differentiable at \mathbf{x}^* with

$$\begin{aligned} \nabla \mathcal{P}_\mathcal{C}(\mathbf{x}^*) &= \mathbf{P}_{\mathbf{U}_\perp}^\perp \mathbf{P}_{\mathbf{V}_\perp}, \\ c_1(\mathbf{x}^*) &= \infty, \quad c_2(\mathbf{x}^*) = \frac{4(1 + \sqrt{2})}{\sigma_r(\mathbf{X}^*)}, \end{aligned}$$

where $\mathbf{X}^* = \text{vec}^{-1}(\mathbf{x}^*)$. Indeed, the constants $c_1(\mathbf{x}^*)$ and $c_2(\mathbf{x}^*)$ are obtained from the matrix inequality form (62). Regarding $\nabla \mathcal{P}_\mathcal{C}(\mathbf{x}^*)$, let $\mathbf{P}_{\mathbf{U}_\perp}$ and $\mathbf{P}_{\mathbf{V}_\perp}$ be the projections onto the left and right null spaces of \mathbf{X}^* , respectively. Denote $\mathbf{P}_{\mathbf{U}_\perp}^\perp \mathbf{P}_{\mathbf{V}_\perp} = \mathbf{P}_{\mathbf{V}_\perp} \otimes \mathbf{P}_{\mathbf{U}_\perp}$ and $\mathbf{P}_{\mathbf{U}_\perp}^\perp \mathbf{P}_{\mathbf{V}_\perp}^\perp = \mathbf{I}_{mn} - \mathbf{P}_{\mathbf{U}_\perp} \mathbf{P}_{\mathbf{V}_\perp}$. Since $\text{vec}(\mathbf{A}\mathbf{B}\mathbf{C}) = (\mathbf{C}^\top \otimes \mathbf{A}) \text{vec}(\mathbf{B})$, for any matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} of compatible dimensions, (61) can be vectorized to obtain $\nabla \mathcal{P}_\mathcal{C}(\mathbf{x}^*)(\delta) = (\mathbf{I}_{mn} - \mathbf{P}_{\mathbf{V}_\perp} \otimes \mathbf{P}_{\mathbf{U}_\perp}) \delta = \mathbf{P}_{\mathbf{U}_\perp}^\perp \mathbf{P}_{\mathbf{V}_\perp} \delta$ for any $\delta \in \mathbb{R}^{mn}$.

Next, the stationarity condition (8) can be represented using $\nabla \mathcal{P}_\mathcal{C}(\mathbf{x}^*)(\delta) = \mathbf{P}_{\mathbf{U}_\perp}^\perp \mathbf{P}_{\mathbf{V}_\perp} \delta$ as $\mathbf{P}_{\mathbf{U}_\perp}^\perp \mathbf{P}_{\mathbf{V}_\perp} \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{x}^* - \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \text{vec}(\mathbf{M}) = 0$. Denote $\mathbf{Q}_\perp \in \mathbb{R}^{mn \times r(m+n-r)}$ the matrix satisfying $\mathbf{Q}_\perp^\top \mathbf{Q}_\perp = \mathbf{I}_{r(m+n-r)}$ and $\mathbf{Q}_\perp \mathbf{Q}_\perp^\top = \mathbf{P}_{\mathbf{U}_\perp}^\perp \mathbf{P}_{\mathbf{V}_\perp}$. Then, we obtain the conditions for \mathbf{x}^* to be a Lipschitz stationary point of (59) are $\mathbf{x}^* \in \mathcal{C}_{=r}$ and

$$\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{x}^* - \text{vec}(\mathbf{M}) = 0. \quad (64)$$

Step 3: The stationarity condition (64) leads to two cases. The first case is when $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega$ has full (row-)rank and hence,

$$\mathbf{S}_\Omega^\top \mathbf{x}^* - \text{vec}(\mathbf{M}) = 0. \quad (65)$$

In matrix form, (65) can be rewritten as $\mathcal{P}_\Omega(\mathbf{X}^*) = \mathcal{P}_\Omega(\mathbf{M})$, which implies \mathbf{X}^* is a global minimizer of (59). Interestingly, this case enjoys the special setting considered in Corollary 1 as

$$\mathbf{z}_\eta^* = \mathbf{x}^* - \eta \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{x}^* - \text{vec}(\mathbf{M}) = \mathbf{x}^*, \quad (66)$$

for any $\eta > 0$. In the second case, if $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega$ has rank strictly less than $r(m+n-r)$, then $\mathbf{S}_\Omega^\top \mathbf{x}^* - \text{vec}(\mathbf{M})$ may not be 0 (e.g., a non-zero right singular vector of $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega$). This implies $\mathbf{z}_\eta^* \neq \mathbf{x}^*$ and one needs to characterize the Lipschitz-continuous differentiability of the projection $\mathcal{P}_\mathcal{C}$ onto the set of

low-rank matrices at \mathbf{z}_η^* that may not have exact rank r . While the derivative of \mathcal{P}_C at a matrix with rank greater than r has been studied in [38], [76], it requires complete development of the error bound on the first-order expansion of this operator to obtain the constants $c_1(\mathbf{z}_\eta^*)$ and $c_2(\mathbf{z}_\eta^*)$. For the purpose of demonstration, we restrict our subsequent analysis to the first case when $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega$ has full (row-)rank. Since $\mathbf{z}_\eta^* = \mathbf{x}^*$ in this case, \mathcal{P}_C is Lipschitz-continuously differentiable at \mathbf{z}_η^* with

$$\nabla \mathcal{P}_C(\mathbf{z}_\eta^*) = \mathbf{P}_{\mathbf{U}_\perp^\top \mathbf{V}_\perp}^\perp, \quad c_1(\mathbf{z}_\eta^*) = \infty, \quad c_2(\mathbf{z}_\eta^*) = \frac{4(1 + \sqrt{2})}{\sigma_r(\mathbf{X}^*)}.$$

Step 4: Since $\nabla \mathcal{P}_C(\mathbf{z}_\eta^*) = \nabla \mathcal{P}_C(\mathbf{x}^*) = \mathbf{P}_{\mathbf{U}_\perp^\top \mathbf{V}_\perp}^\perp$, using (19), we obtain the asymptotic linear rate as

$$\rho = \max\{|1 - \eta\lambda_1|, |1 - \eta\lambda_{r(m+n-r)}|\}, \quad (67)$$

where λ_1 and $\lambda_{r(m+n-r)}$ are the largest and smallest eigenvalues of $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp$, respectively.

Step 5: From (67), we have $\rho < 1$ if and only if $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp$ has full rank and

$$0 < \eta < \frac{2}{\|\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp\|_2}.$$

Here, we would like to point out the condition $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp$ has full rank implies $s \geq r(m+n-r)$, which can be interpreted as a requirement for the number of observations being no less than the degree of freedom in matrix completion. The invertibility of $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp$ is also equivalent to the injectivity of the sampling operator restricted to the tangent space T to $\mathcal{M}_{\leq r}$ at \mathbf{X}^* , denoted by $\mathcal{A}_{\Omega T}$ in [71]-Section 4.2. It is interesting to note that under the standard assumptions on uniform sampling and incoherence property, Candès and Recht [71] showed that $\mathcal{A}_{\Omega T}$ is injective with high probability.

Step 6: Recall that $c_1(\mathbf{x}^*) = c_1(\mathbf{z}_\eta^*) = \infty$. Since $\nabla \mathcal{P}_C(\mathbf{z}_\eta^*) = \nabla \mathcal{P}_C(\mathbf{x}^*) = \mathbf{P}_{\mathbf{U}_\perp^\top \mathbf{V}_\perp}^\perp$, using (20), the region of linear convergence is given by

$$\|\mathbf{x} - \mathbf{x}^*\| \leq \frac{(1 - \rho)\sigma_r(\mathbf{X}^*)}{8(1 + \sqrt{2})}. \quad (68)$$

Remark 8. Similar to (38) and (39), the optimal step size and the optimal convergence rate are given by

$$\eta_{opt} = \frac{2}{\lambda_1(\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp) + \lambda_{r(m+n-r)}(\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp)},$$

$$\rho_{opt} = 1 - \frac{2}{\kappa(\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp) + 1}. \quad (69)$$

We consider the following numerical experiment to verify the analytical rate in (67). The data is generated randomly as follows. First, we sample two matrices \mathbf{A} and \mathbf{B} with *i.i.d* normally distributed entries, of dimensions 50×3 and 40×3 , respectively. Next, we obtain the rank-3 matrix of dimension 50×40 as the product $\mathbf{X}^* = \mathbf{A}\mathbf{B}^\top$. Third, we select 800 observations uniformly at random among the 2000 positions in \mathbf{X}^* . We apply PGD with different step sizes (listed in Fig. 3) including η_{opt} in (69) and record the value of $\|\mathbf{X}^{(k)} - \mathbf{X}^*\|_F$ as a function of k . It can be seen from Fig. 3 that the theoretical rate matches well the empirical rate, reassuring the correctness of our analysis in the previous section.




Fig. 3: (Log-scale) plot of the distance between the current iterate and the local minimizer of the matrix completion problem as a function of the number of iterations. Each solid line corresponds to PGD with a different fixed step size. Each dashed line represents the respective exponential bound ρ^k up to a constant, where the theoretical rate ρ is given by (67). In the experiment, we select $m = 50$, $n = 40$, $r = 3$, and $s = 800$. The optimal step size $\eta_{opt} = 2.2833$ is given by (69), with the corresponding optimal rate $\rho_{opt} = 0.9265$.

Remark 9. The rate in (67) has not been proposed in the literature. However, in the special case of using unit step size, it matches the rate established for the IHTSVD algorithm in [32]. In their work, the authors provide the result relative to the matrix $(\mathbf{S}_\Omega)^\top \mathbf{P}_{\mathbf{U}_\perp^\top \mathbf{V}_\perp} \mathbf{S}_\Omega$ instead of $\mathbf{Q}_\perp^\top \mathbf{S}_\Omega \mathbf{S}_\Omega^\top \mathbf{Q}_\perp$, where $\mathbf{S}_\Omega \in \mathbb{R}^{mn \times (mn-s)}$ is the selection matrix that is complement to \mathbf{S}_Ω . It can be shown that the two matrices share the same set of eigenvalues while may only differ by the eigenvalues at 1. Since IHTSVD uses $\eta = 1$, these unit eigenvalues do not affect the maximization in (67). Compared to the local convergence result in [32], our application in this subsection not only considers PGD with different step sizes but also includes the region of linear convergence in (68).

V. CONCLUSION AND FUTURE WORK

We presented a unified framework to analyze the local convergence of projected gradient descent for constrained least squares. Our analysis provides the asymptotic rate of convergence in a closed-form expression, the number of iterations required to reach certain accuracy, and the local region of convergence. Notably, our technique relies on the Lipschitz-continuous differentiability of the projection operator at two key points: \mathbf{x}^* and \mathbf{z}_η^* . Finally, we demonstrated the application of our proposed framework to local convergence analysis of PGD in four well-known problems: linear equality-constrained least squares, sparse recovery, least squares with a unit norm constraint, and matrix completion.

While the work here focuses on the specific setting of linear converges of the PGD algorithm, we believe it can be expanded in several directions. First, our framework can be utilized to analyze the following cases: (i) adaptive step

size schemes (e.g., the backtracking line search rule), (ii) accelerated methods (e.g., the Nesterov's accelerated gradient and the Heavy Ball method), (iii) general objective functions other than least squares, and (iv) other algorithms for manifold optimization such as Riemannian gradient descent. Another interesting research direction is to sharpen the theoretical bound on the ROC in order to better explain the actual region in which the algorithm converges to the desired solution. Finally, the proposed framework can be used to further study the performance of PGD for a variety of constrained least squares problems arising in the area of phase-only beamforming [10], online power system optimization [77], spectral compressed sensing [78], and linear dimensionality reduction [79].

APPENDIX A PROOF OF LEMMA 1

Our goal in the proof of Lemma 1 is to show that if the fixed point condition $\mathbf{x}^* = \mathcal{P}_C(\mathbf{x}^* - \eta \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}))$ holds, then the stationarity condition $\nabla \mathcal{P}_C(\mathbf{x}^*) \mathbf{A}^\top \mathbf{A}\mathbf{x}^* - \mathbf{b} = \mathbf{0}$ holds. Note that if $\mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}) = \mathbf{0}$, then the stationarity condition holds trivially. Hence, we focus on the proof for $\mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}) \neq \mathbf{0}$. We first show that for any $0 \leq \alpha < 1$, $\mathbf{x}^* = \mathcal{P}_C(\mathbf{x}^* - \eta \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}))$ is a sufficient condition for

$$\Pi_C \mathbf{x}^* - \alpha \eta \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b}) = \{\mathbf{x}^*\}. \quad (70)$$

Then, using (70) and the differentiability of \mathcal{P}_C at \mathbf{x}^* , we prove that $\nabla \mathcal{P}_C(\mathbf{x}^*) \mathbf{A}^\top \mathbf{A}\mathbf{x}^* - \mathbf{b} = \mathbf{0}$. We proceed with the detailed proof.

First, let $\mathbf{v}^* = \mathbf{A}^\top(\mathbf{A}\mathbf{x}^* - \mathbf{b})$ and $\mathbf{z}_{\alpha\eta}^* = \mathbf{x}^* - \alpha\eta\mathbf{v}^*$. On the one hand, for any $0 \leq \alpha < 1$ and $\mathbf{y} \in \Pi_C(\mathbf{z}_{\alpha\eta}^*)$, we have

$$\begin{aligned} \|\mathbf{y} - \mathbf{z}_\eta^*\| &= \|(\mathbf{y} - \mathbf{z}_{\alpha\eta}^*) + (\mathbf{z}_{\alpha\eta}^* - \mathbf{z}_\eta^*)\| \\ &\leq \|\mathbf{y} - \mathbf{z}_{\alpha\eta}^*\| + \|\mathbf{z}_{\alpha\eta}^* - \mathbf{z}_\eta^*\| \\ &= d(\mathbf{z}_{\alpha\eta}^*, \mathcal{C}) + \|\mathbf{z}_{\alpha\eta}^* - \mathbf{z}_\eta^*\| \\ &\leq \|\mathbf{x}^* - \mathbf{z}_{\alpha\eta}^*\| + \|\mathbf{z}_{\alpha\eta}^* - \mathbf{z}_\eta^*\|, \end{aligned} \quad (71)$$

where the first inequality uses the triangle inequality that holds when $\mathbf{y} - \mathbf{z}_{\alpha\eta}^* = \beta(\mathbf{z}_{\alpha\eta}^* - \mathbf{z}_\eta^*)$, for some $\beta \geq 0$. Using the fact that $\mathbf{z}_\eta^* = \mathbf{x}^* - \eta\mathbf{v}^*$ and $\mathbf{z}_{\alpha\eta}^* = \mathbf{x}^* - \alpha\eta\mathbf{v}^*$, we obtain

$$\begin{aligned} \|\mathbf{x}^* - \mathbf{z}_{\alpha\eta}^*\| + \|\mathbf{z}_{\alpha\eta}^* - \mathbf{z}_\eta^*\| &= \|\alpha\eta\mathbf{v}^*\| + \|(1 - \alpha)\eta\mathbf{v}^*\| \\ &= \|\alpha\eta\mathbf{v}^* + (1 - \alpha)\eta\mathbf{v}^*\| \\ &= \|\eta\mathbf{v}^*\| \\ &= \|\mathbf{x}^* - \mathbf{z}_\eta^*\|. \end{aligned} \quad (72)$$

From (71) and (72), we have

$$\|\mathbf{y} - \mathbf{z}_\eta^*\| \leq \|\mathbf{x}^* - \mathbf{z}_\eta^*\|, \quad (73)$$

with the equality holding if and only if

$$\begin{aligned} \mathbf{y} - \mathbf{z}_{\alpha\eta}^* &= \beta(\mathbf{z}_{\alpha\eta}^* - \mathbf{z}_\eta^*), \\ \|\mathbf{y} - \mathbf{z}_{\alpha\eta}^*\| &= \|\mathbf{x}^* - \mathbf{z}_{\alpha\eta}^*\|. \end{aligned} \quad (74)$$

Using the fact that $\mathbf{z}_\eta^* = \mathbf{x}^* - \eta\mathbf{v}^*$ and $\mathbf{z}_{\alpha\eta}^* = \mathbf{x}^* - \alpha\eta\mathbf{v}^*$, (74) holds if and only if

$$\beta = \frac{\alpha}{1 - \alpha} \text{ and } \mathbf{y} = \mathbf{x}^*. \quad (75)$$

On the other hand, since $\mathbf{x}^* = \mathcal{P}_C(\mathbf{z}_\eta^*)$ and $\mathbf{y} \in \mathcal{C}$, we have

$$\|\mathbf{y} - \mathbf{z}_\eta^*\| \geq \|\mathbf{x}^* - \mathbf{z}_\eta^*\|. \quad (76)$$

From (73) and (76), we conclude that $\|\mathbf{y} - \mathbf{z}_\eta^*\| = \|\mathbf{x}^* - \mathbf{z}_\eta^*\|$. Moreover, from (75), the equality holds if and only if $\mathbf{y} = \mathbf{x}^*$. Since this holds for any $\mathbf{y} \in \Pi_C(\mathbf{z}_{\alpha\eta}^*)$, we conclude that $\Pi_C(\mathbf{z}_{\alpha\eta}^*) = \{\mathbf{x}^*\}$ for all $0 \leq \alpha < 1$.

Next, using the differentiability of the projection \mathcal{P}_C at \mathbf{x}^* from Definition 1, we have

$$\lim_{\alpha \rightarrow 0} \sup_{\mathbf{y} \in \Pi_C(\mathbf{x}^* - \alpha\eta\mathbf{v}^*)} \frac{\|\mathbf{y} - \mathcal{P}_C(\mathbf{x}^*) - \nabla \mathcal{P}_C(\mathbf{x}^*)(\alpha\eta\mathbf{v}^*)\|}{\|\alpha\eta\mathbf{v}^*\|} = 0.$$

Substituting $\Pi_C(\mathbf{x}^* - \alpha\eta\mathbf{v}^*) = \Pi_C(\mathbf{z}_{\alpha\eta}^*) = \{\mathbf{x}^*\}$ and $\mathcal{P}_C(\mathbf{x}^*) = \mathbf{x}^*$ into the last equation, we obtain

$$\begin{aligned} 0 &= \lim_{\alpha \rightarrow 0} \frac{\|\mathbf{x}^* - \mathbf{x}^* - \alpha\eta\nabla \mathcal{P}_C(\mathbf{x}^*)\mathbf{v}^*\|}{\|\alpha\eta\mathbf{v}^*\|} \\ &= \lim_{\alpha \rightarrow 0} \frac{\alpha\eta\|\nabla \mathcal{P}_C(\mathbf{x}^*)\mathbf{v}^*\|}{\alpha\eta\|\mathbf{v}^*\|} \\ &= \frac{\|\nabla \mathcal{P}_C(\mathbf{x}^*)\mathbf{v}^*\|}{\|\mathbf{v}^*\|}, \end{aligned}$$

which only holds if $\nabla \mathcal{P}_C(\mathbf{x}^*)\mathbf{v}^* = \mathbf{0}$. This completes our proof of the lemma.

APPENDIX B PROOF OF COROLLARY 1

In the following, under the assumption $\nabla \mathcal{P}_C(\mathbf{z}_\eta^*) = \nabla \mathcal{P}_C(\mathbf{x}^*) = \mathbf{U}_x \mathbf{U}_x^\top$, we show that (i) the asymptotic convergence rate $\rho(\mathbf{H})$ is given by (19), (ii) the sufficient conditions for $\rho(\mathbf{H}) < 1$ are $(\mathbf{A}\mathbf{U}_x)^\top \mathbf{A}\mathbf{U}_x$ is full rank and (18) holds, and (iii) the region of linear convergence can be simplified from (13) to (20).

First, we prove (19) by simplifying the expression of \mathbf{H} in (17) and the fact that $\mathbf{U}_x^\top \mathbf{U}_x = \mathbf{I}_d$ as follows. Substituting $\nabla \mathcal{P}_C(\mathbf{x}^*)$ and $\nabla \mathcal{P}_C(\mathbf{z}_\eta^*)$ by $\mathbf{U}_x \mathbf{U}_x^\top$ into (12) yields

$$\begin{aligned} \mathbf{H} &= \mathbf{U}_x \mathbf{U}_x^\top (\mathbf{I}_n - \eta \mathbf{A}^\top \mathbf{A}) \mathbf{U}_x \mathbf{U}_x^\top \\ &= \mathbf{U}_x (\mathbf{I}_d - \eta \mathbf{U}_x^\top \mathbf{A}^\top \mathbf{A} \mathbf{U}_x) \mathbf{U}_x^\top, \end{aligned}$$

where the second equality stems from $\mathbf{U}_x^\top \mathbf{U}_x = \mathbf{I}_d$. Since \mathbf{H} is symmetric, its spectral radius equals to its spectral norm:

$$\rho(\mathbf{H}) = \|\mathbf{U}_x (\mathbf{I}_d - \eta \mathbf{U}_x^\top \mathbf{A}^\top \mathbf{A} \mathbf{U}_x) \mathbf{U}_x^\top\|_2.$$

Using the fact that the spectral norm is invariant under left-multiplication by matrices with orthonormal columns and right-multiplication by matrices with orthonormal rows (see [63] - Exercise 5.6.9), we further have

$$\rho(\mathbf{H}) = \|\mathbf{I}_d - \eta \mathbf{U}_x^\top \mathbf{A}^\top \mathbf{A} \mathbf{U}_x\|_2. \quad (77)$$

Let $\mathbf{U}_x^\top \mathbf{A}^\top \mathbf{A} \mathbf{U}_x = \hat{\mathbf{U}} \hat{\mathbf{\Lambda}} \hat{\mathbf{U}}^\top$ be an eigendecomposition, where $\hat{\mathbf{U}} \in \mathbb{R}^{d \times d}$ is an orthogonal matrix and $\hat{\mathbf{\Lambda}} = \text{diag}(\lambda_1(\mathbf{U}_x^\top \mathbf{A}^\top \mathbf{A} \mathbf{U}_x), \dots, \lambda_d(\mathbf{U}_x^\top \mathbf{A}^\top \mathbf{A} \mathbf{U}_x))$. Since $\hat{\mathbf{U}}^\top \hat{\mathbf{U}} = \mathbf{I}_d$, (77) can be represented as

$$\rho(\mathbf{H}) = \hat{\mathbf{U}} (\mathbf{I}_d - \hat{\mathbf{\Lambda}}) \hat{\mathbf{U}}^\top = \mathbf{I}_d - \hat{\mathbf{\Lambda}}.$$

Now using the fact that the spectral norm of a diagonal matrix is the maximum of the absolute values of its diagonal entries, we obtain

$$\begin{aligned}\rho(\mathbf{H}) &= \max_{1 \leq i \leq d} |1 - \eta \lambda_i(\mathbf{U}_x^\top \mathbf{A}^\top \mathbf{A} \mathbf{U}_x)| \\ &= \max\{|1 - \eta \lambda_1|, |1 - \eta \lambda_d|\}.\end{aligned}$$

Second, we establish the sufficient conditions for $\rho(\mathbf{H}) < 1$ by bounding each term inside the maximum in (19) as follows. Since $\lambda_d \geq 0$, we have

$$-1 < 1 - \eta \lambda_1 \leq 1 - \eta \lambda_d \leq 1, \text{ for } i = 1, \dots, d,$$

if $0 < \eta < 2/\lambda_1$. It is also noted from the definition of the spectral norm that $\|\mathbf{A} \mathbf{U}_x\|_2^2 = \lambda_1$. Therefore, $\rho(\mathbf{H}) \leq 1$ provided that (18) holds. The equality $\rho(\mathbf{H}) = 1$ holds if and only if $\lambda_d = 0$, i.e., $(\mathbf{A} \mathbf{U}_x)^\top \mathbf{A} \mathbf{U}_x$ is singular. In other words, when $(\mathbf{A} \mathbf{U}_x)^\top \mathbf{A} \mathbf{U}_x$ is full rank and (18) holds, the linear convergence is guaranteed as $\rho(\mathbf{H}) < 1$.

Finally, the region of linear convergence in (20) is determined based on simplifying (13) as follows. First, using Remark 1, we obtain $\kappa(\mathbf{Q}) = 1$. Second, from (17), we have $\|\nabla \mathcal{P}_C(\mathbf{z}_\eta^*)\|_2 = \|\mathcal{P}_{T_x}(c)\|_2 = 1$. Third, substituting $\kappa(\mathbf{Q}) = 1$ and $\|\nabla \mathcal{P}_C(\mathbf{z}_\eta^*)\|_2 = 1$ into (13) yields (20). This completes our proof of the corollary.

APPENDIX C PROOF OF LEMMA 2

Our goal is to show the error vector $\delta^{(k)}$ satisfies the asymptotically-linear quadratic system dynamic in (25) and to bound the norm of the residual \mathbf{q}_2 by (26).

First, our key idea in proving (25) is the Lipschitz-continuous differentiability of \mathcal{P}_C at \mathbf{x}^* and at \mathbf{z}_η^* . Specifically, for any k such that $\delta^{(k)}$ admits a perturbation $(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \delta^{(k)}$ that satisfies

$$\|(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \delta^{(k)}\| < c_1(\mathbf{z}_\eta^*), \quad (78)$$

applying the Lipschitz-continuous differentiability of \mathcal{P}_C at \mathbf{z}_η^* to (23) yields

$$\delta^{(k+1)} = \nabla \mathcal{P}_C(\mathbf{z}_\eta^*)(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \delta^{(k)} + \mathbf{q}_1(\delta^{(k)}), \quad (79)$$

where the residual $\mathbf{q}_1: \mathbb{R}^n \rightarrow \mathbb{R}^n$ satisfies

$$\begin{aligned}\|\mathbf{q}_1(\delta^{(k)})\| &\leq c_2(\mathbf{z}_\eta^*) \|(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \delta^{(k)}\|^2 \\ &\leq c_2(\mathbf{z}_\eta^*) u_\eta^2 \|\delta^{(k)}\|^2.\end{aligned} \quad (80)$$

On the other hand, using the fact that $\mathbf{x}^* = \mathcal{P}_C(\mathbf{x}^*)$, $\mathbf{x}^{(k)} = \mathcal{P}_C(\mathbf{x}^{(k)})$, and the Lipschitz-continuous differentiability of \mathcal{P}_C at \mathbf{x}^* with the perturbation $\delta^{(k)} \in \mathcal{B}(\mathbf{0}, c_1(\mathbf{x}^*))$, we obtain

$$\begin{aligned}\delta^{(k)} &= \mathbf{x}^{(k)} - \mathbf{x}^* \\ &= \mathcal{P}_C(\mathbf{x}^* + \delta^{(k)}) - \mathcal{P}_C(\mathbf{x}^*) \\ &= \nabla \mathcal{P}_C(\mathbf{x}^*) \delta^{(k)} + \mathbf{q}_x(\delta^{(k)}),\end{aligned} \quad (81)$$

where the residual $\mathbf{q}_x: \mathbb{R}^n \rightarrow \mathbb{R}^n$ satisfies $\|\mathbf{q}_x(\delta^{(k)})\| \leq c_2(\mathbf{x}^*) \|\delta^{(k)}\|^2$. We proceed with the proof of (25) by combining the results from (79) and (81) as follows. Since $\|(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \delta^{(k)}\| \leq \|\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}\|_2 \|\delta^{(k)}\| = u_\eta \|\delta^{(k)}\|$, the sufficient condition for (78) is $\|\delta^{(k)}\| < c_1(\mathbf{z}_\eta^*)/u_\eta$. Thus,

$\|\delta^{(k)}\| < \min\{c_1(\mathbf{x}^*), c_1(\mathbf{z}_\eta^*)/u_\eta\}$ is sufficient for both (79) and (81). Substituting (81) into the RHS of (79), we obtain (25) with $\mathbf{q}_2(\delta^{(k)}) = \nabla \mathcal{P}_C(\mathbf{z}_\eta^*)(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \mathbf{q}_x(\delta^{(k)}) + \mathbf{q}_1(\delta^{(k)})$. Next, to bound the norm of the residual \mathbf{q}_2 , we apply the triangle inequality as follows

$$\begin{aligned}\|\mathbf{q}_2(\delta^{(k)})\| &\leq \|\nabla \mathcal{P}_C(\mathbf{z}_\eta^*)(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \mathbf{q}_x(\delta^{(k)})\| \\ &\quad + \|\mathbf{q}_1(\delta^{(k)})\|.\end{aligned} \quad (82)$$

On the one hand, the first term on the RHS of (82) can be bounded by

$$\begin{aligned}\|\nabla \mathcal{P}_C(\mathbf{z}_\eta^*)(\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}) \mathbf{q}_x(\delta^{(k)})\| &\leq \|\nabla \mathcal{P}_C(\mathbf{z}_\eta^*)\|_2 \|\mathbf{I} - \eta \mathbf{A}^\top \mathbf{A}\|_2 \|\mathbf{q}_x(\delta^{(k)})\| \\ &\leq \|\nabla \mathcal{P}_C(\mathbf{z}_\eta^*)\|_2 u_\eta c_2(\mathbf{x}^*) \|\delta^{(k)}\|^2.\end{aligned} \quad (83)$$

On the other hand, the second term on the RHS of (82) can be bounded by (80). Combining the two bounds, we obtain (26).

APPENDIX D PROOF OF LEMMA 3

In this section, our goal is to show the recursion on the transformed error vector (29) holds at any $k \in \mathbb{N}$ provided that the initial error vector lies within the region of linear convergence described by (13). In the first step, we prove that if the current transformed error vector lies within the region of linear convergence

$$\|\tilde{\delta}^{(k)}\| < \min \left\{ \frac{c_1(\mathbf{x}^*)}{\|\mathbf{Q}\|_2}, \frac{c_1(\mathbf{z}_\eta^*)}{\|\mathbf{Q}\|_2 u_\eta}, \frac{1 - \rho(\mathbf{H})}{q/\|\mathbf{Q}^{-1}\|_2} \right\}. \quad (84)$$

then $\tilde{\delta}^{(k+1)} = \tilde{\delta}^{(k)} + \mathbf{q}_3(\tilde{\delta}^{(k)})$ and moreover, the next transformed error vector also lies within the region of linear convergence

$$\|\tilde{\delta}^{(k+1)}\| < \min \left\{ \frac{c_1(\mathbf{x}^*)}{\|\mathbf{Q}\|_2}, \frac{c_1(\mathbf{z}_\eta^*)}{\|\mathbf{Q}\|_2 u_\eta}, \frac{1 - \rho(\mathbf{H})}{q/\|\mathbf{Q}^{-1}\|_2} \right\}. \quad (85)$$

Therefore, by the principle of induction, the initial condition on the transformed error vector, i.e., (84) holds at $k = 0$, is the sufficient condition for (29) to hold at any $k \in \mathbb{N}$. In the second step, we show that (84) holds at $k = 0$ if the initial condition on the error vector (13) holds and hence, completes the proof of lemma. We proceed with our detailed proof below.

First, let us assume that (84) holds. We have

$$\begin{aligned}\|\delta^{(k)}\| &= \|\mathbf{Q} \tilde{\delta}^{(k)}\| \leq \|\mathbf{Q}\|_2 \|\tilde{\delta}^{(k)}\| \\ &< \|\mathbf{Q}\|_2 \min \left\{ \frac{c_1(\mathbf{x}^*)}{\|\mathbf{Q}\|_2}, \frac{c_1(\mathbf{z}_\eta^*)}{\|\mathbf{Q}\|_2 u_\eta}, \frac{1 - \rho(\mathbf{H})}{q/\|\mathbf{Q}^{-1}\|_2} \right\} \\ &\leq \min \left\{ c_1(\mathbf{x}^*), \frac{c_1(\mathbf{z}_\eta^*)}{u_\eta} \right\},\end{aligned} \quad (86)$$

Thus, by Lemma 2, we have $\delta^{(k+1)} = \mathbf{H} \delta^{(k)} + \mathbf{q}_2(\delta^{(k)})$. Substituting $\mathbf{H} = \mathbf{Q} \mathbf{Q}^{-1}$ and multiplying both sides with \mathbf{Q}^{-1} yields $\mathbf{Q}^{-1} \delta^{(k+1)} = \mathbf{Q}^{-1} \delta^{(k)} + \mathbf{Q}^{-1} \mathbf{q}_2(\delta^{(k)})$. Replacing $\mathbf{Q}^{-1} \delta^{(k)}$ by $\tilde{\delta}^{(k)}$ and $\delta^{(k)}$ by $\mathbf{Q} \tilde{\delta}^{(k)}$ in the last equation, we

obtain (29), i.e., $\tilde{\delta}^{(k+1)} = \tilde{\delta}^{(k)} + \mathbf{q}_3(\tilde{\delta}^{(k)})$. Here, the second term \mathbf{q}_3 can be bounded as follows

$$\begin{aligned} \|\mathbf{q}_3 \tilde{\delta}^{(k)}\| &= \|\mathbf{Q}^{-1} \mathbf{q}_2(\mathbf{Q} \tilde{\delta}^{(k)})\| \leq \|\mathbf{Q}^{-1}\|_2 \|\mathbf{q}_2(\mathbf{Q} \tilde{\delta}^{(k)})\| \\ &\leq \|\mathbf{Q}^{-1}\|_2 c_2(\mathbf{x}^*) + \|\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_\eta^*)\|_2 c_2(\mathbf{z}_\eta^*) \|\mathbf{Q} \tilde{\delta}^{(k)}\|^2 \\ &\leq \|\mathbf{Q}^{-1}\|_2 c_2(\mathbf{x}^*) + \|\nabla \mathcal{P}_{\mathcal{C}}(\mathbf{z}_\eta^*)\|_2 c_2(\mathbf{z}_\eta^*) \|\mathbf{Q}\|_2^2 \|\tilde{\delta}^{(k)}\|^2 \\ &= \frac{q}{\|\mathbf{Q}^{-1}\|_2} \|\tilde{\delta}^{(k)}\|^2. \end{aligned} \quad (87)$$

Now, taking the norms of both sides of (29) and applying the triangle inequality yield

$$\begin{aligned} \|\tilde{\delta}^{(k+1)}\| &\leq \|\tilde{\delta}^{(k)}\| + \|\mathbf{q}_3 \tilde{\delta}^{(k)}\| \\ &\leq \rho(\mathbf{H}) \|\tilde{\delta}^{(k)}\| + \frac{q}{\|\mathbf{Q}^{-1}\|_2} \|\tilde{\delta}^{(k)}\|^2 \\ &< \rho(\mathbf{H}) \|\tilde{\delta}^{(k)}\| + 1 - \rho(\mathbf{H}) \|\tilde{\delta}^{(k)}\| \\ &= \|\tilde{\delta}^{(k)}\|, \end{aligned} \quad (88)$$

where the second inequality stems from $\|\tilde{\delta}^{(k)}\| < (1 - \rho(\mathbf{H})) / (q / \|\mathbf{Q}^{-1}\|_2)$. From (84) and (88), we conclude that (85) holds. By the principle of induction, we have (84) holds for all $k \in \mathbb{N}$ provided that it holds at $k = 0$, i.e.,

$$\|\tilde{\delta}^{(0)}\| < \min \left\{ \frac{c_1(\mathbf{x}^*)}{\|\mathbf{Q}\|_2}, \frac{c_1(\mathbf{z}_\eta^*)}{\|\mathbf{Q}\|_2 u_\eta}, \frac{1 - \rho(\mathbf{H})}{q / \|\mathbf{Q}^{-1}\|_2} \right\}. \quad (89)$$

Second, we prove that (13) is sufficient for (89). Using the definition $\tilde{\delta}^{(k)} = \mathbf{Q}^{-1} \delta^{(k)}$, we have

$$\|\tilde{\delta}^{(k)}\| = \|\mathbf{Q}^{-1} \delta^{(k)}\| \leq \|\mathbf{Q}^{-1}\|_2 \|\delta^{(k)}\|. \quad (90)$$

Upper-bounding $\|\delta^{(k)}\|$ by the LHS of (13) and substituting back into (90) yield

$$\|\tilde{\delta}^{(k)}\| < \|\mathbf{Q}^{-1}\|_2 \min \left\{ \frac{c_1(\mathbf{x}^*)}{\kappa(\mathbf{Q})}, \frac{c_1(\mathbf{z}_\eta^*)}{\kappa(\mathbf{Q}) u_\eta}, \frac{1 - \rho(\mathbf{H})}{q} \right\}.$$

Finally, replacing $\kappa(\mathbf{Q})$ by the product $\|\mathbf{Q}\|_2 \|\mathbf{Q}^{-1}\|_2$ and simplifying yield (89). This completes our proof of the lemma.

APPENDIX E

PROOF OF LEMMA 4

In this section, we show the convergence of $\{\|\tilde{\delta}^{(k)}\|\}_{k=0}^\infty$ using Theorem 1 in [37]. Our idea is to consider a surrogate sequence $\{a_k\}_{k=0}^\infty$ that upper-bounds $\{\|\tilde{\delta}^{(k)}\|\}_{k=0}^\infty$:

$$\begin{aligned} a_0 &= \|\tilde{\delta}^{(0)}\|, \\ a_{k+1} &= \rho(\mathbf{H}) a_k + \frac{q}{\|\mathbf{Q}^{-1}\|_2} a_k^2. \end{aligned}$$

First, we prove by induction that

$$\|\tilde{\delta}^{(k)}\| \leq a_k \quad \forall k \in \mathbb{N}. \quad (91)$$

The base case when $k = 0$ holds trivially as $a_0 = \|\tilde{\delta}^{(0)}\|$. In the induction step, given $\|\tilde{\delta}^{(k)}\| \leq a_k$ for some integer $k \geq 0$, we have

$$\begin{aligned} \|\tilde{\delta}^{(k+1)}\| &\leq \rho(\mathbf{H}) \|\tilde{\delta}^{(k)}\| + \frac{q}{\|\mathbf{Q}^{-1}\|_2} \|\tilde{\delta}^{(k)}\|^2 \\ &\leq \rho a_k + \frac{q}{\|\mathbf{Q}^{-1}\|_2} a_k^2 \\ &= a_{k+1}. \end{aligned}$$

By the principle of induction, (91) holds for all $k \in \mathbb{N}$. Next, applying Theorem 1 in [37], under the condition $a_0 = \|\tilde{\delta}^{(0)}\| < (1 - \rho(\mathbf{H})) / (q / \|\mathbf{Q}^{-1}\|_2)$, yields $a_k \leq \tilde{\epsilon} a_0$ for any integer k satisfies (31). From (91), we further have $\|\tilde{\delta}^{(k)}\| \leq a_k \leq \tilde{\epsilon} a_0 = \tilde{\epsilon} \|\tilde{\delta}^{(0)}\|$. This completes our proof of the lemma.

ACKNOWLEDGMENT

The authors would like to thank Prof. Amir Beck of Tel-Aviv University for discussion on projected gradient descent.

REFERENCES

- [1] M. A. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Top. Signal. Process.*, vol. 1, no. 4, pp. 586–597, 2007.
- [2] T. Blumensath and M. E. Davies, "Iterative thresholding for sparse approximations," *J. Fourier Anal. Appl.*, vol. 14, no. 5-6, pp. 629–654, 2008.
- [3] —, "Iterative hard thresholding for compressed sensing," *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.
- [4] B. R. Hunt, "The application of constrained least squares estimation to image restoration by digital computer," *IEEE Trans. Comput.*, vol. 100, no. 9, pp. 805–812, 1973.
- [5] N. P. Galatsanos, A. K. Katsaggelos, R. T. Chin, and A. D. Hillery, "Least squares restoration of multichannel images," *IEEE Trans. Signal Process.*, vol. 39, no. 10, pp. 2222–2236, 1991.
- [6] V. Z. Mesarovic, N. P. Galatsanos, and A. K. Katsaggelos, "Regularized constrained total least squares image restoration," *IEEE Trans. Image Process.*, vol. 4, no. 8, pp. 1096–1108, 1995.
- [7] C. I. Puryear, O. N. Portniaguine, C. M. Cobos, and J. P. Castagna, "Constrained least-squares spectral analysis: Application to seismic data," *Geophysics*, vol. 77, no. 5, pp. V143–V167, 2012.
- [8] Y. Chen, J. Yuan, S. Zu, S. Qu, and S. Gan, "Seismic imaging of simultaneous-source data using constrained least-squares reverse time migration," *J. Appl. Geophys.*, vol. 114, pp. 32–35, 2015.
- [9] W. Menke, *Geophysical data analysis: Discrete inverse theory*. Academic press, 2018.
- [10] J. Tranter, N. D. Sidiropoulos, X. Fu, and A. Swami, "Fast unit-modulus least squares with applications in beamforming," *IEEE Trans. Signal Process.*, vol. 65, no. 11, pp. 2875–2887, 2017.
- [11] M. Zhang, J. Li, S. Zhu, and X. Chen, "Fast and simple gradient projection algorithms for phase-only beamforming," *IEEE Trans. Veh. Technol.*, 2021.
- [12] P. Jain, R. Meka, and I. S. Dhillon, "Guaranteed rank minimization via singular value projection," in *Proc. Adv. Neural Inf. Process. Syst.*, 2010, pp. 937–945.
- [13] Y. Chen and M. J. Wainwright, "Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees," *arXiv preprint arXiv:1509.03025*, 2015.
- [14] R. Khanna and A. Kyrillidis, "IHT dies hard: Provable accelerated iterative hard thresholding," in *Proc. Int. Conf. Artif. Intell. Stat.*, 2018, pp. 188–198.
- [15] C.-J. Lin, "Projected gradient methods for nonnegative matrix factorization," *Neural Comput.*, vol. 19, no. 10, pp. 2756–2779, 2007.
- [16] N. Mohammadiha and A. Leijon, "Nonnegative matrix factorization using projected gradient algorithms with sparseness constraints," in *IEEE Int. Symp. Signal Process. Inf. Technol.* IEEE, 2009, pp. 418–423.
- [17] N. Guan, D. Tao, Z. Luo, and B. Yuan, "NeNMF: An optimal gradient method for nonnegative matrix factorization," *IEEE Trans. Signal Process.*, vol. 60, no. 6, pp. 2882–2898, 2012.
- [18] D. G. Luenberger, Y. Ye *et al.*, *Linear and nonlinear programming*. Springer, 1984, vol. 2.
- [19] D. P. Bertsekas, "Nonlinear programming," *J. Oper. Res. Soc.*, vol. 48, no. 3, pp. 334–334, 1997.
- [20] A. Beck, *First-order methods in optimization*. SIAM, 2017.
- [21] P. Jain, P. Kar *et al.*, "Non-convex optimization for machine learning," *Found. Trends Mach. Learn.*, vol. 10, no. 3-4, pp. 142–336, 2017.
- [22] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-point algorithms for inverse problems in science and engineering*. Springer, 2011, pp. 185–212.
- [23] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

- [24] R. Sun and Z.-Q. Luo, "Guaranteed matrix completion via non-convex factorization," *IEEE Trans. Inf. Theory*, vol. 62, no. 11, pp. 6535–6579, 2016.
- [25] A. Beck and Y. Vaisbourd, "Globally solving the trust region subproblem using simple first-order methods," *SIAM J. Optim.*, vol. 28, no. 3, pp. 1951–1967, 2018.
- [26] D. G. Luenberger, "The gradient projection method along geodesics," *Manag. Sci.*, vol. 18, no. 11, pp. 620–631, 1972.
- [27] A. Lichnewsky, "Minimisation des fonctionnelles définies sur une variété par la méthode du gradient conjugué," Ph.D. dissertation, These de Doctorat d'Etat. Paris: Université de Paris-Sud, 1979.
- [28] D. Gabay, "Minimizing a differentiable function over a differential manifold," *J. Optim. Theory Appl.*, vol. 37, no. 2, pp. 177–219, 1982.
- [29] A. Uschmajew and B. Vandereycken, "Geometric methods on low-rank matrix and tensor manifolds," in *Variational methods for nonlinear geometric data and applications*. Springer Nature, 2020, pp. 261–313.
- [30] B. O'donoghue and E. Candes, "Adaptive restart for accelerated gradient schemes," *Found. Comput. Math.*, vol. 15, no. 3, pp. 715–732, 2015.
- [31] W. G. Strang, "On the Kantorovich inequality," in *Proc. Amer. Math. Soc.*, vol. 11, no. 468, 1960, pp. 0095–09601.
- [32] E. Chunikhina, R. Raich, and T. Nguyen, "Performance analysis for matrix completion via iterative hard-thresholded SVD," in *Proc. IEEE Stat. Signal Process. Workshop*. IEEE, 2014, pp. 392–395.
- [33] T. Vu, R. Raich, and X. Fu, "On convergence of projected gradient descent for minimizing a large-scale quadratic over the unit sphere," in *Proc. IEEE Int. Workshop Mach. Learn. Signal Process.* IEEE, 2019, pp. 1–6.
- [34] T. Vu and R. Raich, "Local convergence of the Heavy Ball method in iterative hard thresholding for low-rank matrix completion," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.* IEEE, 2019, pp. 3417–3421.
- [35] —, "Accelerating iterative hard thresholding for low-rank matrix completion via adaptive restart," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.* IEEE, 2019, pp. 2917–2921.
- [36] —, "Exact linear convergence rate analysis for low-rank symmetric matrix completion via gradient descent," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.* IEEE, 2021, pp. 3240–3244.
- [37] —, "A closed-form bound on the asymptotic linear convergence of iterative methods via fixed point analysis," *Optim. Lett.*, 2022.
- [38] T. Vu, E. Chunikhina, and R. Raich, "Perturbation expansions and error bounds for the truncated singular value decomposition," *Linear Algebra Appl.*, 2021.
- [39] B. T. Polyak, "Some methods of speeding up the convergence of iteration methods," *USSR Comput. Math. Math. Phys.*, vol. 4, no. 5, pp. 1–17, 1964.
- [40] I. Gelfand, "Normierte ringe," *Rec. Math. [Mat. Sbornik]*, vol. 9, no. 1, pp. 3–24, 1941.
- [41] F. Vasilyev and A. Y. Ivanitskiy, *In-depth analysis of linear programming*. Springer Science & Business Media, 2013.
- [42] R. L. Foote, "Regularity of the distance function," *Proc. Am. Math. Soc.*, vol. 92, no. 1, pp. 153–155, 1984.
- [43] E. Dudek and K. Holly, "Nonlinear orthogonal projection," *Ann. Pol. Math.*, vol. 59, no. 1, pp. 1–31, 1994.
- [44] A. S. Lewis and J. Malick, "Alternating projections on manifolds," *Math. Oper. Res.*, vol. 33, no. 1, pp. 216–234, 2008.
- [45] L. Ambrosio and C. Mantegazza, "Curvature and distance function from a manifold," *J. Geom. Anal.*, vol. 8, no. 5, pp. 723–748, 1998.
- [46] P.-A. Absil and J. Malick, "Projection-like retractions on matrix manifolds," *SIAM J. Optim.*, vol. 22, no. 1, pp. 135–158, 2012.
- [47] J. Rataj and M. Zähle, *Curvature measures of singular sets*. Springer, 2019.
- [48] G. Leobacher and A. Steinicke, "Existence, uniqueness and regularity of the projection onto differentiable manifolds," *Ann. Glob. Anal. Geom.*, pp. 1–29, 2021.
- [49] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [50] B. T. Polyak, *Introduction to optimization. Optimization software*. Inc., Publications Division, New York, 1987, vol. 1.
- [51] S.-G. Hwang, "Cauchy's interlace theorem for eigenvalues of Hermitian matrices," *Am. Math. Mon.*, vol. 111, no. 2, pp. 157–159, 2004.
- [52] R. Bellman, *Stability theory of differential equations*. McGraw-Hill, NY, New York, 1953.
- [53] M. Abramowitz and I. A. Stegun, "Handbook of mathematical functions with formulas, graphs, and mathematical tables," *NBS Appl. Math. Ser.*, vol. 55, 1964.
- [54] X. Hong, X. Lai, and R. Zhao, "Matrix-based algorithms for constrained least-squares and minimax designs of 2-d linear-phase FIR filters," *IEEE Trans. Signal Process.*, vol. 61, no. 14, pp. 3620–3631, 2013.
- [55] M. L. de Campos, S. Werner, and J. A. Apolinário, "Constrained adaptive filters," in *Adaptive Antenna Arrays: Trends and Applications*. Springer Berlin Heidelberg, 2004, pp. 46–64.
- [56] O. L. Frost, "An algorithm for linearly constrained adaptive array processing," *Proc. IEEE*, vol. 60, no. 8, pp. 926–935, 1972.
- [57] L. S. Resende, J. M. T. Romano, and M. G. Bellanger, "A fast least-squares algorithm for linearly constrained adaptive filtering," *IEEE Trans. Signal Process.*, vol. 44, no. 5, pp. 1168–1174, 1996.
- [58] C. Van Loan, "On the method of weighting for equality-constrained least-squares problems," *SIAM J. Numer. Anal.*, vol. 22, no. 5, pp. 851–864, 1985.
- [59] J. C. Dunn, "Global and asymptotic convergence rate estimates for a class of projected gradient processes," *SIAM J. Control Optim.*, vol. 19, no. 3, pp. 368–400, 1981.
- [60] D. P. Bertsekas and E. M. Gafni, "Projection methods for variational inequalities with application to the traffic assignment problem," in *Nondifferential and variational techniques in optimization*. Springer, 1982, pp. 139–159.
- [61] Z.-Q. Luo and P. Tseng, "Error bound and reduced-gradient projection algorithms for convex minimization over a polyhedral set," *SIAM J. Optim.*, vol. 3, no. 1, pp. 43–59, 1993.
- [62] B. Johansson, T. Elfving, V. Kozlov, Y. Censor, P.-E. Forssén, and G. Granlund, "The application of an oblique-projected Landweber method to a model of supervised learning," *Math. comput. model.*, vol. 43, no. 7–8, pp. 892–909, 2006.
- [63] C. D. Meyer, *Matrix analysis and applied linear algebra*. SIAM, 2000, vol. 71.
- [64] A. Dogandžić, R. Gu, and K. Qiu, "Mask iterative hard thresholding algorithms for sparse image reconstruction of objects with known contour," in *Conf. Rec. Asilomar Conf. Signals Syst. Comput.* IEEE, 2011, pp. 2111–2116.
- [65] Z. Gao, C. Zhang, Z. Wang, and S. Chen, "Prior-information aided iterative hard threshold: A low-complexity high-accuracy compressive sensing based channel estimation for TDS-OFDM," *IEEE Trans. Wireless Commun.*, vol. 14, no. 1, pp. 242–251, 2014.
- [66] C. Stöckle, J. Munir, A. Mezghani, and J. A. Nossek, "Channel estimation in massive MIMO systems using 1-bit quantization," in *IEEE Workshop Signal Process. Adv. Wirel. Commun.* IEEE, 2016, pp. 1–6.
- [67] C. Stoeckle, J. Munir, A. Mezghani, and J. A. Nossek, "DOA estimation performance and computational complexity of subspace-and compressed sensing-based methods," in *Int. ITG Workshop Smart Antennas*. VDE, 2015, pp. 1–6.
- [68] Y. Yu, Z. Sun, W. Zhu, and J. Gu, "A homotopy iterative hard thresholding algorithm with extreme learning machine for scene recognition," *IEEE Access*, vol. 6, pp. 30424–30436, 2018.
- [69] A. Tarantola, *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.
- [70] W. W. Hager, "Minimizing a quadratic over a sphere," *SIAM J. Optim.*, vol. 12, no. 1, pp. 188–208, 2001.
- [71] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Found. Comput. Math.*, vol. 9, no. 6, p. 717, 2009.
- [72] P. Jain and P. Netrapalli, "Fast exact matrix completion with finite samples," in *Proc. Conf. Learn. Theory*, 2015, pp. 1007–1034.
- [73] L. Ding and Y. Chen, "Leave-one-out approach for matrix completion: Primal and dual analysis," *IEEE Trans. Inf. Theory*, vol. 66, no. 11, pp. 7274–7301, 2020.
- [74] C. Eckart and G. Young, "The approximation of one matrix by another of lower rank," *Psychometrika*, vol. 1, no. 3, pp. 211–218, 1936.
- [75] M. Lee John, "Introduction to smooth manifolds," *Graduate Texts in Mathematics*, vol. 218, 2003.
- [76] F. Feppon and P. F. Lermusiaux, "A geometric approach to dynamical model order reduction," *SIAM J. Matrix Anal. Appl.*, vol. 39, no. 1, pp. 510–538, 2018.
- [77] A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler, "Projected gradient descent on Riemannian manifolds with applications to online power system optimization," in *Annu. Allert. Conf. Commun. Control Comput.* IEEE, 2016, pp. 225–232.
- [78] J.-F. Cai, T. Wang, and K. Wei, "Spectral compressed sensing via projected gradient descent," *SIAM J. Optim.*, vol. 28, no. 3, pp. 2625–2653, 2018.
- [79] J. P. Cunningham and Z. Ghahramani, "Linear dimensionality reduction: Survey, insights, and generalizations," *J. Mach. Learn. Res.*, vol. 16, no. 1, pp. 2859–2900, 2015.

