

On the Suboptimality of Proximal Gradient Descent for ℓ^0 Sparse Approximation

Yingzhen Yang
Snap Research, USA

SUPERYYZG@GMAIL.COM

Jiashi Feng
Department of ECE, National University of Singapore, Singapore

ELEFJIA@NUS.EDU.SG

Nebojsa Jojic
MICROSOFT RESEARCH, USA

JOJIC@MICROSOFT.COM

Jianchao Yang
Snap Research, USA

JIANCHAO.YANG@SNAP.COM

Thomas S. Huang
Beckman Institute, University of Illinois at Urbana-Champaign, USA

T-HUANG1@ILLINOIS.EDU

Editor:

Abstract

We study the proximal gradient descent (PGD) method for ℓ^0 sparse approximation problem as well as its accelerated optimization with randomized algorithms in this paper. We first offer theoretical analysis of PGD showing the bounded gap between the sub-optimal solution by PGD and the globally optimal solution for the ℓ^0 sparse approximation problem under conditions weaker than Restricted Isometry Property widely used in compressive sensing literature. Moreover, we propose randomized algorithms to accelerate the optimization by PGD using randomized low rank matrix approximation (PGD-RMA) and randomized dimension reduction (PGD-RDR). Our randomized algorithms substantially reduces the computation cost of the original PGD for the ℓ^0 sparse approximation problem, and the resultant sub-optimal solution still enjoys provable suboptimality, namely, the sub-optimal solution to the reduced problem still has bounded gap to the globally optimal solution to the original problem.

1. Introduction

In this paper, we consider the ℓ^0 sparse approximation problem, also named the ℓ^0 penalized Least Square Estimation (LSE) problem

$$\min_{\mathbf{z} \in \mathbb{R}^n} L(\mathbf{z}) = \|\mathbf{x} - \mathbf{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_0 \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is a signal in n -dimensional Euclidean space, \mathbf{D} is the design matrix of dimension $d \times n$ which is also called a dictionary with n atoms in the sparse coding literature.

The goal of problem (1) is to approximately represent signal \mathbf{x} by the atoms of the dictionary \mathbf{D} while requiring the representation \mathbf{z} to be sparse. Due to the nonconvexity imposed by the ℓ^0 -norm, extensive existing works resort to solve its ℓ^1 relaxation

$$\min_{\mathbf{z} \in \mathbb{R}^n} \|\mathbf{x} - \mathbf{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_1 \quad (2)$$

(2) is convex and also known as Basis Pursuit Denoising which can be solved efficiently by linear programming or iterative shrinkage algorithms (Daubechies et al., 2004; Elad, 2006; Bredies and Lorenz, 2008). Albeit the nonconvexity of (1), sparse coding methods such as (Mancera and Portilla, 2006; Bao et al., 2014) that directly optimize virtually the same objective as (1) demonstrate compelling performance compared to its ℓ^1 norm counterpart in various application domains such as data mining, applied machine learning and computer vision. Cardinality constraint in terms of ℓ^0 -norm is also studied for M-estimation problems (Jain et al., 2014).

We use the proximal gradient descent (PGD) method to obtain a sub-optimal solution to (1) in an iterative shrinkage manner with theoretical guarantee. Although the Iterative Hard-Thresholding (IHT) algorithm proposed by (Blumensath and Davies, 2008) also features iterative shrinkage, we prove the bound for gap between the sub-optimal solution and the globally optimal solution to (1). Our result of the bounded gap only requires nonsingularity of the submatrix of \mathbf{D} with the columns in the support of the sub-optimal and globally optimal solution (see the subsection ‘‘Assumptions for Our Analysis’’ for details). To the best of our knowledge, this is the first result analyzing the gap between sub-optimal solution and the globally optimal solution for the important ℓ^0 sparse approximation problem under assumptions weaker than Restricted Isometry Property (RIP) (Candes, 2008). The most related research is presented in in (Zhang and Zhang, 2012) where the distance between two local solutions of the concave regularized problems under much more restrictive assumptions including sparse eigenvalues of the dictionary. Moreover, our results suggest the merit of sparse initialization.

Furthermore, we propose to accelerate PGD for the ℓ^0 sparse approximation problem by two randomized algorithms. We propose proximal gradient descent via Randomized Matrix Approximation (PGD-RMA) which employs rank- k approximation of the dictionary via random projection. PGD-RMA reduces the cost of computing the gradient during the gradient descent step of PGD from $\mathcal{O}(dn)$ to $\mathcal{O}(dk + nk)$ by solving the reduced problem instead of the original problem with $k \ll \min\{d, n\}$, for a dictionary of size $d \times n$. The second randomized algorithm is proximal gradient descent via Randomized Dimension Reduction (PGD-RDR) which employs random projection to generate dimensionality-reduced signal and dictionary. PGD-RDR reduces the computational cost of the gradient descent step of PGD from $\mathcal{O}(dn)$ to $\mathcal{O}(mn)$ with $m < d$. While previous research focuses on the theoretical guarantee for convex problems via such randomized optimization (Drineas et al., 2011; Zhang et al., 2016b), we present the gap between the sub-optimal solution to the reduced problem and the globally optimal solution to the original problem. Our result establishes provable and efficient optimization by randomized low rank matrix decomposition and randomized dimension reduction for the nonconvex and nonsmooth ℓ^0 sparse approximation problem, while very few results are available in the literature in this direction.

Notations

Throughout this paper, we use bold letters for matrices and vectors, regular lower letter for scalars. The bold letter with subscript indicates the corresponding element of a matrix or vector, and $\|\cdot\|_p$ denote the ℓ^p -norm of a vector, or the p -norm of a matrix. We let $\beta_{\mathbf{I}}$ denote the vector formed by the elements of β with indices in \mathbf{I} when β is a vector, or matrix formed by columns of β with indices being the nonzero elements of \mathbf{I} when β is a matrix. $\text{supp}(\cdot)$ indicates the support of a vector, i.e. the set of indices of nonzero elements of this vector. $\sigma_{\min}(\cdot)$ and $\sigma_{\max}(\cdot)$ indicate the smallest and largest nonzero singular value of a matrix.

2. Proximal Gradient Descent for ℓ^0 Sparse Approximation

Solving the ℓ^0 sparse approximation problem (1) is NP-hard in general (Natarajan, 1995). Therefore, the literature extensively resorts to approximate algorithms, such as Orthogonal Matching Pursuit (Tropp, 2004), or that using surrogate functions (Hyder and Mahata, 2009), for ℓ^0 problems. In addition, PGD has been used by (Bao et al., 2014) to find the approximate solution to (1) with sublinear convergence to the critical point of the objective of (1), as well as satisfactory empirical results. The success of PGD raises the interesting question that how good the approximate solution by PGD is.

In this section, we first present the algorithm that employs PGD to optimize (1) in an iterative shrinkage manner. Then we show the suboptimality of the sub-optimal solution by PGD in terms of the gap between the sub-optimal solution and the globally optimal solution.

2.1 Algorithm

In t -th iteration of PGD for $t \geq 1$, gradient descent is performed on the squared loss term of $L(\mathbf{z})$, i.e. $Q(\mathbf{z}) \triangleq \|\mathbf{x} - \mathbf{D}\mathbf{z}\|_2^2$, to obtain

$$\tilde{\mathbf{z}}^{(t)} = \mathbf{z}^{(t-1)} - \frac{2}{\tau s} (\mathbf{D}^\top \mathbf{D} \mathbf{z}^{(t-1)} - \mathbf{D}^\top \mathbf{x}) \quad (3)$$

where τ is any constant that is greater than 1. $s > 0$ is usually chosen as the Lipschitz constant for the gradient of function $Q(\cdot)$, namely

$$\|\nabla Q(\mathbf{y}) - \nabla Q(\mathbf{z})\|_2 \leq s \|\mathbf{y} - \mathbf{z}\|_2, \quad \forall \mathbf{y}, \mathbf{z} \in \mathbb{R}^n \quad (4)$$

$\mathbf{z}^{(t)}$ is then the solution to the following the proximal mapping:

$$\mathbf{z}^{(t)} = \arg \min_{\mathbf{v} \in \mathbb{R}^n} \frac{\tau s}{2} \|\mathbf{v} - \tilde{\mathbf{z}}^{(t)}\|_2^2 + \lambda \|\mathbf{v}\|_0 \quad (5)$$

and (5) admits the closed-form solution:

$$\mathbf{z}^{(t)} = h_{\sqrt{\frac{2\lambda}{\tau s}}}(\tilde{\mathbf{z}}^{(t)}) \quad (6)$$

where h_θ is an element-wise hard thresholding operator:

$$[h_\theta(\mathbf{u})]_j = \begin{cases} 0 & : |\mathbf{u}_j| < \theta \\ \mathbf{u}_j & : \text{otherwise} \end{cases}, \quad 1 \leq j \leq n$$

The iterations start from $t = 1$ and continue until the sequence $\{L(\mathbf{z}^{(t)})\}_t$ or $\{\mathbf{z}^{(t)}\}_t$ converges or maximum iteration number is achieved. The optimization algorithm for the ℓ^0 sparse approximation problem (1) by PGD is described in Algorithm 1. In practice, the time complexity of optimization by PGD is $\mathcal{O}(Mdn)$ where M is the number of iterations (or maximum number of iterations) for PGD.

Algorithm 1 Proximal Gradient Descent for the ℓ^0 Sparse Approximation (1)

Input:

The given signal $\mathbf{x} \in \mathbb{R}^d$, the dictionary \mathbf{D} , the parameter λ for the weight of the ℓ^0 -norm, maximum iteration number M , stopping threshold ε , the initialization $\mathbf{z}^{(0)} \in \mathbb{R}^n$.

- 1: Obtain the sub-optimal solution $\tilde{\mathbf{z}}$ by the proximal gradient descent (PGD) method with (3) and (6) starting from $t = 1$. The iteration terminates either $\{\mathbf{z}^{(t)}\}_t$ or $\{L(\mathbf{z}^{(t)})\}_t$ converges under certain threshold or the maximum iteration number is achieved.

Output: Obtain the sparse code $\hat{\mathbf{z}}$ upon the termination of the iterations.

2.2 Theoretical Analysis

In this section we present the bound for the gap between the sub-optimal solution by PGD in Algorithm 1 and the globally optimal solution for the ℓ^0 sparse approximation problem (1). With proper initialization $\mathbf{z}^{(0)}$, we show that the sub-optimal solution by PGD is actually a critical point of $L(\mathbf{z})$ in Lemma 2, namely the sequence $\{\mathbf{z}^{(t)}\}_t$ converges to a critical point of the objective (1). We then show that both this sub-optimal solution and the globally optimal solution to (1) are local solutions of a carefully designed capped- ℓ^1 regularized problem in Lemma 3. The bound for ℓ^2 -distance between the sub-optimal solution and the globally optimal solution is then presented in Theorem 1. In the following analysis, we let $\mathbf{S} = \text{supp}(\mathbf{z}^{(0)})$. Also, since ℓ^0 is invariant to scaling, the original problem (1) is equivalent to $\min_{\mathbf{z} \in \mathbb{R}^n} \|\mathbf{x} - \frac{\mathbf{D}}{m}\mathbf{z}'\|_2^2 + \lambda\|\mathbf{z}'\|_0$ for $\mathbf{z}' = m\mathbf{z}$ with $m > \max_i \|\mathbf{D}^i\|_2$, so it is safe to assume $\max_i \|\mathbf{D}^i\|_2 \leq 1$. Without loss of generality, we let $\|\mathbf{x}\|_2 \leq 1$.

Lemma 1 (*Support shrinkage and sufficient decrease of the objective function*) Choose $\mathbf{z}^{(0)}$ such that $\|\mathbf{x} - \mathbf{D}\mathbf{z}^{(0)}\|_2 \leq 1$. When $s > \max\{2|\mathbf{S}|, \frac{2(1+\lambda|\mathbf{S}|)}{\lambda\tau}\}$, then

$$\text{supp}(\mathbf{z}^{(t)}) \subseteq \text{supp}(\mathbf{z}^{(t-1)}), t \geq 1 \tag{7}$$

namely the support of the sequence $\{\mathbf{z}^{(t)}\}_t$ shrinks. Moreover, the sequence of the objective $\{L(\mathbf{z}^{(t)})\}_t$ decreases, and the following inequality holds for $t \geq 1$:

$$L(\mathbf{z}^{(t)}) \leq L(\mathbf{z}^{(t-1)}) - \frac{(\tau - 1)s}{2} \|\mathbf{z}^{(t)} - \mathbf{z}^{(t-1)}\|_2^2 \tag{8}$$

And it follows that the sequence $\{L(\mathbf{z}^{(t)})\}_t$ converges.

Remark 1 One can always choose $\mathbf{z}^{(0)}$ as the optimal solution to the ℓ^1 regularized problem, so $\|\mathbf{x} - \mathbf{D}\mathbf{z}^{(0)}\|_2^2 + \lambda\|\mathbf{z}^{(0)}\|_1 \leq \|\mathbf{x}\|_2^2 \leq 1$, and it follows that $\|\mathbf{x} - \mathbf{D}\mathbf{z}^{(0)}\|_2^2 \leq 1$. Also, when not every subspace spanned by linearly independent columns of \mathbf{D} is orthogonal to \mathbf{x} (which

is common in practice), we can always find $\mathbf{z}^{(0)}$ such that $\|\mathbf{x} - \mathbf{D}\mathbf{z}^{(0)}\|_2^2 \leq 1$ and $\mathbf{z}^{(0)}$ is a nonzero vector. The support shrinkage property by Lemma 1 shows that the original ℓ^0 sparse approximation (1) is equivalent to a dictionary-reduced version

$$\min_{\mathbf{z} \in \mathbb{R}^n} L(\mathbf{z}) = \|\mathbf{x} - \mathbf{D}\mathbf{S}\mathbf{z}\|_2^2 + \lambda\|\mathbf{z}\|_0 \quad (9)$$

since PGD would not choose dictionary atoms outside of $\mathbf{D}\mathbf{S}$, so the numerical computation of PGD can be improved by using $\mathbf{D}\mathbf{S}$ as the dictionary with $s > \max\{2|\mathbf{S}|, \frac{2(1+\lambda|\mathbf{S}|)}{\lambda\tau}\}$. When $\mathbf{z}^{(0)}$ is sparse, $\mathbf{D}\mathbf{S}$ has a small number of atoms which enables fast optimization of the reduced problem (9). Also note that s can be much smaller than the Lipschitz constant for the gradient of function $Q(\cdot)$ by the choice indicated by Lemma 1, which leads to a larger step size for the gradient descent step (3).

The definition of critical points are defined below which is important for our analysis.

Definition 1 (Critical points) Given the non-convex function $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ which is a proper and lower semi-continuous function.

- for a given $\mathbf{x} \in \text{dom}f$, its Frechet subdifferential of f at \mathbf{x} , denoted by $\tilde{\partial}f(x)$, is the set of all vectors $\mathbf{u} \in \mathbb{R}^n$ which satisfy

$$\limsup_{\mathbf{y} \neq \mathbf{x}, \mathbf{y} \rightarrow \mathbf{x}} \frac{f(\mathbf{y}) - f(\mathbf{x}) - \langle \mathbf{u}, \mathbf{y} - \mathbf{x} \rangle}{\|\mathbf{y} - \mathbf{x}\|} \geq 0$$

- The limiting-subdifferential of f at $\mathbf{x} \in \mathbb{R}^n$, denoted by written $\partial f(x)$, is defined by

$$\partial f(x) = \{\mathbf{u} \in \mathbb{R}^n : \exists \mathbf{x}^k \rightarrow \mathbf{x}, f(\mathbf{x}^k) \rightarrow f(\mathbf{x}), \tilde{\mathbf{u}}^k \in \tilde{\partial}f(\mathbf{x}^k) \rightarrow \mathbf{u}\}$$

The point \mathbf{x} is a critical point of f if $0 \in \partial f(x)$.

If $\mathbf{D}\mathbf{S}$ is nonsingular, Lemma 2 shows that the sequences $\{\mathbf{z}^{(t)}\}_t$ produced by PGD converges to a critical point of $L(\mathbf{z})$, the objective of the ℓ^0 sparse approximation problem (1).

Lemma 2 With $\mathbf{z}^{(0)}$ and s in Lemma 1, if $\mathbf{D}\mathbf{S}$ is nonsingular, then the sequence $\{\mathbf{z}^{(t)}\}_t$ generated by PDG with (3) and (6) converges to a critical point of $L(\mathbf{z})$.

Denote the critical point of $L(\mathbf{z})$ by $\hat{\mathbf{z}}$ that the sequence $\{\mathbf{z}^{(t)}\}_t$ converges to when the assumption of Lemma 2 holds, and denote by \mathbf{z}^* the globally optimal solution to the ℓ^0 sparse approximation problem (1). Also, we consider the following capped- ℓ^1 regularized problem, which replaces the noncontinuous ℓ^0 -norm with the continuous capped- ℓ^1 regularization term R :

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^n} L_{\text{capped-}\ell^1}(\boldsymbol{\beta}) = \|\mathbf{x} - \mathbf{D}\boldsymbol{\beta}\|_2^2 + \mathbf{R}(\boldsymbol{\beta}; b) \quad (10)$$

where $\mathbf{R}(\boldsymbol{\beta}; b) = \sum_{j=1}^n R(\beta_j; b)$, $R(t; b) = \lambda \frac{\min\{|t|, b\}}{b}$ for some $b > 0$. It can be seen that $R(t; b)$

approaches the ℓ^0 -norm when $b \rightarrow 0+$. Our following theoretical analysis aims to obtain the gap between $\hat{\mathbf{z}}$ and \mathbf{z}^* . For the sake of this purpose, the definition of local solution and degree of nonconvexity of a regularizer are necessary and presented below.

Definition 2 (*Local solution*) A vector $\tilde{\boldsymbol{\beta}}$ is a local solution to the problem (10) if

$$\|2\mathbf{D}^\top(\mathbf{D}\tilde{\boldsymbol{\beta}} - \mathbf{x}) + \dot{\mathbf{R}}(\tilde{\boldsymbol{\beta}}; b)\|_2 = 0 \quad (11)$$

where $\dot{\mathbf{R}}(\tilde{\boldsymbol{\beta}}; b) = [\dot{R}(\tilde{\boldsymbol{\beta}}_1; b), \dot{R}(\tilde{\boldsymbol{\beta}}_2; b), \dots, \dot{R}(\tilde{\boldsymbol{\beta}}_n; b)]^\top$.

Note that in the above definition and the following text, $\dot{R}(t; b)$ can be chosen as any value between the right differential $\frac{\partial R}{\partial t}(t+; b)$ (or $\dot{R}(t+; b)$) and left differential $\frac{\partial R}{\partial t}(t-; b)$ (or $\dot{R}(t-; b)$).

Definition 3 (*Degree of nonconvexity of a regularizer*) For $\kappa \geq 0$ and $t \in \mathbb{R}$, define

$$\theta(t, \kappa) := \sup_s \{-\text{sgn}(s - t)(\dot{P}(s; b) - \dot{P}(t; b)) - \kappa|s - t|\}$$

as the degree of nonconvexity for function P . If $\mathbf{u} = (u_1, \dots, u_n)^\top \in \mathbb{R}^n$, $\theta(\mathbf{u}, \kappa) = [\theta(u_1, \kappa), \dots, \theta(u_n, \kappa)]$. sgn is a sign function.

Note that $\theta(t, \kappa) = 0$ if P is a convex function.

Let $\hat{\mathbf{S}} = \text{supp}(\hat{\mathbf{z}})$, $\mathbf{S}^* = \text{supp}(\mathbf{z}^*)$, the following lemma shows that both $\hat{\mathbf{z}}$ and \mathbf{z}^* are local solutions to the capped- ℓ^1 regularized problem (10).

Lemma 3 With $\mathbf{z}^{(0)}$ and s in Lemma 1, if

$$0 < b < \min\left\{\min_{j \in \hat{\mathbf{S}}} |\hat{\mathbf{z}}_j|, \frac{\lambda}{\max_{j \notin \hat{\mathbf{S}}} \left| \frac{\partial Q}{\partial \mathbf{z}_j} \Big|_{\mathbf{z}=\hat{\mathbf{z}}} \right|}, \min_{j \in \mathbf{S}^*} |\mathbf{z}_j^*|, \frac{\lambda}{\max_{j \notin \mathbf{S}^*} \left| \frac{\partial Q}{\partial \mathbf{z}_j} \Big|_{\mathbf{z}=\mathbf{z}^*} \right|}\right\} \quad (12)$$

(if the denominator is 0, $\frac{\lambda}{0}$ is defined to be $+\infty$ in the above inequality), then both $\hat{\mathbf{z}}$ and \mathbf{z}^* are local solutions to the capped- ℓ^1 regularized problem (10).

Theorem 1 (*Sub-optimal solution is close to the globally optimal solution*) With $\mathbf{z}^{(0)}$ and s in Lemma 1, and suppose $\mathbf{D}_{\hat{\mathbf{S}} \cup \mathbf{S}^*}$ is not singular with $\kappa_0 \triangleq \sigma_{\min}(\mathbf{D}_{\hat{\mathbf{S}} \cup \mathbf{S}^*}) > 0$. When $\kappa_0^2 > \kappa > 0$ and b is chosen according to (12) as in Lemma 3, let $\mathbf{F} = (\hat{\mathbf{S}} \setminus \mathbf{S}^*) \cup (\mathbf{S}^* \setminus \hat{\mathbf{S}})$ be the symmetric difference between $\hat{\mathbf{S}}$ and \mathbf{S}^* , then

$$\|\hat{\mathbf{z}} - \mathbf{z}^*\|_2 \leq \frac{\left(\sum_{j \in \mathbf{F} \cap \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\hat{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{F} \setminus \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}}}{2\kappa_0^2 - \kappa} \quad (13)$$

It is worthwhile to connect the assumption on the dictionary in Theorem 1 to the Restricted Isometry Property (RIP) (Candes and Tao, 2005) used frequently in the compressive sensing literature. Before the discussion, the definition of sparse eigenvalues is defined below.

Definition 4 (*Sparse eigenvalues*) The lower and upper sparse eigenvalues of a matrix \mathbf{A} are defined as

$$\kappa_-(m) := \min_{\|\mathbf{u}\|_0 \leq m; \|\mathbf{u}\|_2 = 1} \|\mathbf{A}\mathbf{u}\|_2^2 \quad \kappa_+(m) := \max_{\|\mathbf{u}\|_0 \leq m, \|\mathbf{u}\|_2 = 1} \|\mathbf{A}\mathbf{u}\|_2^2$$

Assumptions for Our Analysis

Typical RIP requires bounds such as $\delta_\tau + \delta_{2\tau} + \delta_{3\tau} < 1$ or $\delta_{2\tau} < \sqrt{2} - 1$ (Candes, 2008) for stably recovering the signal from measurements and τ is the sparsity of the signal, where $\delta_\tau = \max\{\kappa_+(\tau) - 1, 1 - \kappa_-(\tau)\}$. It should be emphasized that our bound (13) only requires nonsingularity of the submatrix of \mathbf{D} with the columns in the support of $\hat{\mathbf{z}}$ and \mathbf{z}^* , which are more general than RIP in the sense of not requiring bounds in terms of δ . To see this point, choose the initialization \mathbf{z}^0 such that $|\text{supp}(\mathbf{z}^0)| \leq \tau$, then the condition $\delta_{2\tau} < \sqrt{2} - 1$ (Candes, 2008) indicates that $\sigma_{\min}(\mathbf{D}_{\hat{\mathbf{S}} \cup \mathbf{S}^*}) > 2 - \sqrt{2}$, which is stronger than our assumption that $\sigma_{\min}(\mathbf{D}_{\hat{\mathbf{S}} \cup \mathbf{S}^*}) > 0$. In addition, our assumption on the nonsingularity of $\mathbf{D}_{\hat{\mathbf{S}} \cup \mathbf{S}^*}$ is much weaker than that of the sparse eigenvalues used in RIP conditions which require minimum eigenvalue of every submatrix of \mathbf{D} of specified number of columns.

Remark 2 *If $\mathbf{z}^{(0)}$ is sparse, $\hat{\mathbf{z}}$ is also sparse by the property of support shrinkage in Lemma 1. We can then expect that $|\hat{\mathbf{S}} \cup \mathbf{S}^*|$ is reasonably small, and a small $|\hat{\mathbf{S}} \cup \mathbf{S}^*|$ often increases the chance of a larger $\sigma_{\min}(\mathbf{D}_{|\hat{\mathbf{S}} \cup \mathbf{S}^*|})$. Also note that the bound for distance between the sub-optimal solution and the globally optimal solution presented in Theorem 1 does not require typical RIP conditions. Moreover, when $\frac{\lambda}{b} - \kappa|\hat{\mathbf{z}}_j - b|$ for nonzero $\hat{\mathbf{z}}_j$ and $\frac{\lambda}{b} - \kappa b$ are no greater than 0, or they are small positive numbers, the sub-optimal solution $\hat{\mathbf{z}}$ is equal to or very close to the globally optimal solution.*

3. Accelerated Proximal Gradient Descent by Randomized Algorithms

In this section, we propose and analyze two randomized algorithms that accelerate PGD for the ℓ^0 sparse approximation problem by random projection. Our first algorithm employs randomized low rank matrix approximation for the dictionary \mathbf{D} , and the second algorithm uses random projection to generate dimensionality-reduced signal and dictionary and then perform PGD on the low-dimensional signal and dictionary. Our theoretical analysis establishes the suboptimality of the solutions obtained by the proposed randomized algorithms.

3.1 Algorithm

While Lemma 1 shows that the ℓ^0 sparse approximation (1) is equivalent to its dictionary-reduced version (9) which leads to improved efficiency, we are still facing the computational challenge incurred by dictionary with large dimension d and size n (or large $|\mathbf{S}|$). The literature has extensively employed randomized algorithms for accelerating the numerical computation of different kinds of matrix optimization problems including low rank approximation and matrix decomposition (Frieze et al., 2004; Drineas et al., 2004; Sarlos, 2006; Drineas et al., 2006, 2008; Mahoney and Drineas, 2009; Drineas et al., 2011; Lu et al., 2013). In order to accelerate the numerical computation involved in PGD, we propose two randomized algorithms. The first algorithm adopts the randomized low rank approximation by random projection (Halko et al., 2011) to obtain a low rank approximation of the dictionary so as to accelerate the computation of gradient for PGD. The second algorithm generates dimensionality-reduced signal and dictionary by random projection and then apply PGD upon the low-dimensional signal and dictionary for improved efficiency. The optimization algorithm for the ℓ^0 sparse approximation problem (1) by PGD with low rank approximation of \mathbf{D} via Randomized Matrix Approximation, termed PGD-RMA in

this paper, is described in Section 3.1.1. Bearing the idea of using randomized algorithm for dimension reduction, the algorithm that accelerates PGD via randomized dimension reduction, termed PGD-RDR, is introduced in Section 3.1.2.

3.1.1 ACCELERATED PROXIMAL GRADIENT DESCENT VIA RANDOMIZED MATRIX APPROXIMATION: PGD-RMA

The procedure of PGD-RMA is described as follows. A random matrix $\Omega \in \mathbb{R}^{n \times k}$ is computed such that each element Ω_{ij} is sampled independently from the Gaussian distribution $\mathcal{N}(0, 1)$. With the QR decomposition of $\mathbf{D}\Omega$, i.e. $\mathbf{D}\Omega = \mathbf{Q}\mathbf{R}$ where $\mathbf{Q} \in \mathbb{R}^{d \times k}$ is an orthogonal matrix of rank k and $\mathbf{R} \in \mathbb{R}^{k \times k}$ is an upper triangle matrix. The columns of \mathbf{Q} form the orthogonal basis for $\mathbf{D}\Omega$. Then \mathbf{D} is approximated by projecting \mathbf{D} onto the range of $\mathbf{D}\Omega$: $\mathbf{Q}\mathbf{Q}^\top \mathbf{D} = \mathbf{Q}\mathbf{W} = \tilde{\mathbf{D}}$ where $\mathbf{W} = \mathbf{Q}^\top \mathbf{D} \in \mathbb{R}^{k \times n}$. Replacing \mathbf{D} with its low rank approximation $\tilde{\mathbf{D}}$, we resort to solve the following reduced ℓ^0 sparse approximation problem (14)

$$\min_{\mathbf{z} \in \mathbb{R}^n} \tilde{L}(\mathbf{z}) = \|\mathbf{x} - \tilde{\mathbf{D}}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_0 \quad (14)$$

And the first step of PGD (3) for the original ℓ^0 sparse approximation problem is reduced to

$$\begin{aligned} \tilde{\mathbf{z}}^{(t)} &= \mathbf{z}^{(t-1)} - \frac{2}{\tau s} (\tilde{\mathbf{D}}^\top \tilde{\mathbf{D}} \mathbf{z}^{(t-1)} - \tilde{\mathbf{D}}^\top \mathbf{x}) \\ &= \mathbf{z}^{(t-1)} - \frac{2}{\tau s} (\mathbf{W}^\top \mathbf{Q}^\top \mathbf{Q} \mathbf{W} \mathbf{z}^{(t-1)} - \mathbf{W}^\top \mathbf{Q}^\top \mathbf{x}) \end{aligned} \quad (15)$$

The complexity of this step is reduced from $\mathcal{O}(dn)$ to $\mathcal{O}(dk + nk)$ wherein $k \ll \min\{d, n\}$ and significant efficiency improvement is achieved. Note that the computational cost of QR decomposition for $\mathbf{D}\Omega$ is less than $2dk^2$, which is acceptable with a small k .

The randomized algorithm PGD-RMA is described in Algorithm 2. The time complexity of PGD-RMA is $\mathcal{O}(M(dk + nk))$ where M is the number of iterations (or maximum number of iterations), compared to the complexity $\mathcal{O}(Mdn)$ for the original PGD.

Algorithm 2 Proximal Gradient Descent via Randomized Matrix Approximation (PGD-RMA) for the ℓ^0 Sparse Approximation (1)

Input:

The given signal $\mathbf{x} \in \mathbb{R}^d$, the dictionary \mathbf{D} , the parameter λ for the weight of the ℓ^0 -norm, maximum iteration number M , stopping threshold ε , the initialization $\mathbf{z}^{(0)} \in \mathbb{R}^n$.

- 1: Sample a random matrix $\Omega \in \mathbb{R}^{n \times k}$ by $\Omega_{ij} \sim \mathcal{N}(0, 1)$.
- 2: Compute the QR decomposition of $\mathbf{D}\Omega$: $\mathbf{D}\Omega = \mathbf{Q}\mathbf{R}$
- 3: Approximate \mathbf{D} by $\tilde{\mathbf{D}} = \mathbf{Q}\mathbf{W}$ where $\mathbf{W} = \mathbf{Q}^\top \mathbf{D}$
- 4: Perform PGD with (15) and (6) starting from $t = 1$. The iteration terminates either $\{\mathbf{z}^{(t)}\}_t$ or $\{L(\mathbf{z}^{(t)})\}_t$ converges under certain threshold or maximum iteration number is achieved.

Output: Obtain the sparse code $\tilde{\mathbf{z}}$ upon the termination of the iterations.

3.1.2 ACCELERATED PROXIMAL GRADIENT DESCENT VIA RANDOM DIMENSION REDUCTION: PGD-RDR

We introduce PGD-RDR which employs random projection to generate low-dimensional signal and dictionary, upon which PGD is applied for improved efficiency. The literature (Frankl and Maehara, 1987; Indyk and Motwani, 1998; Zhang et al., 2016a) extensively considers the random projection that satisfies the following ℓ^2 -norm preserving property, which is closely related to the proof of the Johnson-Lindenstrauss lemma (Dasgupta and Gupta, 2003).

Definition 5 *The linear operator $\mathbf{T}: \mathbb{R}^d \rightarrow \mathbb{R}^m$ satisfies the ℓ^2 -norm preserving property if there exists constant $c > 0$ such that*

$$\Pr \left[(1 - \varepsilon) \|\mathbf{v}\|_2 \leq \|\mathbf{T}\mathbf{v}\|_2 \leq (1 + \varepsilon) \|\mathbf{v}\|_2 \right] \geq 1 - 2e^{-\frac{m\varepsilon^2}{c}} \quad (16)$$

holds for any fixed $\mathbf{v} \in \mathbb{R}^d$ and $0 < \varepsilon \leq \frac{1}{2}$.

The linear operator \mathbf{T} satisfying the ℓ^2 -norm preserving property can be generated randomly according to uncomplicated distributions. With $\mathbf{T}' = \sqrt{m}\mathbf{T}$, it is proved in (Arriaga and Vempala, 2006; Achlioptas, 2003) that \mathbf{T} satisfies the ℓ^2 -norm preserving property, if all the elements of \mathbf{T}' are sampled independently from the Gaussian distribution $\mathcal{N}(0, 1)$, or uniform distribution over ± 1 , or the database-friendly distribution described by

$$\mathbf{T}'_{ij} = \begin{cases} \sqrt{3} & : \text{with probability } \frac{1}{6} \\ \sqrt{0} & : \text{with probability } \frac{2}{3} \\ -\sqrt{3} & : \text{with probability } \frac{1}{6} \end{cases}, 1 \leq i \leq m, 1 \leq j \leq d$$

With $m < d$, PGD-RDR first generate the dimensionality-reduced signal and dictionary by $\bar{\mathbf{x}} = \mathbf{T}\mathbf{x}$ and $\bar{\mathbf{D}} = \mathbf{T}\mathbf{D}$, then solve the following dimensionality-reduced ℓ^0 sparse approximation problem

$$\min_{\mathbf{z} \in \mathbb{R}^n} \bar{L}(\mathbf{z}) = \|\bar{\mathbf{x}} - \bar{\mathbf{D}}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_0 \quad (17)$$

by PGD. The procedure of PGD-RDR for the ℓ^0 sparse approximation problem (1) is described in Algorithm 3. The time complexity of sampling the random matrix \mathbf{T} is $\mathcal{O}(md)$, and the time complexity of the first step of PGD (3) for gradient descent is reduced from $\mathcal{O}(dn)$ to $\mathcal{O}(mn)$. The time complexity of PGD-RDR is $\mathcal{O}(Mmn)$ where M is the number of iterations (or maximum number of iterations), compared to the complexity $\mathcal{O}(Mdn)$ for the original PGD. Improvement on the efficiency is achieved with $m < d$.

3.2 Theoretical Analysis

We analyze the theoretical properties of the proposed PGD-RMA and PGD-RDR in the previous section. For both randomized algorithms, we present the bounded gap between the sub-optimal solution to the reduced ℓ^0 sparse approximation problem (14) or (17) and the globally optimal solution \mathbf{z}^* to the original problem.

3.3 Analysis for PGD-RMA

(Halko et al., 2011) proved that the approximation $\tilde{\mathbf{D}}$ is close to \mathbf{D} in terms of the spectral norm:

Algorithm 3 Proximal Gradient Descent via Randomized Dimension Reduction (PGD-RDR) for the ℓ^0 Sparse Approximation (1)

Input:

The given signal $\mathbf{x} \in R^d$, the dictionary \mathbf{D} , the parameter λ for the weight of the ℓ^0 -norm, maximum iteration number M , stopping threshold ε , the initialization $\mathbf{z}^{(0)} \in \mathbb{R}^n$.

- 1: Sample a random matrix $\mathbf{T} \in \mathbb{R}^{m \times n}$ which satisfies the ℓ^2 -norm preserving property in Definition 5, e.g. $\sqrt{m}\mathbf{T}_{ij} \sim \mathcal{N}(0, 1)$.
- 2: Compute the low-dimensional signal $\bar{\mathbf{x}} = \mathbf{T}\mathbf{x}$ and the dimensionality-reduced dictionary $\bar{\mathbf{D}} = \mathbf{T}\mathbf{D}$.
- 3: Perform PGD with (15) and (6) starting from $t = 1$, with \mathbf{x} and \mathbf{D} replaced by $\bar{\mathbf{x}}$ and $\bar{\mathbf{D}}$. The iteration terminates either $\{\mathbf{z}^{(t)}\}_t$ or $\{L(\mathbf{z}^{(t)})\}_t$ converges under certain threshold or maximum iteration number is achieved.

Output: Obtain the sparse code $\bar{\mathbf{z}}$ upon the termination of the iterations.

Lemma 4 (Corollary 10.9 by (Halko et al., 2011)) Let $k_0 \geq 2$ and $p = k - k_0 \geq 4$, then probability at least $1 - 6e^{-p}$, then the spectral norm of $\mathbf{D} - \hat{\mathbf{D}}$ is bounded by

$$\|\mathbf{D} - \hat{\mathbf{D}}\|_2 \leq C_{k,k_0} \quad (18)$$

where

$$C_{k,k_0} = \left(1 + 17\sqrt{1 + \frac{k_0}{p}}\right)\sigma_{k_0+1} + \frac{8\sqrt{k}}{p+1}\left(\sum_{j>k_0} \sigma_j^2\right)^{\frac{1}{2}} \quad (19)$$

$\sigma_1 \geq \sigma_2 \geq \dots$ are the singular values of \mathbf{D} .

Let $\tilde{\mathbf{z}}$ be the globally optimal solution to (14), $\tilde{\mathbf{S}} = \text{supp}(\tilde{\mathbf{z}})$, $\tilde{Q}(\mathbf{z}) = \|\mathbf{x} - \tilde{\mathbf{D}}\mathbf{z}\|_2^2$, $\mathbf{z}^{(0)}$ be the initialization for PGD for the optimization of the reduced problem (14). We have the following theorem showing the upper bound for the gap between $\tilde{\mathbf{z}}$ and \mathbf{z}^* .

Theorem 2 (Optimal solution to the reduced problem (14) is close to the that to the original problem) Let $\mathbf{G} = \tilde{\mathbf{S}} \cup \mathbf{S}^*$. Suppose $\mathbf{D}_{\mathbf{G}}$ is not singular with $\tau_0 \triangleq \sigma_{\min}(\mathbf{D}_{\mathbf{G}}) > 0$, $2\tau_0^2 > \tau > 0$. Then with probability at least $1 - 6e^{-p}$,

$$\begin{aligned} & \|\mathbf{z}^* - \tilde{\mathbf{z}}\|_2 \\ & \leq \frac{1}{2\tau_0^2 - \tau} \left(\left(\sum_{j \in \mathbf{G} \cap \tilde{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\tilde{\mathbf{z}}_j - b|\})^2 + \right. \right. \\ & \quad \left. \left. \sum_{j \in \mathbf{G} \setminus \tilde{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}} + 2C_{k,k_0}M_0(2\sigma_{\max}(\mathbf{D}) + C_{k,k_0}) + 2C_{k,k_0}\|\mathbf{x}\|_2 \right) \end{aligned} \quad (20)$$

where $M_0 = \frac{\|\mathbf{x}\|_2 + \sqrt{L(\mathbf{z}^{(0)})}}{\tau_0}$, and b satisfies

$$0 < b < \min\left\{ \min_{j \in \tilde{\mathbf{S}}} |\tilde{\mathbf{z}}_j|, \max_{k \notin \tilde{\mathbf{S}}} \frac{\lambda}{\left(\frac{\partial \tilde{Q}}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\tilde{\mathbf{z}}} - \lambda\right)_+}, \min_{j \in \mathbf{S}^*} |\mathbf{z}_j^*|, \max_{k \notin \mathbf{S}^*} \frac{\lambda}{\left(\frac{\partial Q}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\mathbf{z}^*} - \lambda\right)_+} \right\} \quad (21)$$

Let $\mathbf{S} = \text{supp}(\mathbf{z}^{(0)})$. According to Theorem 1 and Theorem 2, we have the bounded gap between the sub-optimal solution to the reduced problem (14) and the globally optimal solution \mathbf{z}^* to the original problem (1).

Theorem 3 (Sub-optimal solution to the reduced problem (14) is close to the globally optimal solution to the original problem) Choose $\mathbf{z}^{(0)}$ such that $\|\mathbf{x} - \tilde{\mathbf{D}}\mathbf{z}^{(0)}\|_2 \leq 1$, and $s > \max\{2|\mathbf{S}|, \frac{2(1+\lambda|\mathbf{S}|)}{\lambda\tau}\}$, if $\tilde{\mathbf{D}}_{\mathbf{S}}$ is nonsingular, then the sequence $\{\mathbf{z}^{(t)}\}_t$ generated by PDG for the reduced problem (14) converges to a critical point of $\tilde{L}(\mathbf{z})$, denoted by $\hat{\mathbf{z}}$. Let $\hat{\mathbf{S}} = \text{supp}(\hat{\mathbf{z}})$, $\mathbf{F} = (\hat{\mathbf{S}} \setminus \tilde{\mathbf{S}}) \cup (\tilde{\mathbf{S}} \setminus \hat{\mathbf{S}})$, $\mathbf{G} = \hat{\mathbf{S}} \cup \mathbf{S}^*$. Suppose $\tilde{\mathbf{D}}_{\hat{\mathbf{S}} \cup \tilde{\mathbf{S}}}$ is not singular with $\kappa_0 \triangleq \sigma_{\min}(\tilde{\mathbf{D}}_{\hat{\mathbf{S}} \cup \tilde{\mathbf{S}}}) > 0$, and $\kappa_0^2 > \kappa > 0$; $\mathbf{D}_{\mathbf{G}}$ is not singular with $\tau_0 \triangleq \sigma_{\min}(\mathbf{D}_{\mathbf{G}}) > 0$, $2\tau_0^2 > \tau > 0$. Then with probability at least $1 - 6e^{-p}$,

$$\|\mathbf{z}^* - \hat{\mathbf{z}}\|_2 \leq b_1 + b_2, \quad (22)$$

where

$$b_1 = \frac{\left(\sum_{j \in \mathbf{F} \cap \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\hat{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{F} \setminus \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}}}{2\kappa_0^2 - \kappa}$$

$$b_2 = \frac{1}{2\tau_0^2 - \tau} \left(\left(\sum_{j \in \mathbf{G} \cap \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\hat{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{G} \setminus \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}} \right. \\ \left. + 2C_{k,k_0} M_0 (2\sigma_{\max}(\mathbf{D}) + C_{k,k_0}) + 2C_{k,k_0} \|\mathbf{x}\|_2 \right)$$

$M_0 = \frac{\|\mathbf{x}\|_2 + \sqrt{\tilde{L}(\mathbf{z}^{(0)})}}{\tau_0}$, and b satisfies

$$0 < b < \min \left\{ \min_{j \in \hat{\mathbf{S}}} |\hat{\mathbf{z}}_j|, \max_{k \notin \hat{\mathbf{S}}} \frac{\lambda}{\left(\frac{\partial \tilde{Q}}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\hat{\mathbf{z}}} - \lambda \right)_+}, \min_{j \in \mathbf{S}^*} |\mathbf{z}_j^*|, \max_{k \notin \mathbf{S}^*} \frac{\lambda}{\left(\frac{\partial Q}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\mathbf{z}^*} - \lambda \right)_+}, \min_{j \in \hat{\mathbf{S}}} |\hat{\mathbf{z}}_j|, \max_{k \notin \hat{\mathbf{S}}} \frac{\lambda}{\left(\frac{\partial \tilde{Q}}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\hat{\mathbf{z}}} - \lambda \right)_+} \right\} \quad (23)$$

3.4 Analysis for PGD-RDR

Let $\bar{\mathbf{z}}$ be the globally optimal solution to (17), $\bar{\mathbf{S}} = \text{supp}(\bar{\mathbf{z}})$, $\bar{Q}(\mathbf{z}) = \|\bar{\mathbf{x}} - \bar{\mathbf{D}}\mathbf{z}\|_2^2$, $\mathbf{z}^{(0)}$ be the initialization for PGD for the optimization of the reduced problem (17). We have the following theorem showing the upper bound for the gap between $\bar{\mathbf{z}}$ and \mathbf{z}^* .

Theorem 4 (Optimal solution to the dimensionality-reduced problem (17) is close to the that to the original problem) Let $\mathbf{H} = \bar{\mathbf{S}} \cup \mathbf{S}^*$. Suppose $\mathbf{D}_{\mathbf{H}}$ is not singular with $\eta_0 \triangleq \sigma_{\min}(\mathbf{D}_{\mathbf{H}}) > 0$, $2\eta_0^2 > \eta > 0$. If \mathbf{T} satisfies the ℓ^2 -norm preserving property in Definition 5, $m \geq 4c \log \frac{4}{\delta}$, then with probability at least $1 - \delta$,

$$\|\mathbf{z}^* - \bar{\mathbf{z}}\|_2 \leq \frac{1}{2\eta_0^2 - \eta} \left(\left(\sum_{j \in \mathbf{H} \cap \bar{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\bar{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{H} \setminus \bar{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}} + 2\|\mathbf{D}\|_F M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} (\sigma_{\max}(\mathbf{D}) + 1) \right) \quad (24)$$

where $M_1 = \frac{\|\mathbf{x}\|_2 + \sqrt{\bar{L}(\mathbf{z}^{(0)})}}{\eta_0}$, and b satisfies

$$0 < b < \min\left\{\min_{j \in \bar{\mathbf{S}}} |\bar{\mathbf{z}}_j|, \max_{k \notin \bar{\mathbf{S}}} \frac{\lambda}{\left(\frac{\partial \bar{Q}}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\bar{\mathbf{z}}} - \lambda\right)_+}, \min_{j \in \mathbf{S}^*} |\mathbf{z}_j^*|, \max_{k \notin \mathbf{S}^*} \frac{\lambda}{\left(\frac{\partial Q}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\mathbf{z}^*} - \lambda\right)_+}\right\} \quad (25)$$

Similar to the analysis for PGD-RMA, we let $\mathbf{S} = \text{supp}(\mathbf{z}^{(0)})$. Combining Theorem 1 and Theorem 4, we have the bounded gap between the sub-optimal solution to the dimensionality-reduced problem (17) and the globally optimal solution \mathbf{z}^* to the original problem (1).

Theorem 5 (Sub-optimal solution to the dimensionality-reduced problem (17) is close to the globally optimal solution to the original problem) Choose $\mathbf{z}^{(0)}$ such that $\|\mathbf{x} - \bar{\mathbf{D}}\mathbf{z}^{(0)}\|_2 \leq 1$, and $s > \max\{2|\mathbf{S}|, \frac{2(1+\lambda|\mathbf{S}|)}{\lambda\eta}\}$, if $\bar{\mathbf{D}}_{\mathbf{S}}$ is nonsingular, then the sequence $\{\mathbf{z}^{(t)}\}_t$ generated by PDG for the dimensionality-reduced problem (17) converges to a critical point of $\bar{L}(\mathbf{z})$, denoted by $\hat{\mathbf{z}}$. Let $\hat{\mathbf{S}} = \text{supp}(\hat{\mathbf{z}})$, $\mathbf{F} = (\hat{\mathbf{S}} \setminus \bar{\mathbf{S}}) \cup (\bar{\mathbf{S}} \setminus \hat{\mathbf{S}})$, $\mathbf{H} = \bar{\mathbf{S}} \cup \mathbf{S}^*$. Suppose $\mathbf{D}_{\hat{\mathbf{S}} \cup \bar{\mathbf{S}}}$ is not singular with $\kappa_0 \triangleq \sigma_{\min}(\bar{\mathbf{D}}_{\hat{\mathbf{S}} \cup \bar{\mathbf{S}}}) > 0$, and $\kappa_0^2 > \kappa > 0$; $\mathbf{D}_{\mathbf{H}}$ is not singular with $\eta_0 \triangleq \sigma_{\min}(\mathbf{D}_{\mathbf{H}}) > 0$, $2\eta_0^2 > \eta > 0$. If \mathbf{T} satisfies the ℓ^2 -norm preserving property in Definition 5, $m \geq 4c \log \frac{4}{\delta}$, then with probability at least $1 - \delta$,

$$\|\mathbf{z}^* - \hat{\mathbf{z}}\|_2 \leq b_1 + b_2, \quad (26)$$

where

$$b_1 = \frac{\left(\sum_{j \in \mathbf{F} \cap \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\hat{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{F} \setminus \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2\right)^{\frac{1}{2}}}{2\kappa_0^2 - \kappa}$$

$$b_2 = \frac{1}{2\eta_0^2 - \eta} \left(\left(\sum_{j \in \mathbf{H} \cap \bar{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\bar{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{H} \setminus \bar{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2\right)^{\frac{1}{2}} + 2\|\mathbf{D}\|_F M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} (\sigma_{\max}(\mathbf{D}) + 1) \right)$$

$M_0 = \frac{\|\mathbf{x}\|_2 + \sqrt{\bar{L}(\mathbf{z}^{(0)})}}{\eta_0}$, and b satisfies

$$0 < b < \min\left\{\min_{j \in \bar{\mathbf{S}}} |\bar{\mathbf{z}}_j|, \max_{k \notin \bar{\mathbf{S}}} \frac{\lambda}{\left(\frac{\partial \bar{Q}}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\bar{\mathbf{z}}} - \lambda\right)_+}, \min_{j \in \mathbf{S}^*} |\mathbf{z}_j^*|, \max_{k \notin \mathbf{S}^*} \frac{\lambda}{\left(\frac{\partial Q}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\mathbf{z}^*} - \lambda\right)_+}, \min_{j \in \hat{\mathbf{S}}} |\hat{\mathbf{z}}_j|, \max_{k \notin \hat{\mathbf{S}}} \frac{\lambda}{\left(\frac{\partial \bar{Q}}{\partial \mathbf{z}_k} \Big|_{\mathbf{z}=\hat{\mathbf{z}}} - \lambda\right)_+}\right\} \quad (27)$$

The detailed proofs of the theorems and lemmas are included in Section 5. Note that we slightly abuse the notation of \mathbf{F} , $\hat{\mathbf{S}}$, κ and κ_0 in the analysis for PGD-RMA and PGD-RDR with no confusion. To the best of our knowledge, our theoretical results are among the very few results for the provable randomized efficient algorithms for the nonsmooth and nonconvex ℓ^0 sparse approximation problem.

4. Conclusions

We propose to use proximal gradient descent (PGD) to obtain a sub-optimal solution to the ℓ^0 sparse approximation problem. Our theoretical analysis renders the bound for the

ℓ^2 -distance between the sub-optimal solution and the globally optimal solution, under conditions weaker than Restricted Isometry Property (RIP). To the best of our knowledge, this is the first time that such gap between sub-optimal solution and globally optimal solution is obtained under our mild conditions. Moreover, we propose provable randomized algorithms, namely proximal gradient descent via Randomized Matrix Approximation (PGD-RMA) and proximal gradient descent via Random Dimension Reduction (PGD-RDR), to accelerate the ordinary optimization by PGD.

References

- Dimitris Achlioptas. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *J. Comput. Syst. Sci.*, 66(4):671–687, June 2003. ISSN 0022-0000. doi: 10.1016/S0022-0000(03)00025-4. URL [http://dx.doi.org/10.1016/S0022-0000\(03\)00025-4](http://dx.doi.org/10.1016/S0022-0000(03)00025-4).
- Rosa I. Arriaga and Santosh Vempala. An algorithmic theory of learning: Robust concepts and random projection. *Machine Learning*, 63(2):161–182, 2006. ISSN 1573-0565. doi: 10.1007/s10994-006-6265-7. URL <http://dx.doi.org/10.1007/s10994-006-6265-7>.
- Chenglong Bao, Hui Ji, Yuhui Quan, and Zuowei Shen. L0 norm based dictionary learning by proximal methods with global convergence. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 3858–3865, 2014. doi: 10.1109/CVPR.2014.493.
- Thomas Blumensath and Mike E. Davies. Iterative thresholding for sparse approximations. *Journal of Fourier Analysis and Applications*, 14(5):629–654, 2008. ISSN 1531-5851. doi: 10.1007/s00041-008-9035-z.
- Jérôme Bolte, Shoham Sabach, and Marc Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Math. Program.*, 146(1-2):459–494, August 2014. ISSN 0025-5610. doi: 10.1007/s10107-013-0701-9.
- Kristian Bredies and Dirk A. Lorenz. Iterated hard shrinkage for minimization problems with sparsity constraints. *SIAM Journal on Scientific Computing*, 30(2):657–683, 2008. doi: 10.1137/060663556.
- E.J. Candes and T. Tao. Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12):4203–4215, 2005.
- Emmanuel J. Candes. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathématique*, 346(910):589 – 592, 2008. ISSN 1631-073X.
- Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Struct. Algorithms*, 22(1):60–65, January 2003. ISSN 1042-9832. doi: 10.1002/rsa.10073. URL <http://dx.doi.org/10.1002/rsa.10073>.
- I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004. ISSN 1097-0312. doi: 10.1002/cpa.20042.

- P. Drineas, A. Frieze, R. Kannan, S. Vempala, and V. Vinay. Clustering large graphs via the singular value decomposition. *Machine Learning*, 56(1):9–33, 2004. ISSN 1573-0565. doi: 10.1023/B:MACH.0000033113.59016.96. URL <http://dx.doi.org/10.1023/B:MACH.0000033113.59016.96>.
- Petros Drineas, Ravi Kannan, and Michael W. Mahoney. Fast monte carlo algorithms for matrices ii: Computing a low-rank approximation to a matrix. *SIAM Journal on Computing*, 36(1):158–183, 2006. doi: 10.1137/S0097539704442696.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error $\$cur\$$ matrix decompositions. *SIAM Journal on Matrix Analysis and Applications*, 30(2):844–881, 2008. doi: 10.1137/07070471X.
- Petros Drineas, Michael W. Mahoney, S. Muthukrishnan, and Tamás Sarlós. Faster least squares approximation. *Numerische Mathematik*, 117(2):219–249, 2011. ISSN 0945-3245. doi: 10.1007/s00211-010-0331-6. URL <http://dx.doi.org/10.1007/s00211-010-0331-6>.
- M. Elad. Why simple shrinkage is still relevant for redundant representations? *IEEE Transactions on Information Theory*, 52(12):5559–5569, Dec 2006. ISSN 0018-9448. doi: 10.1109/TIT.2006.885522.
- P. Frankl and H. Maehara. The johnson-lindenstrauss lemma and the sphericity of some graphs. *J. Comb. Theory Ser. A*, 44(3):355–362, June 1987. ISSN 0097-3165. URL <http://dl.acm.org/citation.cfm?id=48184.48193>.
- Alan Frieze, Ravi Kannan, and Santosh Vempala. Fast monte-carlo algorithms for finding low-rank approximations. *J. ACM*, 51(6):1025–1041, November 2004. ISSN 0004-5411. doi: 10.1145/1039488.1039494. URL <http://doi.acm.org/10.1145/1039488.1039494>.
- N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, May 2011. ISSN 0036-1445. doi: 10.1137/090771806.
- M. Hyder and K. Mahata. An approximate l_0 norm minimization algorithm for compressed sensing. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*, pages 3365–3368, April 2009.
- Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 604–613, New York, NY, USA, 1998. ACM. ISBN 0-89791-962-9. doi: 10.1145/276698.276876. URL <http://doi.acm.org/10.1145/276698.276876>.
- Prateek Jain, Ambuj Tewari, and Purushottam Kar. On iterative hard thresholding methods for high-dimensional m -estimation. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 685–693, 2014.

- Yichao Lu, Paramveer S. Dhillon, Dean Foster, and Lyle Ungar. Faster ridge regression via the subsampled randomized hadamard transform. In *Proceedings of the 26th International Conference on Neural Information Processing Systems*, NIPS'13, pages 369–377, USA, 2013. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=2999611.2999653>.
- Michael W. Mahoney and Petros Drineas. Cur matrix decompositions for improved data analysis. *Proceedings of the National Academy of Sciences*, 106(3):697–702, 2009. doi: 10.1073/pnas.0803205106. URL <http://www.pnas.org/content/106/3/697.abstract>.
- L. Mancera and J. Portilla. L0-norm-based sparse representation through alternate projections. In *Image Processing, 2006 IEEE International Conference on*, pages 2089–2092, Oct 2006.
- B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995. doi: 10.1137/S0097539792240406.
- T. Sarlos. Improved approximation algorithms for large matrices via random projections. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 143–152, Oct 2006. doi: 10.1109/FOCS.2006.37.
- Joel A. Tropp. Greed is good: algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, 2004. doi: 10.1109/TIT.2004.834793.
- Cun-Hui Zhang and Tong Zhang. A general theory of concave regularization for high-dimensional sparse estimation problems. *Statist. Sci.*, 27(4):576–593, 11 2012.
- Lijun Zhang, Tianbao Yang, Rong Jin, and Zhi-Hua Zhou. Sparse learning for large-scale and high-dimensional data: A randomized convex-concave optimization approach. In *Algorithmic Learning Theory - 27th International Conference, ALT 2016, Bari, Italy, October 19-21, 2016, Proceedings*, pages 83–97, 2016a. doi: 10.1007/978-3-319-46379-7_6. URL https://doi.org/10.1007/978-3-319-46379-7_6.
- Weizhong Zhang, Lijun Zhang, Rong Jin, Deng Cai, and Xiaofei He. Accelerated sparse linear regression via random projection. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2337–2343, 2016b.

5. Proofs

5.1 Proof of Lemma 1

Proof We prove this Lemma by mathematical induction.

When $t = 1$, we first show that $\text{supp}(\mathbf{z}^{(1)}) \subseteq \text{supp}(\mathbf{z}^{(0)})$, i.e. the support of \mathbf{z} shrinks after the first iteration. To see this, $\tilde{\mathbf{z}}^{(t)} = \mathbf{z}^{(t-1)} - \frac{2}{\tau s}(\mathbf{D}^\top \mathbf{D} \mathbf{z}^{(t-1)} - \mathbf{D}^\top \mathbf{x})$.

Since $\|\mathbf{x} - \mathbf{D} \mathbf{z}^{(t-1)}\|_2^2 \leq 1$, let $\mathbf{g}^{(t-1)} = -\frac{2}{\tau s}(\mathbf{D}^\top \mathbf{D} \mathbf{z}^{(t-1)} - \mathbf{D}^\top \mathbf{x})$, then

$$|\tilde{z}_j^{(t)}| \leq \|\mathbf{g}^{(t-1)}\|_\infty \leq \frac{2}{\tau s} \|\mathbf{D}^\top (\mathbf{D} \mathbf{z}^{(t-1)} - \mathbf{x})\|_\infty \leq \frac{2}{\tau s}$$

where j is the index for any zero element of $\mathbf{z}^{(t-1)}$, $1 \leq j \leq n, j \notin \text{supp}(\mathbf{z}^{(t-1)})$. Now $|\tilde{\mathbf{z}}_j^{(t)}| < \sqrt{\frac{2\lambda}{\tau s}}$, and it follows that $\mathbf{z}_j^{(t)} = 0$ due to the update rule (6). Therefore, the zero elements of $\mathbf{z}^{(t-1)}$ remain unchanged in $\mathbf{z}^{(t)}$, and $\text{supp}(\mathbf{z}^{(t)}) \subseteq \text{supp}(\mathbf{z}^{(t-1)})$ for $t = 1$.

Let $Q_{\mathbf{S}}(\mathbf{y}) = \|\mathbf{x} - \mathbf{D}_{\mathbf{S}}\mathbf{y}\|_2^2$ for $\mathbf{y} \in \mathbb{R}^{|\mathbf{S}|}$, then we show that $s > 2|\mathbf{S}|$ is the Lipschitz constant for the gradient of function $Q_{\mathbf{S}}$. To see this, we have

$$\sigma_{\max}(\mathbf{D}_{\mathbf{S}}^{\top} \mathbf{D}_{\mathbf{S}}) = (\sigma_{\max}(\mathbf{D}_{\mathbf{S}}))^2 \leq \text{Tr}(\mathbf{D}_{\mathbf{S}}^{\top} \mathbf{D}_{\mathbf{S}}) = |\mathbf{S}|$$

Also, $\nabla Q_{\mathbf{S}}(\mathbf{y}) = 2(\mathbf{D}_{\mathbf{S}}^{\top} \mathbf{D}_{\mathbf{S}}\mathbf{y} - \mathbf{D}_{\mathbf{S}}^{\top} \mathbf{x})$, and

$$\begin{aligned} \|\nabla Q_{\mathbf{S}}(\mathbf{y}) - \nabla Q_{\mathbf{S}}(\mathbf{z})\|_2 &= 2\|\mathbf{D}_{\mathbf{S}}^{\top} \mathbf{D}_{\mathbf{S}}(\mathbf{y} - \mathbf{z})\|_2 \\ &\leq 2\sigma_{\max}(\mathbf{D}_{\mathbf{S}}^{\top} \mathbf{D}_{\mathbf{S}}) \cdot \|\mathbf{y} - \mathbf{z}\|_2 \\ &\leq 2|\mathbf{S}|\|\mathbf{y} - \mathbf{z}\|_2 < s\|\mathbf{y} - \mathbf{z}\|_2 \end{aligned} \quad (28)$$

Note that when $t = 1$, since

$$\mathbf{z}^{(t)} = \arg \min_{\mathbf{v} \in \mathbb{R}^n} \frac{\tau s}{2} \|\mathbf{v} - \tilde{\mathbf{z}}^{(t)}\|_2^2 + \lambda \|\mathbf{v}\|_0$$

we have

$$\begin{aligned} &\frac{\tau s}{2} \|\mathbf{z}^{(t)} - \tilde{\mathbf{z}}^{(t)}\|_2^2 + \lambda \|\mathbf{z}^{(t)}\|_0 \\ &\leq \frac{\tau s}{2} \left\| \frac{\nabla Q(\mathbf{z}^{(t-1)})}{\tau s} \right\|_2^2 + \lambda \|\mathbf{z}^{(t-1)}\|_0 \end{aligned} \quad (29)$$

which is equivalent to

$$\begin{aligned} &\langle \nabla Q_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}}^{(t-1)}), \mathbf{z}_{\mathbf{S}}^{(t)} - \mathbf{z}_{\mathbf{S}}^{(t-1)} \rangle + \frac{\tau s}{2} \|\mathbf{z}^{(t)} - \mathbf{z}^{(t-1)}\|_2^2 \\ &+ \lambda \|\mathbf{z}^{(t)}\|_0 \leq \lambda \|\mathbf{z}^{(t-1)}\|_0 \end{aligned} \quad (30)$$

due to the fact that

$$\langle \nabla Q(\mathbf{z}^{(t-1)}), \mathbf{z}^{(t)} - \mathbf{z}^{(t-1)} \rangle = \langle \nabla Q_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}}^{(t-1)}), \mathbf{z}_{\mathbf{S}}^{(t)} - \mathbf{z}_{\mathbf{S}}^{(t-1)} \rangle$$

Also, since s is the Lipschitz constant for $\nabla Q_{\mathbf{S}}$,

$$\begin{aligned} Q_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}}^{(t)}) &\leq Q_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}}^{(t-1)}) + \langle \nabla Q_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}}^{(t-1)}), \mathbf{z}_{\mathbf{S}}^{(t)} - \mathbf{z}_{\mathbf{S}}^{(t-1)} \rangle \\ &+ \frac{s}{2} \|\mathbf{z}_{\mathbf{S}}^{(t)} - \mathbf{z}_{\mathbf{S}}^{(t-1)}\|_2^2 \end{aligned} \quad (31)$$

Combining (30) and (31) and note that $\|\mathbf{z}_{\mathbf{S}}^{(t)} - \mathbf{z}_{\mathbf{S}}^{(t-1)}\|_2 = \|\mathbf{z}^{(t)} - \mathbf{z}^{(t-1)}\|_2$, $Q_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}}^{(t)}) = Q(\mathbf{z}^{(t)})$ and $Q_{\mathbf{S}}(\mathbf{z}_{\mathbf{S}}^{(t-1)}) = Q(\mathbf{z}^{(t-1)})$, we have

$$\begin{aligned} Q(\mathbf{z}^{(t)}) + \lambda \|\mathbf{z}^{(t)}\|_0 &\leq Q(\mathbf{z}^{(t-1)}) + \lambda \|\mathbf{z}^{(t-1)}\|_0 \\ &- \frac{(\tau - 1)s}{2} \|\mathbf{z}^{(t)} - \mathbf{z}^{(t-1)}\|_2^2 \end{aligned} \quad (32)$$

Now (7) and (8) are verified for $t = 1$. Suppose (7) and (8) hold for all $t \geq t_0$ with $t_0 \geq 1$. Since $\{L(\mathbf{z}^{(t)})\}_{t=1}^{t_0}$ is decreasing, we have

$$\begin{aligned} L(\mathbf{z}^{(t_0)}) &= \|\mathbf{x} - \mathbf{D}\mathbf{z}^{(t_0)}\|_2^2 + \lambda \|\mathbf{z}^{(t_0)}\|_0 \\ &\leq \|\mathbf{x} - \mathbf{D}\mathbf{z}^{(0)}\|_2^2 + \lambda \|\mathbf{z}^{(0)}\|_0 \leq 1 + \lambda|\mathbf{S}| \end{aligned}$$

which indicates that $\|\mathbf{x} - \mathbf{D}\mathbf{z}^{(t_0)}\|_2 \leq \sqrt{1 + \lambda|\mathbf{S}|}$. When $t = t_0 + 1$,

$$\begin{aligned} |\tilde{\mathbf{z}}_j^{(t)}| &\leq \|\mathbf{g}^{(t-1)}\|_\infty \leq \frac{2}{\tau s} \|\mathbf{D}^\top(\mathbf{D}\mathbf{z}^{(t-1)} - \mathbf{x})\|_\infty \\ &\leq \frac{2}{\tau s} \sqrt{1 + \lambda|\mathbf{S}|} \end{aligned}$$

where j is the index for any zero element of $\mathbf{z}^{(t-1)}$, $1 \leq j \leq n, j \notin \text{supp}(\mathbf{z}^{(t-1)})$. Now $|\tilde{\mathbf{z}}_j^{(t)}| < \sqrt{\frac{2\lambda}{\tau s}}$, and it follows that $\mathbf{z}_j^{(t)} = 0$ due to the update rule in (6). Therefore, the zero elements of $\mathbf{z}^{(t-1)}$ remain unchanged in $\mathbf{z}_j^{(t)}$, and $\text{supp}(\mathbf{z}^{(t)}) \subseteq \text{supp}(\mathbf{z}^{(t-1)}) \subseteq \mathbf{S}$ for $t = t_0 + 1$. Moreover, similar to the case when $t = 1$, we can derive (30), (31) and (32), so that the support shrinkage (7) and decline of the objective (8) are verified for $t = t_0 + 1$. It follows that the claim of this lemma holds for all $t \geq 1$.

Since the sequence $\{L(\mathbf{z}^{(t)})\}_t$ is decreasing with lower bound 0, it must converge. \blacksquare

5.2 Proof of Lemma 2

Proof We first prove that the sequences $\{\mathbf{z}^{(t)}\}_t$ is bounded for any $1 \leq i \leq n$. In the proof of Lemma 1, it is proved that

$$\begin{aligned} L(\mathbf{z}^{(t)}) &= \|\mathbf{x} - \mathbf{D}\mathbf{z}^{(t)}\|_2^2 + \lambda\|\mathbf{z}^{(t)}\|_0 \\ &\leq \|\mathbf{x} - \mathbf{D}\mathbf{z}^{(0)}\|_2^2 + \lambda\|\mathbf{z}^{(0)}\|_0 \leq 1 + \lambda|\mathbf{S}| \end{aligned}$$

for $t \geq 1$. Therefore, $\|\mathbf{x} - \mathbf{D}\mathbf{z}^{(t)}\|_2 \leq \sqrt{1 + \lambda|\mathbf{S}|}$ and it follows that $\|\mathbf{D}\mathbf{z}^{(t)}\|_2^2 \leq (1 + \sqrt{1 + \lambda|\mathbf{S}|})^2$. Since $\text{supp}(\mathbf{z}^{(t)}) \subseteq \mathbf{S}$ for $t \geq 0$ due to Lemma 1,

$$\begin{aligned} (1 + \sqrt{1 + \lambda|\mathbf{S}|})^2 &\geq \|\mathbf{D}\mathbf{z}^{(t)}\|_2^2 = \|\mathbf{D}_\mathbf{S}\mathbf{z}_\mathbf{S}^{(t)}\|_2^2 \\ &\geq \sigma_{\min}(\mathbf{D}_\mathbf{S}^\top \mathbf{D}_\mathbf{S}) \|\mathbf{z}_\mathbf{S}^{(t)}\|_2^2 = \sigma_{\min}(\mathbf{D}_\mathbf{S}^\top \mathbf{D}_\mathbf{S}) \|\mathbf{z}^{(t)}\|_2^2 \end{aligned}$$

Since $\mathbf{D}_\mathbf{S}$ is nonsingular, we have $\sigma_{\min}(\mathbf{D}_\mathbf{S}^\top \mathbf{D}_\mathbf{S}) = (\sigma_{\min}(\mathbf{D}))^2$ and it follows that $\mathbf{z}^{(t)}$ is bounded: $\|\mathbf{z}^{(t)}\|_2^2 \leq \frac{(1 + \sqrt{1 + \lambda|\mathbf{S}|})^2}{(\sigma_{\min}(\mathbf{D}))^2}$. In addition, since ℓ^0 -norm function $\|\cdot\|_0$ is a semi-algebraic function, therefore, according to Theorem 1 in Bolte et al. (2014), $\{\mathbf{z}^{(t)}\}_t$ converges to a critical point of $L(\mathbf{z})$, denoted by $\hat{\mathbf{z}}$. \blacksquare

5.3 Proof of Lemma 3

Proof Let $\hat{\mathbf{v}} = 2\mathbf{D}^\top(\mathbf{D}\hat{\mathbf{z}} - \mathbf{x}) + \dot{\mathbf{R}}(\hat{\mathbf{z}}; b)$. For for $j \in \hat{\mathbf{S}}$, since $\hat{\mathbf{z}}$ is a critical point of $L(\mathbf{z}) = \|\mathbf{x} - \mathbf{D}\mathbf{z}\|_2^2 + \lambda\|\mathbf{z}\|_0$. then $\frac{\partial Q}{\partial \mathbf{z}_j}|_{\mathbf{z}=\hat{\mathbf{z}}} = 0$ because $\frac{\partial \|\mathbf{z}\|_0}{\partial \mathbf{z}_j}|_{\mathbf{z}=\hat{\mathbf{z}}} = 0$. Note that $\min_{j \in \hat{\mathbf{S}}} |\hat{\mathbf{z}}_j| > b$, so $\frac{\partial \mathbf{R}}{\partial \mathbf{z}_j}|_{\mathbf{z}=\hat{\mathbf{z}}} = 0$, and it follows that $\hat{\mathbf{v}}_j = 0$.

For $j \notin \hat{\mathbf{S}}$, since $\frac{dR}{dz_j}(\hat{\mathbf{z}}_j +; b) = \frac{\lambda}{b}$ and $\frac{dR}{dz_j}(\hat{\mathbf{z}}_j -; b) = -\frac{\lambda}{b}$, $\frac{\lambda}{b} > \max_{j \notin \hat{\mathbf{S}}} |\frac{\partial Q}{\partial \mathbf{z}_j}|_{\mathbf{z}=\hat{\mathbf{z}}}$, we can choose the j -th element of $\dot{\mathbf{R}}(\hat{\mathbf{z}}; b)$ such that $\hat{\mathbf{v}}_j = 0$. Therefore, $\|\hat{\mathbf{v}}\|_2 = 0$, and $\hat{\mathbf{z}}$ is a local solution to the problem (10).

Now we prove that \mathbf{z}^* is also a local solution to (10). Let $\mathbf{v}^* = 2\mathbf{D}^\top(\mathbf{D}\mathbf{z}^* - \mathbf{x}) + \dot{\mathbf{R}}(\mathbf{z}^*; b)$, and Q is defined as before. For $j \in \mathbf{S}^*$, since \mathbf{z}^* is the global optimal solution to problem

(1), we also have $\frac{\partial Q}{\partial \mathbf{z}_j}|_{\mathbf{z}=\mathbf{z}^*} = 0$. If it is not the case and $\frac{\partial Q}{\partial \mathbf{z}_j}|_{\mathbf{z}=\mathbf{z}^*} \neq 0$, then we can change \mathbf{z}_j by a small amount in the direction of the gradient $\frac{\partial Q}{\partial \mathbf{z}_j}$ at the point $\mathbf{z} = \mathbf{z}^*$ and still make $\mathbf{z}_j \neq 0$, leading to a smaller value of the objective $L(\mathbf{z})$.

Note that $\min_{j \in \mathbf{S}^*} |\mathbf{z}_j^*| > b$, so $\frac{\partial \mathbf{R}}{\partial \mathbf{z}_j}|_{\mathbf{z}=\hat{\mathbf{z}}} = 0$, and it follows that $\mathbf{v}_j^* = 0$.

For $j \notin \mathbf{S}^*$, since $\frac{\lambda}{b} > \max_{j \notin \hat{\mathbf{S}}} |\frac{\partial Q}{\partial \mathbf{z}_j}|_{\mathbf{z}=\mathbf{z}^*}|$, we can choose the j -th element of $\hat{\mathbf{R}}(\mathbf{z}^*; b)$ such that $\mathbf{v}_j^* = 0$. It follows that $\|\mathbf{v}^*\|_2 = 0$, and \mathbf{z}^* is also a local solution to the problem (10). \blacksquare

5.4 Proof of Theorem 1

Proof According to Lemma 3, both $\hat{\mathbf{z}}$ and \mathbf{z}^* are local solutions to problem (10). In the following text, let $\beta_{\mathbf{I}}$ indicates a vector whose elements are those of β with indices in \mathbf{I} . Let $\Delta = \mathbf{z}^* - \hat{\mathbf{z}}$, $\tilde{\Delta} = \hat{\mathbf{P}}(\mathbf{z}^*) - \hat{\mathbf{P}}(\hat{\mathbf{z}})$. By Lemma 3, we have

$$\|2\mathbf{D}^\top \mathbf{D} \Delta + \tilde{\Delta}\|_2 = 0$$

It follows that

$$2\Delta^\top \mathbf{D}^\top \mathbf{D} \Delta + \Delta^\top \tilde{\Delta} \leq \|\Delta\|_2 \|2\mathbf{D}^\top \mathbf{D} \Delta + \tilde{\Delta}\|_2 = 0$$

Also, by the proof of Lemma 3, for $k \in \hat{\mathbf{S}} \cap \mathbf{S}^*$, since $(\mathbf{D}^\top \mathbf{D} \Delta)_k = 0$ we have $\tilde{\Delta}_k = 0$. We now present another property on any nonconvex function P using the degree of nonconvexity in Definition 3: $\theta(t, \kappa) := \sup_s \{-\text{sgn}(s-t)(\dot{P}(s; b) - \dot{P}(t; b)) - \kappa|s-t|\}$ on the regularizer \mathbf{P} . For any $s, t \in \mathbb{R}$, we have

$$-\text{sgn}(s-t)(\dot{P}(s; b) - \dot{P}(t; b)) - \kappa|s-t| \leq \theta(t, \kappa)$$

by the definition of θ . It follows that

$$\begin{aligned} \theta(t, \kappa)|s-t| &\geq -(s-t)(\dot{P}(s; b) - \dot{P}(t; b)) - \kappa(s-t)^2 \\ &\quad - (s-t)(\dot{P}(s; b) - \dot{P}(t; b)) \leq \theta(t, \kappa)|s-t| + \kappa(s-t)^2 \end{aligned} \quad (33)$$

Applying (33) with $P = P_j$ for $j = 1, \dots, n$, we have

$$\begin{aligned} 2\Delta^\top \mathbf{D}^\top \mathbf{D} \Delta &\leq -\Delta^\top \tilde{\Delta} = -\Delta_{\mathbf{F}}^\top \tilde{\Delta}_{\mathbf{F}} - \Delta_{\hat{\mathbf{S}} \cap \mathbf{S}^*}^\top \tilde{\Delta}_{\hat{\mathbf{S}} \cap \mathbf{S}^*} \\ &\leq |\mathbf{z}_{\mathbf{F}}^* - \hat{\mathbf{z}}_{\mathbf{F}}|^\top \theta(\hat{\mathbf{z}}_{\mathbf{F}}, \kappa) + \kappa \|\mathbf{z}_{\mathbf{F}}^* - \hat{\mathbf{z}}_{\mathbf{F}}\|_2^2 + \|\Delta_{\hat{\mathbf{S}} \cap \mathbf{S}^*}\|_2 \|\tilde{\Delta}_{\hat{\mathbf{S}} \cap \mathbf{S}^*}\|_2 \\ &\leq \|\theta(\hat{\mathbf{z}}_{\mathbf{F}}, \kappa)\|_2 \|\mathbf{z}_{\mathbf{F}}^* - \hat{\mathbf{z}}_{\mathbf{F}}\|_2 + \kappa \|\mathbf{z}_{\mathbf{F}}^* - \hat{\mathbf{z}}_{\mathbf{F}}\|_2^2 + \|\Delta\|_2 \|\tilde{\Delta}_{\hat{\mathbf{S}} \cap \mathbf{S}^*}\|_2 \\ &\leq \|\theta(\hat{\mathbf{z}}_{\mathbf{F}}, \kappa)\|_2 \|\Delta\|_2 + \kappa \|\Delta\|_2^2 + \|\Delta\|_2 \|\tilde{\Delta}_{\hat{\mathbf{S}} \cap \mathbf{S}^*}\|_2 \end{aligned} \quad (34)$$

On the other hand, $\Delta^\top \mathbf{D}^\top \mathbf{D} \Delta \geq \kappa_0^2 \|\Delta\|_2^2$. It follows from (34) that

$$2\kappa_0^2 \|\Delta\|_2^2 \leq \|\theta(\hat{\mathbf{z}}_{\mathbf{F}}, \kappa)\|_2 \|\Delta\|_2 + \kappa \|\Delta\|_2^2 + \|\Delta\|_2 \|\tilde{\Delta}_{\hat{\mathbf{S}} \cap \mathbf{S}^*}\|_2$$

When $\|\Delta\|_2 \neq 0$, we have

$$\begin{aligned} 2\kappa_0^2 \|\Delta\|_2 &\leq \|\theta(\hat{\mathbf{z}}_{\mathbf{F}}, \kappa)\|_2 + \kappa \|\Delta\|_2 + \|\tilde{\Delta}_{\hat{\mathbf{S}} \cap \mathbf{S}^*}\|_2 \\ \Rightarrow \|\Delta\|_2 &\leq \frac{\|\theta(\hat{\mathbf{z}}_{\mathbf{F}}, \kappa)\|_2 + \|\tilde{\Delta}_{\hat{\mathbf{S}} \cap \mathbf{S}^*}\|_2}{2\kappa_0^2 - \kappa} \end{aligned} \quad (35)$$

According to the definition of θ , it can be verified that $\theta(t, \kappa) = \max\{0, \frac{\lambda}{b} - \kappa|t - b|\}$ for $|t| > b$, and $\theta(0, \kappa) = \max\{0, \frac{\lambda}{b} - \kappa b\}$. Therefore,

$$\|\theta(\hat{\mathbf{z}}_{\mathbf{F}}, \kappa)\|_2 = \left(\sum_{j \in \mathbf{F} \cap \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\hat{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{F} \setminus \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}} \quad (36)$$

And it follows that

$$\|\Delta\|_2 \leq \frac{1}{2\kappa_0^2 - \kappa} \left(\left(\sum_{j \in \mathbf{F} \cap \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\hat{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{F} \setminus \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}} \right) \quad (37)$$

This proves the result of this theorem. \blacksquare

5.5 Proof of Theorem 2

Proof By the proof of Lemma 2, we have

$$\|2\tilde{\mathbf{D}}^\top \tilde{\mathbf{D}}\tilde{\mathbf{z}} - 2\tilde{\mathbf{D}}^\top \mathbf{x} + \dot{\mathbf{R}}(\tilde{\mathbf{z}})\|_2 = 0$$

It follows that

$$\begin{aligned} & \|2\mathbf{D}^\top \mathbf{D}\tilde{\mathbf{z}} - 2\mathbf{D}^\top \mathbf{x} + \dot{\mathbf{R}}(\tilde{\mathbf{z}})\|_2 \\ &= \|2\mathbf{D}^\top \mathbf{D}\tilde{\mathbf{z}} - 2\tilde{\mathbf{D}}^\top \tilde{\mathbf{D}}\tilde{\mathbf{z}} + 2\tilde{\mathbf{D}}^\top \tilde{\mathbf{D}}\tilde{\mathbf{z}} - 2\mathbf{D}^\top \mathbf{x} + 2\tilde{\mathbf{D}}^\top \mathbf{x} - 2\tilde{\mathbf{D}}^\top \mathbf{x} + \dot{\mathbf{R}}(\tilde{\mathbf{z}})\|_2 \\ &\leq \|2\mathbf{D}^\top \mathbf{D}\tilde{\mathbf{z}} - 2\tilde{\mathbf{D}}^\top \tilde{\mathbf{D}}\tilde{\mathbf{z}}\|_2 + \|2\mathbf{D}^\top \mathbf{x} - 2\tilde{\mathbf{D}}^\top \mathbf{x}\|_2 + \|2\tilde{\mathbf{D}}^\top \tilde{\mathbf{D}}\tilde{\mathbf{z}} - 2\tilde{\mathbf{D}}^\top \mathbf{x} + \dot{\mathbf{R}}(\tilde{\mathbf{z}})\|_2 \\ &= \|2\mathbf{D}^\top \mathbf{D}\tilde{\mathbf{z}} - 2\tilde{\mathbf{D}}^\top \tilde{\mathbf{D}}\tilde{\mathbf{z}}\|_2 + \|2\mathbf{D}^\top \mathbf{x} - 2\tilde{\mathbf{D}}^\top \mathbf{x}\|_2 \\ &\leq 2\|\mathbf{D}^\top (\mathbf{D} - \tilde{\mathbf{D}})\tilde{\mathbf{z}}\|_2 + 2\|(\mathbf{D} - \tilde{\mathbf{D}})^\top \tilde{\mathbf{D}}\tilde{\mathbf{z}}\|_2 + 2\|\mathbf{D}^\top \mathbf{x} - \tilde{\mathbf{D}}^\top \mathbf{x}\|_2 \end{aligned} \quad (38)$$

By $\tilde{L}(\tilde{\mathbf{z}}) \leq \tilde{L}(\mathbf{z}^{(0)})$, we have $\|\tilde{\mathbf{z}}\|_2 \leq M_0$. By Lemma 4, with probability at least $1 - 6e^{-p}$, $\|\mathbf{D} - \tilde{\mathbf{D}}\|_2 \leq C_{k, k_0}$. It follows from (41) that

$$\begin{aligned} & \|2\mathbf{D}^\top \mathbf{D}\tilde{\mathbf{z}} - 2\mathbf{D}^\top \mathbf{x} + \dot{\mathbf{R}}(\tilde{\mathbf{z}})\|_2 \\ &\leq 2\sigma_{\max}(\mathbf{D})C_{k, k_0}M_0 + 2C_{k, k_0}(\sigma_{\max}(\mathbf{D}) + C_{k, k_0})M_0 + 2C_{k, k_0}\|\mathbf{x}\|_2 \\ &= 2C_{k, k_0}M_0(2\sigma_{\max}(\mathbf{D}) + C_{k, k_0}) + 2C_{k, k_0}\|\mathbf{x}\|_2 \end{aligned}$$

Also, by the proof of Lemma 2,

$$\|2\mathbf{D}^\top \mathbf{D}\mathbf{z}^* - 2\mathbf{D}^\top \mathbf{x} + \dot{\mathbf{R}}(\mathbf{z}^*)\|_2 = 0$$

Let $\Delta = \mathbf{z}^* - \tilde{\mathbf{z}}$, $\tilde{\Delta} = \dot{\mathbf{R}}(\mathbf{z}^*) - \dot{\mathbf{R}}(\tilde{\mathbf{z}})$,

$$\|2\mathbf{D}^\top \mathbf{D}\Delta + \tilde{\Delta}\|_2 \leq 2C_{k, k_0}M_0(2\sigma_{\max}(\mathbf{D}) + C_{k, k_0}) + 2C_{k, k_0}\|\mathbf{x}\|_2$$

Now following the proof of Theorem 1, we have

$$\begin{aligned} & \|\mathbf{z}^* - \tilde{\mathbf{z}}\|_2 = \|\Delta\|_2 \\ &\leq \frac{1}{2\tau_0^2 - \tau} \left(\left(\sum_{j \in \mathbf{G} \cap \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\tilde{\mathbf{z}}_j - b|\})^2 + \sum_{j \in \mathbf{G} \setminus \hat{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}} + 2C_{k, k_0}M_0(2\sigma_{\max}(\mathbf{D}) + C_{k, k_0}) + 2C_{k, k_0}\|\mathbf{x}\|_2 \right) \end{aligned} \quad (39)$$

■

5.6 Proof of Theorem 4

We have the following lemma before proving Theorem 4.

Lemma 5 *Suppose \mathbf{T} satisfies the ℓ^2 -norm preserving property in Definition 5. If $m \geq 4c \log \frac{4}{\delta}$, then for any matrix $\mathbf{A} \in \mathbb{R}^{p \times d}$, $\mathbf{B} \in \mathbb{R}^{d \times q}$, with probability at least $1 - \delta$,*

$$\|\mathbf{A}\mathbf{T}^\top\mathbf{T}\mathbf{B} - \mathbf{A}\mathbf{B}\| \leq \|\mathbf{A}\|_F \|\mathbf{B}\|_F \sqrt{\frac{c}{m} \log \frac{4}{\delta}} \quad (40)$$

Lemma 5 can be proved using the definition of the ℓ^2 -norm preserving property in the same way that Lemma 6 in (Sarlos, 2006) or (Zhang et al., 2016a) is proved.

Proof [Proof of Theorem 4] By the proof of Lemma 2, we have

$$\|2\bar{\mathbf{D}}^\top \bar{\mathbf{D}}\bar{\mathbf{z}} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{x}} + \dot{\mathbf{R}}(\bar{\mathbf{z}})\|_2 = 0$$

It follows that

$$\begin{aligned} & \|2\mathbf{D}^\top \mathbf{D}\bar{\mathbf{z}} - 2\mathbf{D}^\top \mathbf{x} + \dot{\mathbf{R}}(\bar{\mathbf{z}})\|_2 \\ &= \|2\mathbf{D}^\top \mathbf{D}\bar{\mathbf{z}} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{D}}\bar{\mathbf{z}} + 2\bar{\mathbf{D}}^\top \bar{\mathbf{D}}\bar{\mathbf{z}} - 2\mathbf{D}^\top \mathbf{x} + 2\bar{\mathbf{D}}^\top \bar{\mathbf{x}} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{x}} + \dot{\mathbf{R}}(\bar{\mathbf{z}})\|_2 \\ &\leq \|2\mathbf{D}^\top \mathbf{D}\bar{\mathbf{z}} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{D}}\bar{\mathbf{z}}\|_2 + \|2\mathbf{D}^\top \mathbf{x} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{x}}\|_2 + \|2\bar{\mathbf{D}}^\top \bar{\mathbf{D}}\bar{\mathbf{z}} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{x}} + \dot{\mathbf{R}}(\bar{\mathbf{z}})\|_2 \\ &= \|2\mathbf{D}^\top \mathbf{D}\bar{\mathbf{z}} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{D}}\bar{\mathbf{z}}\|_2 + \|2\mathbf{D}^\top \mathbf{x} - 2\bar{\mathbf{D}}^\top \bar{\mathbf{x}}\|_2 \\ &= 2\|\mathbf{D}^\top (\mathbf{I} - \mathbf{T}^\top \mathbf{T})\mathbf{D}\bar{\mathbf{z}}\|_2 + 2\|\mathbf{D}^\top (\mathbf{I} - \mathbf{T}^\top \mathbf{T})\mathbf{x}\|_2 \end{aligned} \quad (41)$$

By $\bar{L}(\bar{\mathbf{z}}) \leq \bar{L}(\mathbf{z}^{(0)})$, we have $\|\bar{\mathbf{z}}\|_2 \leq M_1$. According to Lemma 5, with probability at least $1 - \delta$,

$$\begin{aligned} & 2\|\mathbf{D}^\top (\mathbf{I} - \mathbf{T}^\top \mathbf{T})\mathbf{D}\bar{\mathbf{z}}\|_2 + 2\|\mathbf{D}^\top (\mathbf{I} - \mathbf{T}^\top \mathbf{T})\mathbf{x}\|_2 \\ &\leq 2\|\mathbf{D}\|_F \sigma_{\max}(\mathbf{D}) M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} + 2\|\mathbf{D}\|_F M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} \\ &\leq 2\|\mathbf{D}\|_F M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} (\sigma_{\max}(\mathbf{D}) + 1) \end{aligned} \quad (42)$$

Combining (41) and (42), we have

$$\|2\mathbf{D}^\top \mathbf{D}\bar{\mathbf{z}} - 2\mathbf{D}^\top \mathbf{x} + \dot{\mathbf{R}}(\bar{\mathbf{z}})\|_2 \leq 2\|\mathbf{D}\|_F M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} (\sigma_{\max}(\mathbf{D}) + 1)$$

Also, by the proof of Lemma 2,

$$\|2\mathbf{D}^\top \mathbf{D}\mathbf{z}^* - 2\mathbf{D}^\top \mathbf{x} + \dot{\mathbf{R}}(\mathbf{z}^*)\|_2 = 0$$

Let $\Delta = \mathbf{z}^* - \bar{\mathbf{z}}$, $\bar{\Delta} = \dot{\mathbf{R}}(\mathbf{z}^*) - \dot{\mathbf{R}}(\bar{\mathbf{z}})$,

$$\|2\mathbf{D}^\top \mathbf{D}\Delta + \bar{\Delta}\|_2 \leq 2\|\mathbf{D}\|_F M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} (\sigma_{\max}(\mathbf{D}) + 1)$$

Now following the proof of Theorem 1, with probability at least $1 - \delta$,

$$\begin{aligned}
 \|\mathbf{z}^* - \bar{\mathbf{z}}\|_2 &= \|\Delta\|_2 \\
 &\leq \frac{1}{2\eta_0^2 - \eta} \left(\left(\sum_{j \in \mathbf{H} \cap \bar{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa|\bar{z}_j - b|\})^2 + \right. \right. \\
 &\quad \left. \left. \sum_{j \in \mathbf{H} \setminus \bar{\mathbf{S}}} (\max\{0, \frac{\lambda}{b} - \kappa b\})^2 \right)^{\frac{1}{2}} + 2\|\mathbf{D}\|_F M_1 \sqrt{\frac{c}{m} \log \frac{4}{\delta}} (\sigma_{\max}(\mathbf{D}) + 1) \right) \tag{43}
 \end{aligned}$$

■