

7-2-2014

Linear Convergence of Stochastic Iterative Greedy Algorithms with Sparse Constraints

Nam Nguyen

Deanna Needell
Claremont McKenna College

Tina Woolf

Recommended Citation

Nguyen, N., Needell, D., Woolf, T., "Linear Convergence of Stochastic Iterative Greedy Algorithms with Sparse Constraints", Submitted, arXiv preprint arXiv:1407.0088, 2014.

This Article - preprint is brought to you for free and open access by the CMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

Linear Convergence of Stochastic Iterative Greedy Algorithms with Sparse Constraints

Nam Nguyen, Deanna Needell and Tina Woolf

July 2, 2014

Abstract

Motivated by recent work on stochastic gradient descent methods, we develop two stochastic variants of greedy algorithms for possibly non-convex optimization problems with sparsity constraints. We prove linear convergence¹ in expectation to the solution within a specified tolerance. This generalized framework applies to problems such as sparse signal recovery in compressed sensing, low-rank matrix recovery, and covariance matrix estimation, giving methods with provable convergence guarantees that often outperform their deterministic counterparts. We also analyze the settings where gradients and projections can only be computed approximately, and prove the methods are robust to these approximations. We include many numerical experiments which align with the theoretical analysis and demonstrate these improvements in several different settings.

1 Introduction

Over the last decade, the problem of high-dimensional data inference from limited observations has received significant consideration, with many applications arising from signal processing, computer vision, and machine learning. In these problems, it is not unusual that the data often lies in hundreds of thousands or even million dimensional spaces while the number of collected samples is sufficiently smaller. Exploiting the fact that data arising in real world applications often has very low intrinsic complexity and dimensionality, such as sparsity and low-rank structure, recently developed statistical models have been shown to perform accurate estimation and inference. These models often require solving the following optimization with the constraint that the model parameter is sparse:

$$\min_w F(w) \quad \text{subject to} \quad \|w\|_0 \leq k. \quad (1)$$

Here, $F(w)$ is the objective function that measures the model discrepancy, $\|w\|_0$ is the ℓ_0 -norm that counts the number of non-zero elements of w , and k is a parameter that controls the sparsity of w .

In this paper, we study a more unified optimization that can be applied to a broader class of sparse models. First, we define a more general notion of sparsity. Given the set $\mathcal{D} = \{d_1, d_2, \dots\}$ consisting of vectors or matrices d_i , which we call atoms, we say that the model parameter is sparse if it can be described as a combination of only a few elements from the atomic set \mathcal{D} . Specifically, let $w \in \mathbb{R}^n$ be represented as

$$w = \sum_{i=1}^k \alpha_i d_i, \quad d_i \in \mathcal{D}, \quad (2)$$

¹Linear convergence is sometime called exponential convergence

where α_i are called coefficients of w ; then, w is called sparse with respect to \mathcal{D} if k is relatively small compared to the ambient dimension n . Here, \mathcal{D} could be a finite set (e.g. $\mathcal{D} = \{e_i\}_{i=1}^n$ where e_i 's are basic vectors in Euclidean space), or \mathcal{D} could be infinite (e.g. $\mathcal{D} = \{u_i v_i^*\}_{i=1}^\infty$ where $u_i v_i^*$'s are unit-norm rank-one matrices). This notion is general enough to handle many important sparse models such as group sparsity and low rankness (see [12], [33] for some examples).

Our focus in this paper is to develop algorithms for the following optimization:

$$\min_w \underbrace{\frac{1}{M} \sum_{i=1}^M f_i(w)}_{F(w)} \quad \text{subject to} \quad \|w\|_{0,\mathcal{D}} \leq k, \quad (3)$$

where $f_i(w)$'s, $w \in \mathbb{R}^n$, are smooth functions which can be *non-convex*; $\|w\|_{0,\mathcal{D}}$ is defined as the norm that captures the sparsity level of w . In particular, $\|w\|_{0,\mathcal{D}}$ is the smallest number of atoms in \mathcal{D} such that w can be represented by them:

$$\|w\|_{0,\mathcal{D}} = \min_k \{k : w = \sum_{i \in T} \alpha_i d_i \quad \text{with} \quad |T| = k\}. \quad (4)$$

Also in (3), k is a user-defined parameter that controls the sparsity of the model. The formulation (3) arises in many signal processing and machine learning problems, for instance, compressed sensing (e.g. [16], [8]), Lasso ([39]), sparse logistic regression, and sparse graphical model estimation (e.g. [41]). In the following, we provide some examples to demonstrate the generality of the optimization (3).

1) *Compressed sensing*: The goal is to recover a signal w^* from the set of observations $y_i = \langle a_i, w^* \rangle + \epsilon_i$ for $i = 1, \dots, m$. Assuming that the unknown signal w^* is sparse, we minimize the following to recover w^* :

$$\min_{w \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m (y_i - \langle a_i, w \rangle)^2 \quad \text{subject to} \quad \|w\|_0 \leq k.$$

In this problem, the set \mathcal{D} consists of n basic vectors, each of size n in Euclidean space. This problem can be seen as a special case of (3) with $f_i(w) = (y_i - \langle a_i, w \rangle)^2$ and $M = m$. An alternative way to write the above objective function is

$$\frac{1}{m} \sum_{i=1}^m (y_i - \langle a_i, w \rangle)^2 = \frac{1}{M} \sum_{j=1}^M \frac{1}{b} \left(\sum_{i=(j-1)b+1}^{jb} (y_i - \langle a_i, w \rangle)^2 \right),$$

where $M = m/b$. Thus, we can treat each function $f_j(w)$ as $f_j(w) = \frac{1}{b} \sum_{i=(j-1)b+1}^{jb} (y_i - \langle a_i, w \rangle)^2$. In this setting, each $f_j(w)$ accounts for a collection (or *block*) of observations of size b , rather than only one observation. This setting will be useful later for our proposed stochastic algorithms.

2) *Matrix recovery*: Given m observations $y_i = \langle A_i, W^* \rangle + \epsilon_i$ for $i = 1, \dots, m$ where the unknown matrix $W^* \in \mathbb{R}^{d_1 \times d_2}$ is assumed low-rank, we need to recover the original matrix W^* . To do so, we perform the following minimization:

$$\min_{W \in \mathbb{R}^{d_1 \times d_2}} \frac{1}{m} \sum_{i=1}^m (y_i - \langle A_i, W \rangle)^2 \quad \text{subject to} \quad \text{rank}(W) \leq k.$$

In this problem, the set \mathcal{D} consists of infinitely many unit-normed rank-one matrices and the functions $f_i(W) = (y_i - \langle A_i, W \rangle)^2$. We can also write functions in the block form $f_i(W) = \frac{1}{b} \sum (y_i - \langle A_i, W \rangle)^2$ as above.

3) *Covariance matrix estimation*: Let x be a Gaussian random vector of size n with covariance matrix W^* . The goal is to estimate W^* from m independent copies x_1, \dots, x_m of x . A useful way to find W^* is via solving the maximum log-likelihood function with respect to a sparse constraint on the precision matrix $\Sigma^* = (W^*)^{-1}$. The sparsity of Σ^* encourages the independence between entries of x . The minimization formula is as follows:

$$\min_{\Sigma} \frac{1}{m} \sum_{i=1}^m \langle x_i x_i^T, \Sigma \rangle - \log \det \Sigma \quad \text{subject to} \quad \|\Sigma_{\text{off}}\|_0 \leq k,$$

where Σ_{off} is the matrix Σ with diagonal elements set to zero. In this problem, \mathcal{D} is the finite collection of unit-normed $n \times n$ matrices $\{e_i e_j^*\}$ and the functions $f_i(\Sigma) = \langle x_i x_i^T, \Sigma \rangle - \frac{1}{m} \log \det \Sigma$.

Our paper is organized as follows. In the remainder of Section 1 we discuss related work in the literature and highlight our contributions; we also describe notations used throughout the paper and assumptions employed to analyze the algorithms. We present our stochastic algorithms in Sections 2 and 3, where we theoretically show the linear convergence rate of the algorithms and include a detailed discussion. In Section 4, we explore various extensions of the two proposed algorithms and also provide the theoretical result regarding the convergence rate. We apply our main theoretical results in Section 5, in the context of sparse linear regression and low-rank matrix recovery. In Section 6, we demonstrate several numerical simulations to validate the efficiency of the proposed methods and compare them with existing deterministic algorithms. Our conclusions are given in Section 7. We reserve Section 8 for our theoretical analysis.

1.1 Related work and our contribution

Sparse estimation has a long history and during its development there have been many great ideas along with efficient algorithms to solve (not exactly) the optimization problem (3). We sketch here some main lines which are by no means exhaustive.

Convex relaxation. Optimization based techniques arose as a natural convex relaxation to the problem of sparse recovery (3). There is now a massive amount of work in the field of Compressive Sensing and statistics [7, 13] that demonstrates these methods can accurately recover sparse signals from a small number of noisy linear measurements. Given noisy measurements $y = Aw^* + e$, one can solve the ℓ_1 -minimization problem

$$\hat{w} = \underset{w}{\operatorname{argmin}} \|w\|_1 \quad \text{such that} \quad \|Aw - y\|_2 \leq \varepsilon,$$

where ε is an upper bound on the noise $\|e\|_2 \leq \varepsilon$. Candès, Romberg and Tao [11, 9] prove that under a deterministic condition on the matrix A , this method accurately recovers the signal,

$$\|w^* - \hat{w}\|_2 \leq \varepsilon + \frac{\|w^* - w_k^*\|_2}{\sqrt{k}}, \quad (5)$$

where w_k^* denotes the k largest entries in magnitude of the signal w^* . The deterministic condition is called the Restricted Isometry Property (RIP) [11] and requires that the matrix A behave nicely on sparse vectors:

$$(1 - \delta)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta)\|x\|_2^2 \quad \text{for all } k\text{-sparse vectors } x,$$

for some small enough $\delta < 1$.

The convex approach is also extended beyond the quadratic objectives. In particular, the convex relaxation of the optimization (3) is the following:

$$\min_w F(w) + \lambda \|w\|, \quad (6)$$

where the regularization $\|w\|$ is used to promote sparsity, for instance, it can be the ℓ_1 norm (vector case) or the nuclear norm (matrix case). Many methods have been developed to solve these problems including interior point methods and other first-order iterative methods such as (proximal) gradient descent and coordinate gradient descent (e.g. [27, 18, 14]). The theoretical analyses of these algorithms have also been studied with either linear or sublinear rate of convergence, depending on the assumption imposed on the function $F(w)$. In particular, sublinear convergence rate is obtained if $F(w)$ exhibits a convex and smooth function, whereas the linear convergence rate is achieved when $F(w)$ is the smooth and strongly convex function. For problems such as compressed sensing, although the loss function $F(w)$ does not possess the strong convexity property, experiments still show the linear convergence behavior of the gradient descent method. In the recent work [1], the authors develop theory to explain this behavior. They prove that as long as the function $F(w)$ obeys the restricted strong convexity and restricted smoothness, a property similar to the RIP, then gradient descent algorithm can obtain the linear rate.

Greedy pursuits. More in line with our work are greedy approaches. These algorithms reconstruct the signal by identifying elements of the support iteratively. Once an accurate support set is located, a simple least-squares problem recovers the signal accurately. Greedy algorithms like Orthogonal Matching Pursuit (OMP) [40] and Regularized OMP (ROMP) [30] offer a much faster runtime than the convex relaxation approaches but lack comparable strong recovery guarantees. Recent work on greedy methods like Compressive Sampling Matching Pursuit (CoSaMP) and Iterative Hard Thresholding (IHT) offer both the advantage of a fast runtime and essentially the same recovery guarantees as (5) (e.g. [29, 5, 45, 19]). However, these algorithms are only applied for problems in compressed sensing where the least square loss is used to measure the discrepancy. There certainly exists many loss functions that are commonly used in statistical machine learning and do not exhibit quadratic structure such as log-likelihood loss. Therefore, it is necessary to develop efficient algorithms to solve (3).

There are several methods proposed to solve special instances of (3). [36] and [35] propose the forward selection method for sparse vector and low-rank matrix recovery. The method selects each nonzero entry or each rank-one matrix in an iterative fashion. [43] generalizes this algorithm to the more general dictionary \mathcal{D} . [44] proposes the forward-backward method in which an atom can be added or removed from the set, depending on how much it contributes to decrease the loss function. [25] extends this algorithms beyond the quadratic loss studied in [44]. [3] extends the CoSaMP algorithm for a more general loss function. Very recently, [33] further generalizes CoSaMP and proposes the Gradient Matching Pursuit (GradMP) algorithm to solve (3). This is perhaps the first greedy algorithm for (3) - the very general form of sparse recovery. They show that under a restricted convexity assumption of the objective function, the algorithm linearly converges to the optimal solution. This desirable property is also possessed by CoSaMP. We note that there are other algorithms having also been extended to the setting of sparsity in arbitrary \mathcal{D} but only limited to the quadratic loss setting, see e.g. [15, 20, 21, 22].

We outline the GradMP method here, since it will be used as motivation for the work we propose. GradMP [33] is a generalization of the CoSaMP [29] that solves a wider class of sparse reconstruction problems. Like OMP, these methods consist of four main steps: i) form a signal proxy, ii) select a set of large entries of the proxy, iii) use those as the support estimation and estimate the signal via least-squares, and iv) prune the estimation and repeat. Methods like OMP and CoSaMP use the proxy $A^*(y - Aw^t)$; more general methods like GradMP use the gradient $\nabla F(w^t)$ (see [33] for details). The analysis of GradMP depends on the restricted strong convexity and restricted strong smoothness properties as in Definitions 1 and 2 below, the first of which is motivated by a similar property introduced in [31]. Under these assumptions, the authors prove linear convergence to the noise floor.

The IHT, another algorithm that motivates our work, is a simple method that begins with an estimation $w^0 = 0$ and computes the next estimation using the recursion

$$w^{t+1} = H_k(w^t + A^*(y - Aw^t)),$$

where H_k is the thresholding operator that sets all but the largest (in magnitude) k coefficients of its argument to zero. Blumensath and Davies [5] prove that under the RIP, IHT provides a recovery bound comparable to (5). [24] extends IHT to the matrix recovery. Very recently, [42] proposes the Gradient Hard Thresholding Pursuit (GraHTP), an extension of IHT to solve a special vector case of (3).

Stochastic convex optimization. Methods for stochastic convex optimization have been developed in a very related but somewhat independent large body of work. We discuss only a few here which motivated our work, and refer the reader to e.g. [37, 6] for a more complete survey. Stochastic Gradient Descent (SGD) aims to minimize a convex objective function using unbiased stochastic gradient estimates, typically of the form $\nabla f_i(w)$ where i is chosen stochastically. For the optimization (3) with no constraint, this can be summarized concisely by the update rule

$$w^{t+1} = w^t - \alpha \nabla f_i(w^t),$$

for some step size α . For smooth objective functions $F(w)$, classical results demonstrate a $1/t$ convergence rate with respect to the objective difference $F(w^t) - F(w^*)$. In the strongly convex case, Bach and Moulines [2] improve this convergence to a linear rate, depending on the average squared condition number of the system. Recently, Needell et. al. draw on connections to the Kaczmarz method (see [26, 38] and references therein), and improve this to a linear dependence on the uniform condition number [28]. Another line of work is the Stochastic Coordinate Descent (SCD) beginning with the work of [32]. Extension to minimization of composite functions in (6) is described in [34].

Contribution. In this paper, we exploit ideas from IHT [5], CoSaMP [29] and GradMP [33] as well as the recent results in stochastic optimization [38, 28], and propose two new algorithms to solve (3). The IHT and CoSaMP algorithms have been remarkably popular in the signal processing community due to their simplicity and computational efficiency in recovering sparse signals from incomplete linear measurements. However, these algorithms are mostly used to solve problems in which the objective function is quadratic and it would be beneficial to extend the algorithmic ideas to the more general objective function.

We propose in this paper stochastic versions of the IHT and GradMP algorithms, which we term Stochastic IHT (StoIHT) and Stochastic GradMP (StoGradMP). These algorithms possess favorable properties toward large scale problems:

- The algorithms do not need to compute the full gradient of $F(w)$. Instead, at each iteration, they only sample one index $i \in [M] = \{1, 2, \dots, M\}$ and compute its associated gradient of $f_i(w)$. This property is particularly efficient in large scale settings in which the gradient computation is often prohibitively expensive.
- The algorithms do not need to perform an optimal projection at each iteration as required by the IHT and CoSaMP. Approximated projection is generally sufficient to guarantee linear convergence while the algorithms enjoy significant computational improvement.
- Under the restricted strong convexity assumption of $F(w)$ and the restricted strong smoothness assumption of $f_i(w)$ (defined below), the two proposed algorithms are guaranteed to converge linearly to the optimal solution.

- The algorithms and proofs can be extended further to consider other variants such as inexact gradient computations and inexact estimation.

1.2 Notations and assumptions

Notation: For a set Ω , let $|\Omega|$ denote its cardinality and Ω^c denote its complement. We will write D as the matrix whose columns consist of elements of \mathcal{D} , and denote D_Ω as the submatrix obtained by extracting the columns of D corresponding to the indices in Ω . We denote by $\mathcal{R}(D_\Omega)$ the space spanned by columns of the matrix D_Ω . Also denote by $\mathcal{P}_\Omega w$ the orthogonal projection of w onto $\mathcal{R}(D_\Omega)$. Given a vector $w \in \mathbb{R}^n$ that can be decomposed as $w = \sum_{i \in \Omega} \alpha_i d_i$, we say that the support of w with respect to \mathcal{D} is Ω , denoted by $\text{supp}_{\mathcal{D}}(w) = \Omega$. We denote by $[M]$ the set $\{1, 2, \dots, M\}$. We also define \mathbb{E}_i as the expectation with respect to i where i is drawn randomly from the set $[M]$. For a matrix A , we use conventional notations: $\|A\|$ and $\|A\|_F$ are the spectral norm and Frobenius norms of the matrix A . For the linear operator $\mathcal{A} : W \in \mathbb{R}^{n_1 \times n_2} \rightarrow y \in \mathbb{R}^m$, $\mathcal{A}^* : y \in \mathbb{R}^m \rightarrow W \in \mathbb{R}^{n_1 \times n_2}$ is the transpose of \mathcal{A} .

Denote $F(w) \triangleq \frac{1}{M} \sum_{i=1}^M f_i(w)$ and let $p(1), \dots, p(M)$ be the probability distribution of an index i selected at random from the set $[M]$. Note that $\sum_{i=1}^M p(i) = 1$. Another important observation is that if we select an index i from the set $[M]$ with probability $p(i)$, then

$$\mathbb{E}_i \frac{1}{Mp(i)} f_i(w) = F(w) \quad \text{and} \quad \mathbb{E}_i \frac{1}{Mp(i)} \nabla f_i(w) = \nabla F(w), \quad (7)$$

where the expectation is with respect to the index i .

Define $\text{approx}_k(w, \eta)$ as the operator that constructs a set Γ of cardinality k such that

$$\|\mathcal{P}_\Gamma w - w\|_2 \leq \eta \|w - w_k\|_2, \quad (8)$$

where w_k is the best k -sparse approximation of w with respect to the dictionary \mathcal{D} , that is, $w_k = \text{argmin}_{y \in \mathcal{D}_\Gamma, |\Gamma| \leq k} \|w - y\|_2$. Put another way, denote

$$\Gamma^* = \text{argmin}_{|\Gamma| \leq k} \|w - \mathcal{P}_\Gamma w\|_2.$$

Then, we require that

$$\|w - \mathcal{P}_\Gamma w\|_2 \leq \eta \|w - \mathcal{P}_{\Gamma^*} w\|_2. \quad (9)$$

An immediate consequence is the following inequality:

$$\|w - \mathcal{P}_\Gamma w\|_2 \leq \eta \|w - \mathcal{P}_R w\|_2. \quad (10)$$

for any set of $|R| \leq k$ atoms of \mathcal{D} . This follows because $\|w - \mathcal{P}_{\Gamma^*} w\|_2 \leq \|w - \mathcal{P}_R w\|_2$. In addition, taking the square on both sides of the above inequality and manipulating yields

$$\|\mathcal{P}_R w\|_2^2 \leq \frac{1}{\eta^2} \|\mathcal{P}_\Gamma w\|_2^2 + \frac{\eta^2 - 1}{\eta^2} \|w\|_2^2 = \|\mathcal{P}_\Gamma w\|_2^2 + \frac{\eta^2 - 1}{\eta^2} \|\mathcal{P}_{\Gamma^c} w\|_2^2.$$

Taking the square root gives us an important inequality for our analysis later. For any set of $|R| \leq k$ atoms of \mathcal{D} ,

$$\|\mathcal{P}_R w\|_2 \leq \|\mathcal{P}_\Gamma w\|_2 + \sqrt{\frac{\eta^2 - 1}{\eta^2}} \|\mathcal{P}_{\Gamma^c} w\|_2. \quad (11)$$

Assumptions: Before describing the two algorithms in the next section, we provide assumptions for the functions $f_i(w)$ as well as $F(w)$. The first assumption requires that $F(w)$ is restricted strongly convex with respect to the set \mathcal{D} . Although we do not require $F(w)$ to be globally convex, it is necessary that $F(w)$ is convex in certain directions to guarantee the linear convergence of our proposed algorithms. The intuition is that our greedy algorithms only drive along certain directions and seek for the optimal solution. Thus, a global convexity assumption is not necessary.

Definition 1 (\mathcal{D} -restricted strong convexity (\mathcal{D} -RSC)). *The function $F(w)$ satisfies the \mathcal{D} -RSC if there exists a positive constant ρ_k^- such that*

$$F(w') - F(w) - \langle \nabla F(w), w' - w \rangle \geq \frac{\rho_k^-}{2} \|w' - w\|_2^2, \quad (12)$$

for all vectors w and w' of size n such that $|\text{supp}_{\mathcal{D}}(w) \cup \text{supp}_{\mathcal{D}}(w')| \leq k$.

We notice that the left-hand side of the above inequality relates to the Hessian matrix of $F(w)$ (provided $F(w)$ is smooth) and the assumption essentially implies the positive definiteness of the $k \times k$ Hessian submatrices. We emphasize that this assumption is much weaker than the strong convexity assumption imposed on the full n dimensional space where the latter assumption implies the positive definiteness of the full Hessian matrix. In fact, when $k = n$, $F(w)$ exhibits a strongly convex function with parameter ρ_k^- , and when $\rho_k^- = 0$, $F(w)$ is a convex function. We also highlight that the \mathcal{D} -RSC assumption is particularly relevant when studying statistical estimation problems in the high-dimensional setting. In this setting, the number of observations is often much less than the dimension of the model parameter and therefore, the Hessian matrix of the loss function $F(w)$ used to measure the data fidelity is highly ill-posed.

In addition, we require that $f_i(w)$ satisfies the so-called \mathcal{D} -restricted strong smoothness which is defined as follows:

Definition 2 (\mathcal{D} -restricted strong smoothness (\mathcal{D} -RSS)). *The function $f_i(w)$ satisfies the \mathcal{D} -RSS if there exists a positive constant $\rho_k^+(i)$ such that*

$$\|\nabla f_i(w') - \nabla f_i(w)\|_2 \leq \rho_k^+(i) \|w' - w\|_2 \quad (13)$$

for all vectors w and w' of size n such that $|\text{supp}_{\mathcal{D}}(w) \cup \text{supp}_{\mathcal{D}}(w')| \leq k$.

Variants of these two assumptions have been used to study the convergence of the projected gradient descent algorithm [1]. In fact, the names restricted strong convexity and restricted strong smoothness are adopted from [1].

In this paper, we assume that the functions $f_i(w)$ satisfy \mathcal{D} -RSS with constants $\rho_k^+(i)$ for all $i = 1, \dots, M$ and $F(w)$ satisfies \mathcal{D} -RSC with constant ρ_k^- . The following quantities will be used extensively throughout the paper:

$$\alpha_k \triangleq \max_i \frac{\rho_k^+(i)}{Mp(i)}, \quad \rho_k^+ \triangleq \max_i \rho_k^+(i), \quad \text{and} \quad \bar{\rho}_k^+ \triangleq \frac{1}{M} \sum_{i=1}^M \rho_k^+(i). \quad (14)$$

2 Stochastic Iterative Hard Thresholding (StoIHT)

In this section, we describe the Stochastic Iterative Hard Thresholding (StoIHT) algorithm to solve (3). The algorithm is provided in Algorithm 1. At each iteration, the algorithm performs the following standard steps:

- Select an index i from the set $[M]$ with probability $p(i)$.
- Compute the gradient associated with the index just selected and move the solution along the gradient direction.
- Project the solution onto the constraint space via the approx operator defined in (8).

Ideally, we would like to compute the exact projection onto the constraint space or equivalently the best k -sparse approximation of b^t with respect to \mathcal{D} . However, the exact projection is often hard to evaluate or is computationally expensive in many problems. Take an example of the large scale matrix recovery problem, where computing the best matrix approximation would require an intensive Singular Value Decomposition (SVD) which often costs $\mathcal{O}(kmn)$, where m and n are the matrix dimensions. On the other hand, recent linear algebraic advances allow computing an approximate SVD in only $\mathcal{O}(k^2 \max\{m, n\})$. Thus, approximate projections could have a significant computational gain in each iteration. Of course, the price paid for fast approximate projections is a slower convergence rate. In Theorem 1 we will show this trade-off.

Algorithm 1 StoIHT algorithm

input: $k, \gamma, \eta, p(i)$, and stopping criterion
initialize: w^0 and $t = 0$
repeat
 randomize: select an index i_t from $[M]$ with probability $p(i_t)$
 proxy: $b^t = w^t - \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^t)$
 identify: $\Gamma^t = \text{approx}_k(b^t, \eta)$
 estimate: $w^{t+1} = \mathcal{P}_{\Gamma^t}(b^t)$
 $t = t + 1$
until halting criterion *true*
output: $\hat{w} = w^t$

Denote w^* as a feasible solution of (3). Our main result provides the convergence rate of the StoIHT algorithm via characterizing the ℓ_2 -norm error of t -th iterate w^t with respect to w^* . We first define some quantities necessary for a precise statement of the theorem. First, we denote the *contraction coefficient*

$$\kappa \triangleq 2\sqrt{(1 - \gamma(2 - \gamma\alpha_{3k})\rho_{3k}^-)} + \sqrt{(\eta^2 - 1)(1 + \gamma^2\alpha_{3k}\bar{\rho}_{3k}^+ - 2\gamma\rho_{3k}^-)}, \quad (15)$$

where the quantities α_{3k} , $\bar{\rho}_{3k}^+$, ρ_{3k}^- and η are defined in (14), (12), and (8). As will become clear later, the contraction coefficient κ controls the algorithm's rate of convergence and is required to be less than unity. This κ is intuitively dependent on the characteristics of the objective function (via \mathcal{D} -RSC and \mathcal{D} -RSS constants ρ_{3k}^+ and ρ_{3k}^-), the user-defined step size, the probability distribution, and the approximation error. The price paid for allowing a larger approximation error η is a slower convergence rate, since κ will also become large; however, η should not be allowed too large since κ must still be less than one.

We also define the *tolerance parameter*

$$\sigma_{w^*} \triangleq \frac{\gamma}{\min_i Mp(i)} \left(2\mathbb{E}_i \max_{|\Omega| \leq 3k} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2 + \sqrt{\eta^2 - 1} \mathbb{E}_i \|\nabla f_i(w^*)\|_2 \right), \quad (16)$$

where i is an index selected from $[M]$ with probability $p(i)$. Of course when w^* minimizes all components f_i , we have $\sigma_{w^*} = 0$, and otherwise σ_{w^*} measures (a modified version) of the usual noise variance in stochastic optimization.

In terms of these two ingredients, we now state our first main result. The proof is deferred to Section 8.2.

Theorem 1. *Let w^* be a feasible solution of (3) and w^0 be the initial solution. At the $(t + 1)$ -th iteration of Algorithm 1, the expectation of the recovery error is bounded by*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \frac{\sigma_{w^*}}{(1 - \kappa)} \quad (17)$$

where σ_{w^*} is defined by (16), κ is defined by (15) and is assumed to be strictly less than unity, and expectation is taken over all choices of random variables i_0, \dots, i_t .

The theorem demonstrates a linear convergence for the StoIHT even though the full gradient computation is not available. This is a significant computational advantage in large-scale settings where computing the full gradient often requires performing matrix multiplications with matrix dimensions in the millions. In addition, a stochastic approach may also gain advantages from parallel implementation. We emphasize that the result of Theorem 1 holds for any feasible solution w^* and the error of the $(t + 1)$ -th iterate is mainly governed by the second term involving the gradient of $\{f_i(w^*)\}_{i=1, \dots, M}$. For certain optimization problems, we expect that the energy of these gradients associated with the global optimum is small. For statistical estimation problems, the gradient of the true model parameter often involves only the statistical noise, which is small. Thus, after a sufficient number of iterations, the error between w^{t+1} and the true statistical parameter is only controlled by the model noise.

The result is significantly simpler when the optimal projection is available at each iteration. That is, the algorithm is always able to find the set Γ^t such that w^{t+1} is the best k -sparse approximation of b^t . In this case, $\eta = 1$ and the contraction coefficient κ in (15) is simplified to

$$\kappa = 2\sqrt{(1 - \gamma(2 - \gamma\alpha_{3k})\rho_{3k}^-)},$$

with $\alpha_{3k} = \max_i \frac{\rho_{3k}^+(i)}{Mp(i)}$ and $\sigma_{w^*} = \frac{2\gamma}{\min_i Mp(i)} \mathbb{E}_i \max_{|\Omega| \leq 3k} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2$. In order for $\kappa < 1$, we need $\rho_{3k}^- \geq \frac{3}{4}\alpha_{3k} = \frac{3}{4} \max_i \frac{\rho_{3k}^+(i)}{Mp(i)}$ and

$$\gamma < \frac{1 + \sqrt{1 - \frac{3\alpha_{3k}}{4\rho_{3k}^-}}}{\alpha_{3k}}.$$

The following corollary provides an interesting particular choice of the parameters in which Theorem 1 is easier to access.

Corollary 1. *Suppose that $\rho_{3k}^- \geq \frac{3}{4}\rho_{3k}^+$. Select $\gamma = \frac{1}{\alpha_{3k}}$, $\eta = 1$ and the probability distribution $p(i) = \frac{1}{M}$ for all $i = 1, \dots, M$. Then using the quantities defined by (14),*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \frac{2\gamma}{(1 - \kappa) \min_i Mp(i)} \mathbb{E}_i \max_{|\Omega| \leq 3k} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2,$$

where $\kappa = 2\sqrt{1 - \frac{\rho_{3k}^-}{\rho_{3k}^+}}$.

When the exact projection is not available, we would want to see how big η is such that the StoIHT still allows linear convergence. It is clear from (15) that for a given step size γ , bigger η

leads to bigger κ , or slower convergence rate. It is required by the algorithm that $\kappa < 1$. Therefore, η^2 must at least satisfy

$$\eta^2 \leq 1 + \frac{1}{1 + \gamma^2 \alpha_{3k} \bar{\rho}_{3k}^+ - 2\gamma \rho_{3k}^-}. \quad (18)$$

As $\gamma = \frac{1}{\rho_{3k}^+}$ and $p(i) = \frac{1}{M}$, $i = 1, \dots, M$, the bound is simplified to $\eta^2 \leq 1 + \frac{1}{2(1 - \rho_{3k}^-)}$. This bound implies that the approximation error in (9) should be at most $(1 + \epsilon)$ away from the exact projection error where $\epsilon \in (0, 1)$.

In Algorithm 1, the projection tolerance η is fixed during the iterations. However, there is a flexibility in changing it every iteration. The advantage of this flexibility is that this parameter can be set small during the first few iterations where the convergence is slow and gradually increased for the later iterations. Denoting the projection tolerance at the j -th iteration by η^j , we define the *contraction coefficient* at the j -th iteration:

$$\kappa_j \triangleq 2\sqrt{(1 - \gamma(2 - \gamma\alpha_{3k})\rho_{3k}^-)} + \sqrt{((\eta^j)^2 - 1)(1 + \gamma^2 \alpha_{3k} \bar{\rho}_{3k}^+ - 2\gamma \rho_{3k}^-)}, \quad (19)$$

and the *tolerance parameter* $\sigma_{w^*} \triangleq \max_{j \in [t]} \sigma_{w^*}^j$ where

$$\sigma_{w^*}^j \triangleq \frac{\gamma}{\min_i M p(i)} \left(2\mathbb{E}_i \max_{|\Omega| \leq 3k} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2 + \sqrt{(\eta^j)^2 - 1} \max_i \mathbb{E}_i \|\nabla f_i(w^*)\|_2 \right). \quad (20)$$

The following corollary shows the convergence of the StoIHT algorithm in the case where the projection tolerance is allowed to vary at each iteration:

Corollary 2. *At the $(t + 1)$ -th iteration of Algorithm 1, the recovery error is bounded by*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \|w^0 - w^*\|_2 \prod_{j=0}^{t+1} \kappa_j + \sigma_{w^*} \sum_{i=0}^t \prod_{j=t-i}^t \kappa_j, \quad (21)$$

where κ_j is defined by (19), and $\sigma_{w^*} = \max_{j \in [t]} \sigma_{w^*}^j$ is defined via (20).

3 Stochastic Gradient Matching Pursuit (StoGradMP)

CoSaMP [29] has been a very popular algorithm to recover a sparse signal from its linear measurements. In [33], the authors generalize the idea of CoSaMP and provide the GradMP algorithm that solves a broader class of sparsity-constrained problems. In this paper, we develop a stochastic version of the GradMP, namely StoGradMP, in which at each iteration only the evaluation of the gradient of a function f_i is required. The StoGradMP algorithm is described in Algorithm 2 which consists of following steps at each iteration:

- Randomly select an index i with probability $p(i)$.
- Compute the gradient of $f_i(w)$ with associated index i .
- Choose the subspace of dimension at most $2k$ to which the gradient vector is closest, then merge with the estimated subspace from previous iteration.
- Solve a sub-optimization problem with the search restricted on this subspace.
- Find the subspace of dimension k which is closest to the solution just found. This is the estimated subspace which is hopefully close to the true subspace.

At a high level, StoGradMP can be interpreted as at each iteration, the algorithm looks for a subspace based on the previous estimate and then seeks a new solution via solving a low-dimensional sub-optimization problem. Due to the \mathcal{D} -RSC assumption, the sub-optimization is convex and thus it can be efficiently solved by many off-the-shelf algorithms. StoGradMP stops when a halting criterion is satisfied.

Algorithm 2 StoGradMP algorithm

input: $k, \eta_1, \eta_2, p(i)$, and stopping criterion
initialize: $w^0, \Lambda = 0$, and $t = 0$
repeat
 randomize: select an index i_t from $[M]$ with probability $p(i_t)$
 proxy: $r^t = \nabla f_{i_t}(w^t)$
 identify: $\Gamma = \text{approx}_{2k}(r^t, \eta_1)$
 merge: $\widehat{\Gamma} = \Gamma \cup \Lambda$
 estimate: $b^t = \text{argmin}_w F(w) \quad w \in \text{span}(D_{\widehat{\Gamma}})$
 prune: $\Lambda = \text{approx}_k(b^t, \eta_2)$
 update: $w^{t+1} = \mathcal{P}_{\Lambda}(b^t)$
 $t = t+1$
until halting criterion *true*
output: $\hat{w} = w^t$

Denote w^* as a feasible solution of the optimization (3). We will present our main result for the StoGradMP algorithm. As before, our result controls the convergence rate of the recovery error at each iteration. We define the *contraction coefficient*

$$\kappa \triangleq (1 + \eta_2) \sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} \left(\max_i \sqrt{Mp(i)} \sqrt{\frac{\frac{2\eta_1^2 - 1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-}{\rho_{4k}^-}} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \right), \quad (22)$$

where the quantities $\alpha_{4k}, \rho_{4k}^+, \rho_{4k}^-, \eta_1$, and η_2 are defined in (14), (12), and (8). As will be provided in the following theorem, κ characterizes the convergence rate of the algorithm. This quantity depends on many parameters that play a role in the algorithm.

In addition, we define analogously as before the *tolerance parameter*

$$\sigma_{w^*} \triangleq C(1 + \eta_2) \frac{1}{\min_{i \in [M]} Mp(i)} \max_{|\Omega| \leq 4k, i \in [M]} \|\mathcal{P}_{\Omega} \nabla f_i(w^*)\|_2, \quad (23)$$

where C is defined as $C \triangleq \frac{1}{\rho_{4k}^-} \left(2 \max_{i \in [M]} Mp(i) \sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} + 3 \right)$.

We are now ready to state our result for the StoGradMP algorithm. The error bound has the same structure as that of StoIHT but with a different convergence rate.

Theorem 2. *Let w^* be a feasible solution of (3) and w^0 be the initial solution. At the $(t + 1)$ -th iteration of Algorithm 2, the recovery error is bounded by*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \frac{\sigma_{w^*}}{1 - \kappa} \quad (24)$$

where σ_{w^*} is defined by (23), κ is defined by (22) and is assumed to be strictly less than unity, and expectation is taken over all choices of random variables i_0, \dots, i_t .

When $p(i) = \frac{1}{M}$, $i = 1, \dots, M$, and $\eta_1 = \eta_2 = 1$ (exact projections are obtained), the contraction coefficient κ has a very simple representation: $\kappa = 2\sqrt{\frac{\rho_{4k}^+}{\rho_{4k}} \left(\frac{\rho_{4k}^+}{\rho_{4k}} - 1 \right)}$. This expression of κ is the same as that of the GradMP. In this situation, the requirement $\kappa < 1$ leads to the condition $\rho_{4k}^+ < \frac{2+\sqrt{6}}{4}\rho_{4k}^-$. The following corollary provides the explicit form of the recovery error.

Corollary 3. *Using the parameters described by (14), suppose that $\rho_{4k}^- > \frac{4}{2+\sqrt{6}}\rho_{4k}^+$. Select $\eta_1 = \eta_2 = 1$, and the probability distribution $p(i) = \frac{1}{M}$, $i = 1, \dots, M$. Then,*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \left(2\sqrt{\frac{\rho_{4k}^+(\rho_{4k}^+ - \rho_{4k}^-)}{(\rho_{4k}^-)^2}} \right)^{t+1} \|w^0 - w^*\|_2 + \sigma_{w^*},$$

where $\sigma_{w^*} = \frac{2}{\rho_{4k}} \left(2\sqrt{\frac{\rho_{4k}^+}{\rho_{4k}}} + 3 \right) \max_{|\Omega| \leq 4k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2$.

Similar to the StoIHT, the theorem demonstrates the linear convergence of the StoGradMP to the feasible solution w^* . The expected recovery error naturally consists of two components: one relates to the convergence rate and the other concerns the tolerance factor. As long as the contraction coefficient is small (less than unity), the first component is negligible, whereas the second component can be very large depending on the feasible solution we measure. We expect that the gradients of f_i 's associated with the global optimum to be small, as shown true in many statistical estimation problems such as sparse linear estimation and low-rank matrix recovery, so that the StoGradMP converges linearly to the optimum. We note that the linear rate here is precisely consistent with the linear rate of the original CoSaMP algorithm applied to compressed sensing problems [29]. Furthermore, StoGradMP gains significant computation over CoSaMP and GradMP since the full gradient evaluation is not required at each iteration.

In Algorithm 2, the parameters η_1 and η_2 are fixed during the iterations. However, they can be changed at each iteration. Denoting the projection tolerances at the j -th iteration by η_1^j and η_2^j , we define the *contraction coefficient* at the j -th iteration as

$$\kappa_j \triangleq (1 + \eta_2^j) \sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} \left(\max_i \sqrt{Mp(i)} \sqrt{\frac{2(\eta_1^j)^2 - 1}{(\eta_1^j)^2} \rho_{4k}^+ - \rho_{4k}^-} + \frac{\sqrt{(\eta_1^j)^2 - 1}}{\eta_1^j} \right). \quad (25)$$

Also define the *tolerance parameter* $\sigma_{w^*} \triangleq \max_{j \in [t]} \sigma_{w^*}^j$ where

$$\sigma_{w^*}^j \triangleq C(1 + \eta_2^j) \frac{1}{\min_{i \in [M]} Mp(i)} \max_{|\Omega| \leq 4k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2 \quad (26)$$

and C is defined as $C \triangleq 2 \max_{i \in [M]} Mp(i) \sqrt{\frac{\alpha_{4k}}{\rho_{4k}}} + 3$. The following corollary shows the convergence of the algorithm.

Corollary 4. *At the $(t + 1)$ -th iteration of Algorithm 2, the recovery error is bounded by*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \|w^0 - w^*\|_2 \prod_{j=0}^{t+1} \kappa_j + \sigma_{w^*} \sum_{i=0}^t \prod_{j=t-i}^t \kappa_j, \quad (27)$$

where κ_j is defined by (25), and $\sigma_{w^*} = \max_{j \in [t]} \sigma_{w^*}^j$ is defined via (26).

4 StoIHT and StoGradMP with inexact gradients

In this section, we investigate the StoIHT and StoGradMP algorithms in which the gradient might not be exactly estimated. This issue occurs in many practical problems such as distributed network optimization in which gradients are corrupted by noise during the communication on the network. In particular, in both algorithms, the gradient selected at each iteration is contaminated by a noise vector e^t where t indicates the iteration number. We assume $\{e^t\}_{t=1,2,\dots}$ are deterministic noise with bounded energies.

4.1 StoIHT with inexact gradients

In the StoIHT algorithm, the update b^t at the proxy step has to take into account the noise appearing in the gradient. In particular, at the t -th iteration,

$$b^t = w^t - \frac{\gamma}{Mp(i_t)} (\nabla f_{i_t}(w^t) + e^t).$$

Denote the quantity

$$\sigma_e \triangleq \frac{\gamma}{\min_i Mp(i)} \max_{j \in [t]} \left(2 \max_{|\Omega| \leq 3k} \|\mathcal{P}_\Omega e^j\|_2 + \sqrt{\eta^2 - 1} \|e^j\|_2 \right). \quad (28)$$

We state our result in the following theorem. The proof is deferred to Section 8.4.

Theorem 3. *Let w^* be a feasible solution of (3). At the $(t+1)$ -th iteration of Algorithm 1 with inexact gradients, the expectation of the recovery error is bounded by*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \frac{1}{(1 - \kappa)} (\sigma_{w^*} + \sigma_e), \quad (29)$$

where κ is defined in (15) and is assumed to be strictly less than unity and expectation is taken over all choices of random variables i_1, \dots, i_t . The quantities σ_{w^*} and σ_e are defined in (16) and (28), respectively.

Theorem 3 provides the linear convergence of StoIHT even in the setting of an inexact gradient computation. The error bound shares a similar structure as that of the StoIHT with only an additional term related to the gradient noise. An interesting property is that the noise does not accumulate over iterations. Rather, it only depends on the largest noise level.

4.2 StoGradMP with inexact gradients

In the StoGradMP algorithm, accounting for noise in the gradient appears in the proxy step; the expression of r^t , with an additional noise term, becomes

$$r^t = \nabla f_{i_t}(w^t) + e^t.$$

Denote the quantity

$$\sigma_e \triangleq \frac{\max_i p(i)}{\rho_{4k}^- \min_i p(i)} \max_{j \in [t]} \|e^j\|_2. \quad (30)$$

We have the following theorem.

Theorem 4. *Let w^* be a feasible solution of (3). At the $(t + 1)$ -th iteration of Algorithm 2 with inexact gradients, the expectation of the recovery error is bounded by*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \frac{1}{(1 - \kappa)}(\sigma_{w^*} + \sigma_e), \quad (31)$$

where κ is defined in (22) and is assumed to be strictly less than unity and expectation is taken over all choices of random variables i_1, \dots, i_t . The quantities σ_{w^*} and σ_e are defined in (23) and (30), respectively.

Similar to the StoIHT, StoGradMP is stable under the contamination of gradient noise. Stability means that the algorithm is still able to obtain the linear convergence rate. The gradient noise only affects the tolerance rate and not the contraction factor. Furthermore, the recovery error only depends on the largest gradient noise level, implying that the noise does not accumulate over iterations.

4.3 StoGradMP with inexact gradients and approximated estimation

In this section, we extend the theory of the StoGradMP algorithm further to consider the sub-optimality of optimization at the estimation step. Specifically, we assume that at each iteration, the algorithm only obtains an approximated solution of the sub-optimization. Denote

$$b_{\text{opt}}^t = \underset{w}{\operatorname{argmin}} F(w) \quad \text{subject to} \quad w \in \operatorname{span}(D_{\hat{\Gamma}}), \quad (32)$$

as the optimal solution of this convex optimization, where $\hat{\Gamma} = \Gamma \cup \Lambda$ may also give rise to an approximation at the identification step. Write b^t as the approximated solution available at the estimation step. Then b^t is linked to b_{opt}^t via the relationship: $\|b^t - b_{\text{opt}}^t\|_2 \leq \epsilon^t$. This consideration is realistic in two aspects: first, the optimization (32) can be too slow to converge to the optimal solution, hence we might want to stop the algorithm after a sufficient number of steps or whenever the solution is close to the optimum; second, even if (32) has a closed-form solution as the least-squares problem, it is still beneficial to solve it approximately in order to reduce the computational complexity caused by the pseudo-inverse process (see [17] for an example of randomized least-squares approximation). Denoting the quantity

$$\sigma_\epsilon = \max_{j \in [t]} \epsilon^j, \quad (33)$$

we have the following theorem.

Theorem 5. *Let w^* be a feasible solution of (3). At the $(t + 1)$ -th iteration of Algorithm 2 with inexact gradients and approximated estimations, the expectation of the recovery error is bounded by*

$$\mathbb{E} \|w^{t+1} - w^*\|_2 \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \frac{1}{(1 - \kappa)}(\sigma_{w^*} + \sigma_e + \sigma_\epsilon), \quad (34)$$

where κ is defined in (22) and is assumed to be strictly less than unity and expectation is taken over all choices of random variables i_1, \dots, i_t . The quantities σ_{w^*} , σ_e , and σ_ϵ are defined in (23), (30), and (33), respectively.

Theorem 5 shows the stability of StoGradMP under both the contamination of gradient noise at the proxy step and the approximate optimization at the estimation step. Furthermore, StoGradMP still achieves a linear convergence rate even in the presence of these two sources of noise. Similar to the artifacts of gradient noise, the approximated estimation affects the tolerance rate and not the contraction factor, and the recovery is only impacted by the largest approximated estimation bound (rather than an accumulation over all of the iterations).

5 Some estimates

In this section we investigate some specific problems which require solving an optimization with a sparse constraint and transfer results of Theorems 1 and 2.

5.1 Sparse linear regression

The first problem of interest is the well-studied sparse recovery in which the goal is to recover a k_0 -sparse vector w_0 from noisy observations of the following form:

$$y = Aw_0 + \xi.$$

Here, the $m \times n$ matrix A is called the design matrix and ξ is the m dimensional vector noise. A natural way to recover w_0 from the observation vector y is via solving

$$\min_{w \in \mathbb{R}^n} \frac{1}{2m} \|y - Aw\|_2^2 \quad \text{subject to} \quad \|w\|_0 \leq k, \quad (35)$$

where k is the user-defined parameter which is assumed greater than k_0 . Clearly, this optimization is a special form of (3) with \mathcal{D} being the collection of standard vectors in \mathbb{R}^n and $F(w) = \frac{1}{2m} \|y - Aw\|_2^2$. Decompose the vector y into non-overlapping vectors y_{b_i} of size b and denote A_{b_i} as the $b_i \times n$ submatrix of A . We can then rewrite $F(w)$ as

$$F(w) = \frac{1}{M} \sum_{i=1}^M \frac{1}{2b} \|y_{b_i} - A_{b_i}w\|_2^2 \triangleq \frac{1}{M} \sum_{i=1}^M f_i(w),$$

where $M = m/b$. In order to apply Theorems 1 and 2 for this problem, we need to compute the contraction coefficient and tolerance parameter which involve the \mathcal{D} -RSC and \mathcal{D} -RSS conditions. It is easy to see that these two properties of $F(w)$ and $\{f_i(w)\}_{i=1}^M$ are equivalent to the RIP studied in [9]. In particular, we require that the matrix A satisfies

$$\frac{1}{m} \|Aw\|_2^2 \geq (1 - \delta_k) \|w\|_2^2$$

for all k -sparse vectors w . In addition, the matrices A_{b_i} , $i = 1, \dots, M$, are also required to obey

$$\frac{1}{b} \|A_{b_i}w\|_2^2 \leq (1 + \delta_k) \|w\|_2^2$$

for all k -sparse vectors w . Here, $(1 + \delta_k)$ and $(1 - \delta_k)$ with $\delta_k \in (0, 1]$ play the role of $\rho_k^+(i)$ and ρ_k^- in Definitions 1 and 2, respectively. For the Gaussian matrix A (entries are i.i.d. $\mathcal{N}(0, 1)$), it is well-known that these two assumptions hold as long as $m \geq \frac{Ck \log n}{\delta_k}$ and $b \geq \frac{ck \log n}{\delta_k}$. By setting the block size $b = ck \log n$, the number of blocks M is thus proportional to $\frac{m}{k \log n}$.

Now using StoIHT to solve (35) and applying Theorem 1, we set the step size $\gamma = 1$, the approximation error $\eta = 1$, and $p(i) = 1/M$, $i = 1, \dots, M$, for simplicity. Thus, the quantities in (14) are all the same and equal to $1 + \delta_k$. It is easy to verify that the contraction coefficient defined in (15) is $\kappa = 2\sqrt{2\delta_{3k} - \delta_{3k}^2}$. One can obtain $\kappa \leq 3/4$ when $\delta_{3k} \leq 0.07$, for example. In addition, since w_0 is the feasible solution of (35), the tolerance parameter σ_{w_0} defined in (16) can be rewritten as

$$\sigma_{w_0} = 2\mathbb{E}_i \max_{|\Omega| \leq 3k} \frac{1}{b} \|\mathcal{P}_\Omega A_{b_i}^* \xi_{b_i}\|_2 \leq \frac{2}{b} \sqrt{3k} \max_{i \in [M]} \max_{j \in [n]} |\langle A_{b_i, j}, \xi_{b_i} \rangle|,$$

where $A_{b_i, j}$ is the j -th column of the matrix A_{b_i} . For stochastic noise $\xi \sim \mathcal{N}(0, \sigma^2 I_m)$, it is easy to verify that $\sigma_{w_0} \leq c' \sqrt{\frac{\sigma^2 k \log n}{b}}$ with probability at least $1 - n^{-1}$.

Using StoGradMP to solve (35) with the same setting as above, we write the contraction coefficient in (22) as $\kappa = 2\sqrt{\frac{2\delta_{4k}(1+\delta_{4k})}{(1-\delta_{4k})^2}}$, which is less than $3/4$ if $\delta_{4k} \leq 0.05$. The tolerance parameter σ_{w_0} in (23) can be simplified similarly as in StoIHT. We now provide the following corollary based on what we have discussed.

Corollary 5 (for StoIHT and StoGradMP). *Assume $A \in \mathbb{R}^{m \times n}$ satisfies the \mathcal{D} -RSC and \mathcal{D} -RSS assumptions and $\xi \sim \mathcal{N}(0, \sigma^2)$. Then with probability at least $1 - n^{-1}$, the error at the $(t + 1)$ -th iterate of the StoIHT and StoGradMP algorithms is bounded by*

$$\mathbb{E} \|w^{t+1} - w_0\|_2 \leq (3/4)^{t+1} \|w_0\|_2 + c\sqrt{\frac{\sigma^2 k_0 \log n}{b}}.$$

We describe the convergence result of the two algorithms in one corollary since their results share the same form with the only difference in the constant c . One can see that for a sufficient number of iterations, the first term involving $\|w_0\|_2$ is negligible and the recovery error only depends on the second term. When the noise is absent, both algorithms guarantee recovery of the exact w_0 . The recovery error also depends on the block size b . When b is small, more error is expected, and the error decreases as b increases. This of course matches our intuition. We emphasize that the deterministic IHT and GradMP algorithms deliver the same recovery error with b replaced by m .

5.2 Low-rank matrix recovery

We consider the high-dimensional matrix recovery problem in which the observation model has the form

$$y_j = \langle A_j, W_0 \rangle + \xi_j \quad , \quad j = 1, \dots, m,$$

where W_0 is the $n_1 \times n_2$ unknown rank- k_0 matrix, each measurement A_j is an $m \times n_1$ matrix, and the noise ξ_j is assumed $N(0, \sigma^2)$. Noting that $\frac{1}{2m} \sum_{j=1}^m (y_j - \langle A_j, W \rangle)^2 = \frac{1}{2m} \|y - \mathcal{A}(W)\|_2^2$, the standard approach to recover W_0 is to solve the minimization

$$\min_{W \in \mathbb{R}^{n_1 \times n_2}} \frac{1}{2m} \|y - \mathcal{A}(W)\|_2^2 \quad \text{subject to} \quad \text{rank}(W) \leq k, \quad (36)$$

with k assumed greater than k_0 . Here, \mathcal{A} is the linear operator. In this problem, the set \mathcal{D} consists of infinitely many unit-normed rank-one matrices and the objective function can be written as a summation of sub-functions:

$$F(W) = \frac{1}{M} \sum_{i=1}^M f_i(W) = \frac{1}{M} \sum_{i=1}^M \left(\frac{1}{2b} \sum_{j=(i-1)b+1}^{ib} (y_j - \langle A_j, W \rangle)^2 \right) \triangleq \frac{1}{M} \sum_{i=1}^M \frac{1}{2b} \|y_{b_i} - \mathcal{A}_i(W)\|_2^2,$$

where $m = Mb$ (assume b is integer). Each $f_i(W)$ accounts for a collection (or block) of observations y_{b_i} of size b . In this case, the \mathcal{D} -RSC and \mathcal{D} -RSS properties are equivalent to the matrix-RIP [10], which holds for a wide class of random operators \mathcal{A} . In particular, we require

$$\frac{1}{m} \|\mathcal{A}(W)\|_2^2 \geq (1 - \delta_k) \|W\|_F^2$$

for all rank- k matrices W . In addition, the linear operators \mathcal{A}_i are required to obey

$$\frac{1}{b} \|\mathcal{A}_i(W)\|_2^2 \leq (1 + \delta_k) \|W\|_F^2$$

for all rank- k matrices W . Here, $(1 + \delta_k)$ and $(1 - \delta_k)$ with $\delta_k \in (0, 1]$ play the role of $\rho_k^+(i)$ and ρ_k^- in Definitions 1 and 2, respectively. For the random Gaussian linear operator \mathcal{A} (vectors A_i are i.i.d. $\mathcal{N}(0, I)$), it is well-known that these two assumptions hold as long as $m \geq \frac{Ck(n_1+n_2)}{\delta_k}$ and $b \geq \frac{ck(n_1+n_2)}{\delta_k}$. By setting the block size $b = ck(n_1 + n_2)$, the number of blocks M is thus proportional to $\frac{m}{k(n_1+n_2)}$.

In this section, we consider applying results of Theorems 3 and 5. To do so, we need to compute the contraction coefficients and tolerance parameter. We begin with Theorem 3, which holds for the StoIHT algorithm. Similar to the previous section, κ in (15) can have a similar form. However, in the matrix recovery problem, SVD computations are required at each iteration which is often computationally expensive. There has been a vast amount of research focusing on approximation methods that perform nearly as good as exact SVD but with much faster computation. Among them are the randomized SVD [23] that we will employ in the experimental section. For simplicity, we set the step size $\gamma = 1$ and $p(i) = 1/M$ for all i . Thus, quantities in (14) are the same and equal to $1 + \delta_k$. Rewriting κ in (15), we have

$$\kappa = 2\sqrt{2\delta_{3k} - \delta_{3k}^2} + \sqrt{(\eta^2 - 1)(\delta_{3k}^2 + 4\delta_{3k})},$$

where we recall η is the projection error. Setting $\kappa \leq 3/4$ by allowing the first term to be less than $1/2$ and the second term less than $1/4$, we obtain $\delta_{4k} \leq 0.03$ and the approximation error η is allowed up to 1.19.

The next step is to evaluate the tolerance parameter σ_{W_0} in (16). The parameter σ_{W_0} can be read as

$$\begin{aligned} \sigma_{W_0} &= 2\mathbb{E}_i \max_{|\Omega| \leq 4k} \frac{1}{b} \|\mathcal{P}_\Omega \mathcal{A}_i^*(\xi_{b_i})\|_F + \sqrt{\eta^2 - 1} \mathbb{E}_i \frac{1}{b} \|\mathcal{A}_i^*(\xi_{b_i})\|_F \\ &\leq \frac{2}{b} \sqrt{4k} \max_i \|\mathcal{A}_i^*(\xi_{b_i})\| + \frac{1}{b} \sqrt{\eta^2 - 1} \sqrt{n} \max_i \|\mathcal{A}_i^*(\xi_{b_i})\|. \end{aligned}$$

For stochastic noise $\xi \sim \mathcal{N}(0, \sigma^2 I)$, it is shown in [10], Lemma 1.1 that $\|\mathcal{A}_i^*(\xi_{b_i})\| \leq c\sqrt{\sigma^2 n b}$ with probability at least $1 - n^{-1}$ where $n = \max\{n_1, n_2\}$. Therefore, $\sigma_{W_0} \leq c \left(\sqrt{\frac{\sigma^2 k n}{b}} + \sqrt{\frac{(\eta^2 - 1)\sigma^2 n^2}{b}} \right)$. In addition, the parameter σ_e in (28) is estimated as

$$\sigma_e \leq \max_j \left(2 \max_{|\Omega| \leq 3k} \|\mathcal{P}_\Omega E^j\|_F + \sqrt{\eta^2 - 1} \|E^j\|_F \right) \leq \max_j \left(2\sqrt{3k} \|E^j\| + \sqrt{\eta^2 - 1} \|E^j\|_F \right),$$

where we recall that E^j is the noise matrix that might contaminate the gradient at the j -th iteration. Applying Theorem 3 leads to the following corollary.

Corollary 6 (for StoIHT). *Assume the linear operator \mathcal{A} satisfies the \mathcal{D} -RSC and \mathcal{D} -RSS assumptions and $\xi \sim \mathcal{N}(0, \sigma^2)$. Set $p(i) = 1/M$ for $i = 1, \dots, M$ and $\gamma = 1$. Then with probability at least $1 - n^{-1}$, the error at the $(t + 1)$ -th iterate of the StoIHT algorithm is bounded by*

$$\begin{aligned} \mathbb{E} \|W^{t+1} - W_0\|_F &\leq (3/4)^{t+1} \|W_0\|_F + c \left(\sqrt{\frac{\sigma^2 k n}{b}} + \sqrt{\frac{(\eta^2 - 1)\sigma^2 n^2}{b}} \right) \\ &\quad + 4 \max_{j \in [t]} \left(2\sqrt{3k} \|E^j\| + \sqrt{\eta^2 - 1} \|E^j\|_F \right). \end{aligned}$$

The discussion after Corollary 5 for the vector recovery can also be applied here. For a sufficient number of iterations, the recovery error is naturally controlled by three factors: the measurement noise σ , the approximation projection parameter η , and the largest gradient noise E^j . In the absence of these three parameters, the recovery is exact. When $\eta = 1$ and $E^j = 0$, $j = 0, \dots, t$, the error has the same structure as the convex nuclear norm minimization method [10] which has been shown to be optimal.

Moving to the StoGradMP algorithm, setting $p(i) = 1/M$ for $i = 1, \dots, M$ again for simplicity, we can write the contraction coefficient in (22) as

$$\kappa = (1 + \eta_2) \sqrt{\frac{1 + \delta_{4k}}{1 - \delta_{4k}}} \left(\sqrt{\frac{\frac{2\eta_1^2 - 1}{\eta_1^2} (1 + \delta_{4k}) - (1 - \delta_{4k})}{1 - \delta_{4k}}} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \right).$$

If we allow for example the projection error $\eta_1 = 1.01$ and $\eta_2 = 1.01$ and require $\kappa \leq 0.9$, simple algebra gives us $\delta_{4k} \leq 0.03$. In addition, for stochastic noise $\xi \sim \mathcal{N}(0, \sigma^2 I)$, the tolerance parameter σ_{W_0} in (23) can be read as

$$\sigma_{W_0} = c(1 + \eta_2) \max_{i \in [M], |\Omega| \leq 4k} \|\mathcal{P}_\Omega \mathcal{A}_i^*(\xi_{b_i})\|_F \leq c_1(1 + \eta_2) \sqrt{4k} \max_{i \in [M]} \|\mathcal{A}_i^*(\xi_{b_i})\| \leq c_2(1 + \eta_2) \sqrt{\frac{\sigma k n}{b}}$$

with probability at least $1 - n^{-1}$. Again, the last inequality is due to [10]. Now applying Theorem 5, where we recall the parameter σ_e in (30) is $\sigma_e = \max_j \|E^j\|_F$ and ϵ^j is the optimization error at the estimation step, we have the following corollary.

Corollary 7 (for StoGradMP). *Assume the linear operator \mathcal{A} satisfies the \mathcal{D} -RSC and \mathcal{D} -RSS assumptions and $\xi \sim \mathcal{N}(0, \sigma^2)$. Set $p(i) = 1/M$ for $i = 1, \dots, M$. Then with probability at least $1 - n^{-1}$, the error at the $(t + 1)$ -th iterate of the StoGradMP algorithm is bounded by*

$$\mathbb{E} \|W^{t+1} - W_0\|_F \leq (0.9)^{t+1} \|W_0\|_F + c(1 + \eta_2) \sqrt{\frac{\sigma^2 k n}{b}} + 4 \max_{j \in [t]} \|E^j\|_F + 4 \max_{j \in [t]} \epsilon^j.$$

6 Numerical experiments

In this section we present some experimental results comparing our proposed stochastic methods to their deterministic counterparts. Our goal is to explore several interesting aspects of improvements and trade-offs; we have not attempted to optimize algorithm parameters. Unless otherwise specified, all experiments are run with at least 50 trials, “exact recovery” is obtained when the signal recovery error $\|w - \hat{w}\|_2$ drops below 10^{-6} , and the plots illustrate the 10% trimmed mean. For the approximation error versus epoch (or iteration) plots, the trimmed mean is calculated at each epoch (or iteration) value by excluding the highest and lowest 5% of the error values (rounding when necessary). For the approximation error versus CPU time plots, the trimmed mean computation is the same, except the CPU time values corresponding to the excluded error values are also excluded from the mean CPU time. We begin with experiments in the compressed sensing setting, and follow with application to the low-rank matrix recovery problem.

6.1 Sparse vector recovery

The first setting we explored is standard compressed sensing which has been studied in Subsection 5.1. Unless otherwise specified, the vector has dimension 256, and its non-zero entries are i.i.d.

standard Gaussian. The signal is measured with an $m \times 256$ i.i.d. standard Gaussian measurement matrix. First, we compare signal recovery as a function of the number of measurements used, for various sparsity levels s . Each algorithm terminates upon convergence or upon a maximum of 500 epochs². We used a block size of $b = \min(s, m)$ for sparsity level k_0 and number of measurements m , except when $k_0 = 4$ and $m > 5$ we used $b = 8$ in order to obtain consistent convergence. Specifically, this means that b measurements were used to form the signal proxy at each iteration of the stochastic algorithms. For the IHT and StoIHT algorithms, we use a step size of $\gamma = 1$. The results for IHT and StoIHT are shown in Figure 1 and GradMP and StoGradMP in Figure 2. Here we see that with these parameters, StoIHT requires far fewer measurements than IHT to recover the signal, whereas GradMP and StoGradMP are comparable.




Figure 1: Sparse Vector Recovery: Percent recovery as a function of the number of measurements for IHT (left) and StoIHT (right) for various sparsity levels k_0 .

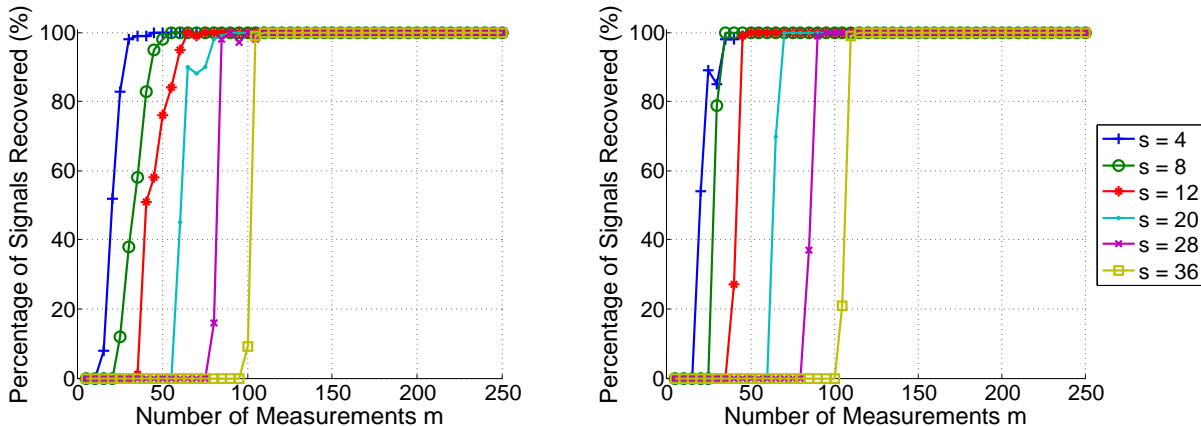


Figure 2: Sparse Vector Recovery: Percent recovery as a function of the number of measurements for GradMP (left) and StoGradMP (right) for various sparsity levels k_0 .

Next we explore how the choice of block size affects performance. We employ the same setup as described above, only now we fix the number of measurements m ($m = 180$ for the IHT methods

²We refer to an epoch as the number of iterations needed to use m rows. Thus for deterministic non-blocked methods, an epoch is one iteration, whereas for a method using blocks of size b , an epoch is m/b iterations.

and $m = 80$ for the GradMP methods), allow a maximum of 100 epochs, and use various block sizes in the stochastic algorithms. The sparsity of the signal is $k_0 = 8$. The results are depicted for both methods in Figure 3. Here we see that in both cases, the deterministic methods seem to offer intermediate performance, outperforming some block sizes and underperforming others. It is interesting that the StoIHT method seems to prefer larger block sizes whereas the StoGradMP seems to prefer smaller sizes. This is likely because StoGradMP, even using only a few gradients may still estimate the support accurately, and thus the signal accurately.




Figure 3: Sparse Vector Recovery: Recovery error as a function of epochs and various block sizes b for HT methods (left) and GradMP methods (right).

Next we repeat the same experiments but examine the recovery error as a function of the number of measurements for various block sizes (note that if the block size exceeds the number of measurements, we simply use the entire matrix as one block). Figure 4 shows these results. Because the methods exhibit graceful decrease in recovery error, here we plot the number of measurements (as a function of block size) required in order for the estimation error $\|w - \hat{w}\|_2$ to drop and remain below 10^{-6} . Although block size is not a parameter for the deterministic methods IHT and GradMP, a red horizontal line at the number of measurements required is included for comparison. We see that the fewest measurements are required when the block sizes are about 10 (recall the signal dimension is 256). We also note that StoIHT requires fewer measurements than IHT for large blocks, whereas StoGradMP requires the same as GradMP for large blocks, which is not surprising. However, we see that both methods offer improvements over their deterministic counterparts if the block sizes are chosen correctly.

6.2 Robustness to measurement noise

We next repeat the above sparse vector recovery experiments in the presence of noise in the measurements. All experiment parameters remain as in the previous setup, but a vector e of Gaussian noise with $\|e\|_2 = 0.5$ is added to the measurement vector. We again compare the recovery error against the number of epochs and measurements needed. The results are shown in Figures 5 and 6 for the IHT and GradMP algorithms, respectively. The right hand plots show the number of measurements required for the error to drop below the noise level 0.5 as a function of block size. Overall, the methods are robust to noise and demonstrate the same improvements and heuristics as in the noiseless experiments.




Figure 4: Sparse Vector Recovery: Number of measurements required for signal recovery as a function of block size (blue marker) for StoIHT (left) and StoGradMP (right). Number of measurements required for deterministic method shown as red solid line.




Figure 5: Sparse Vector Recovery: A comparison of IHT and StoIHT in the presence of noise. Recovery error versus epoch (left) and measurements required versus block size (right).

6.3 The choice of step size in StoIHT

Our last experiment in the sparse vector recovery setting explores the role of the step size γ in StoIHT. Keeping the dimension of the signal at 256, the sparsity $k_0 = 8$, the number of measurements $m = 80$, no noise, and fixing the block size $b = 8$, we test the algorithm using various values of the step size γ . The results are shown in Figure 7. We see that the value of γ clearly plays a role, but the range of successful values is quite large. Not surprisingly, too small of a step size leads to extremely slow convergence, and too large of one leads to divergence (at least initially).

6.4 Low-Rank Matrix Recovery

We now turn to the setting where we wish to recover a low-rank matrix W_0 from m linear measurements as studied in Subsection 5.2. Here W_0 is the 10×10 matrix with rank k_0 and we take m linear Gaussian measurements of the form $y_i = \langle A_i, W_0 \rangle$, where each A_i is a 10×10 matrix with i.i.d. standard Gaussian entries. As before, we first compare the percentage of exact recovery




Figure 6: Sparse Vector Recovery: A comparison of GradMP and StoGradMP in the presence of noise. Recovery error versus epoch (left) and measurements required versus block size (right).

(where again we deem the signal is recovered exactly when the error $\|W_0 - \hat{W}\|_F$ is below 10^{-6}) against the number of measurements required, for various rank levels. For the matrix case, we use a step size of $\gamma = 0.5$ for both the IHT and StoIHT methods, which seems to work well in this setting. The results for IHT and StoIHT are shown in Figure 8 and for GradMP and StoGradMP in Figure 9. For this choice of parameters, we see that both StoIHT and StoGradMP tend to require fewer measurements to recover the signal.

Next we examine the signal recovery error as a function of epoch, for various block sizes and against the deterministic methods. We fix the rank to be $k_0 = 2$ in these experiments. Because both block size and number of measurements affect the convergence, we see different behavior in the low measurement regime and the high measurement regime. This is apparent in Figure 10, where $m = 90$ measurements are used in the plot on the left and $m = 140$ measurements are used in the plot on the right, which shows the convergence of the IHT methods per epoch for various block sizes. We again see that for proper choices of block sizes, the StoIHT method outperforms IHT. It is also interesting to note that IHT seems to reach a higher noise floor than StoIHT. Of course we again point out that we have not optimized any of the algorithm parameters for either method. Results for the GradMP methods are shown in Figure 11, again where $m = 90$ measurements are used in the plot on the left and $m = 140$ measurements are used in the plot on the right. Similar to the IHT results, proper choices of block sizes allows StoGradMP to require much fewer epochs than GradMP to achieve convergence.

Figure 12 compares the block size and the number of measurements required for exact signal recovery for the IHT methods and the GradMP methods, again for a fixed rank of $k_0 = 2$. We again see that StoIHT and StoGradMP prefer small block sizes.

6.5 Recovery with approximations

Finally, we consider the important case where the identification, estimation, and pruning steps can only be performed approximately. In particular, we consider the case of low-rank matrix recovery in which these steps utilize only an approximate Singular Value Decomposition (SVD) of the matrix. This may be something that is unavoidable in certain applications, or may be desirable in others for computational speedup. For our first experiments of this kind, we use $N = 1024$ and generate a $N \times N$ rank $k_0 = 40$ matrix. We take m permuted rows of the $N \times N$ discrete Fourier transform as the measurement operator, use 2 blocks in the stochastic algorithms, and run

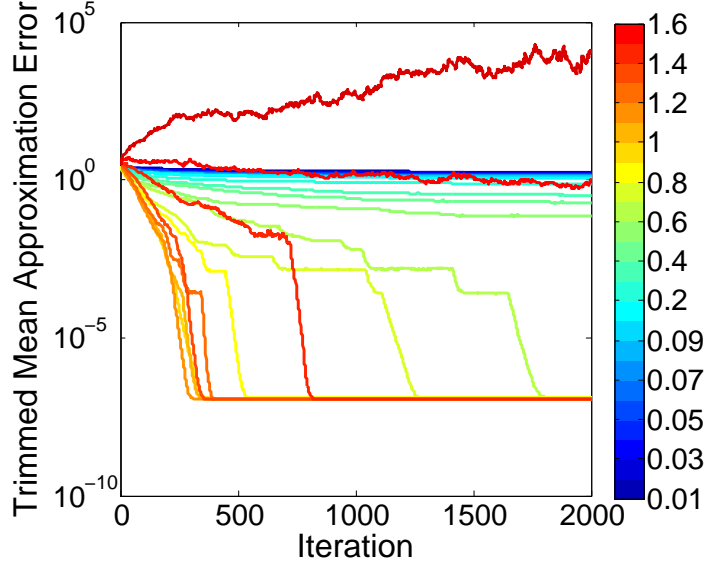


Figure 7: Sparse Vector Recovery: A comparison of StoIHT for various values of the step size γ (shown in the colorbar).

40 trials. In the StoIHT experiments we take $m = 0.3N^2$, and in the StoGradMP experiments we take $m = 0.35N^2$; these values for m empirically seemed to work well with the two algorithms. For each trial of the approximate SVD, we also run 5 sub-trials, to account for the randomness used in the approximate SVD algorithm. Here we use the randomized method described in [23] to compute the approximate SVD of a matrix. Briefly, to obtain a rank- s approximation of a matrix X and compute its approximated SVD, one applies the matrix to a randomly generated $N \times (s+d)$ matrix Ω to obtain the product $Y = X\Omega$ and constructs an orthonormal basis Q for the column space of Y . Here, d is an *over-sampling factor* that can be tuned to balance the tradeoff between accuracy and computation time. Using this basis, one computes the SVD of the product $B = Q^*X = U\Sigma V^*$, and approximates the SVD of X by $X \approx (QU)\Sigma V^*$. Because $(s+d)$ is typically much less than N , significant speedup can be gained. In addition, [23] proves that the approximation error is bounded by

$$\|X - X_s\|_F \leq \left(1 + \sqrt{\frac{s}{s+d}}\right) \|X - X_s^{\text{best}}\|_F,$$

where X_s^{best} is the best rank- s approximation of X and X_s is the approximate rank- s matrix produced from the above procedure. Here, the multiplicative error is associated with the quantity η in the approximation operator $\text{approx}_s(w, \eta)$ defined in (8).

Figure 13 shows the approximation error as a function of epoch and runtime for the StoIHT algorithm, for various over-sampling factors d as well as the full SVD computation for comparison. We again use a 10% trimmed mean over all the trials, and a step-size of $\gamma = 0.5$. We see that in terms of epochs, for reasonably sized over-sampling factors, the convergence using the SVD approximation is very similar to that of using the full SVD. In terms of runtime, we see a significant speedup for moderate choices of over-sampling factor, as expected. Recall that 2 blocks were used for this experiment, but we have observed a very similar relationship between the curves when increasing the number of blocks to 10.

The analogous results for StoGradMP are very similar, and are shown in Figure 14. We again




Figure 8: Low-Rank Matrix Recovery: Percent recovery as a function of the number of measurements for IHT (left) and StoIHT (right) for various rank levels s .

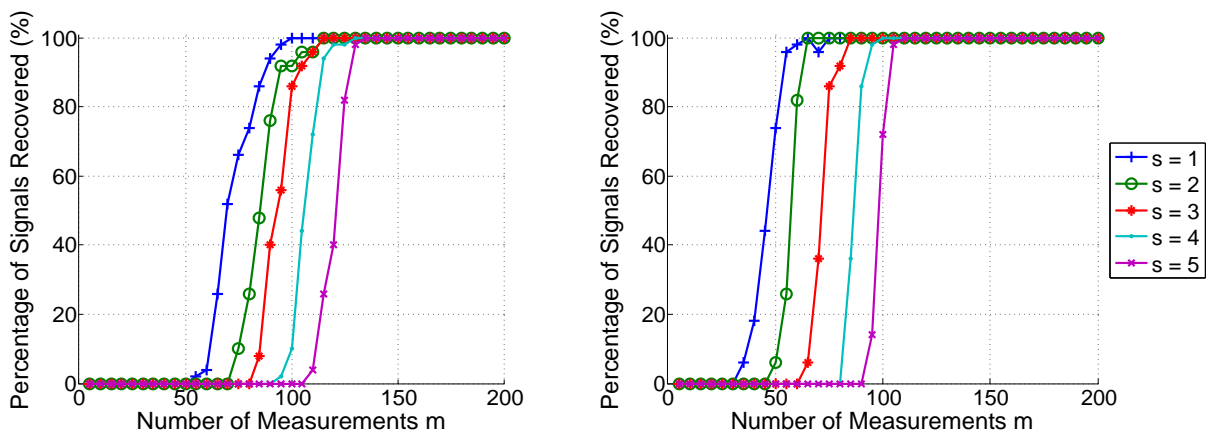


Figure 9: Low-Rank Matrix Recovery: Percent recovery as a function of the number of measurements for GradMP (left) and StoGradMP (right) for various rank levels k_0 .

see that for certain over-sampling factors, the convergence of the approximation error as a function of epoch is similar when using the approximate SVD and the full SVD. We also see a very significant speedup when using the approximate SVD; in this case, all the over-sampling factors used in this experiment offer an improved runtime over the full SVD computation.

7 Conclusion

We study in this paper two stochastic algorithms to solve a possibly non-convex optimization problem with the constraint that the solution has a simple representation with respect to a predefined atom set. This type of optimization has found tremendous applications in signal processing, machine learning, and statistics such as sparse signal recovery and low-rank matrix estimation. Our proposed algorithms, called StoIHT and StoGradMP, have their roots back to the celebrated IHT and CoSaMP algorithms, from which we have made several significant extensions. The first extension is to transfer algorithmic ideas of IHT and CoSaMP to the stochastic setting and the second extension is the allowance of approximate projections at each iteration of the algorithms. More




Figure 10: Low-Rank Matrix Recovery: Recovery error as a function of the number of epochs for StoIHT methods using $m = 90$ (left) and $m = 140$ (right) for various block sizes b .




Figure 11: Low-Rank Matrix Recovery: Recovery error as a function of the number of epochs for the StoGradMP algorithm using $m = 90$ (left) and $m = 140$ (right) for various block sizes b .

importantly, we theoretically prove that the stochastic versions with inexact projections enjoy the same linear convergence rate as their deterministic counterparts. We also show that the algorithms behave predictably even when the gradients are contaminated by noise. Experimentally, stochastic approaches have shown particular advantages over the deterministic counterparts in many problems of interest such as linear regression and matrix recovery.

8 Proofs

8.1 Consequences of the \mathcal{D} -RSC and \mathcal{D} -RSS

The first corollary provides a useful upper bound for the gradient, which we call co-coercivity.

Corollary 8. *Assume the function $f(w)$ satisfies the \mathcal{D} -RSS property, then*

$$\langle w' - w, \nabla f(w') - \nabla f(w) \rangle \leq \rho_s^+ \|w' - w\|_2^2 \quad (37)$$

for all vectors w and w' of size n such that $|\text{supp}_{\mathcal{D}}(w) \cup \text{supp}_{\mathcal{D}}(w')| \leq s$. In addition, let $\Omega =$




Figure 12: Low-Rank Matrix Recovery: Number of measurements required for signal recovery as a function of block size (blue marker) for StoIHT (left) and StoGradMP (right). Number of measurements required for deterministic method shown as red solid line.




Figure 13: Low-Rank Matrix Recovery with Approximations: Trimmed mean recovery error as a function of epochs (left) and runtime (right) for various over-sampling factors d using the StoIHT algorithm. Performance using full SVD computation shown as dashed line.

$\text{supp}_{\mathcal{D}}(w) \cup \text{supp}_{\mathcal{D}}(w')$; then we have

$$[\mathbf{Co-coercivity}] \quad \|\mathcal{P}_{\Omega}(\nabla f(w') - \nabla f(w))\|_2^2 \leq \rho_s^+ \langle w' - w, \nabla f(w') - \nabla f(w) \rangle. \quad (38)$$

Proof. From the definition of \mathcal{D} -RSS, we can show that

$$f(w') - f(w) - \langle \nabla f(w), w' - w \rangle \leq \frac{\rho_s^+}{2} \|w' - w\|_2^2. \quad (39)$$

Similarly, interchanging the role of w and w' , we have

$$f(w) - f(w') - \langle \nabla f(w'), w - w' \rangle \leq \frac{\rho_s^+}{2} \|w' - w\|_2^2.$$

Taking the summation of these two inequalities leads to the first claim.




Figure 14: Low-Rank Matrix Recovery with Approximations: Trimmed mean recovery error as a function of epochs (left) and runtime (right) for various over-sampling factors d using the StoGradMP algorithm. Performance using full SVD computation shown as dashed line.

To prove the second claim, we define a function $G(x) \triangleq f(x) - \langle \nabla f(w), x \rangle$, it is easy to see that for any x and y with $\text{supp}_{\mathcal{D}}(x) \cup \text{supp}_{\mathcal{D}}(y) \in \Omega$, we have

$$\|\nabla G(x) - \nabla G(y)\|_2 = \|\nabla f(x) - \nabla f(y)\|_2 \leq \rho_s^+ \|x - y\|_2.$$

This implies that $G(x)$ has \mathcal{D} -RSS with constant ρ_s^+ . In particular, we get a similar inequality as in (39)

$$G(x) - G(y) - \langle \nabla G(y), x - y \rangle \leq \frac{\rho_s^+}{2} \|x - y\|_2^2. \quad (40)$$

We also observe that

$$G(x) - G(w) = f(x) - f(w) - \langle \nabla f(w), x - w \rangle \geq 0$$

for all x such that $\text{supp}_{\mathcal{D}}(x) \in \Omega$. Let $x \triangleq w' - \frac{1}{\rho_s^+} \mathcal{P}_{\Omega} \nabla G(w')$, then it is clear that $\text{supp}_{\mathcal{D}}(x) \in \Omega$. Thus, by \mathcal{D} -RSS property of $G(x)$, we have

$$\begin{aligned} G(w) &\leq G\left(w' - \frac{1}{\rho_s^+} \mathcal{P}_{\Omega} \nabla G(w')\right) \\ &\leq G(w') + \left\langle \nabla G(w'), -\frac{1}{\rho_s^+} \mathcal{P}_{\Omega} \nabla G(w') \right\rangle + \frac{1}{2\rho_s^+} \|\mathcal{P}_{\Omega} \nabla G(w')\|_2^2 \\ &= G(w') - \frac{1}{2\rho_s^+} \|\mathcal{P}_{\Omega} \nabla G(w')\|_2^2. \end{aligned}$$

Replacing the function $G(x)$ in this equality we get

$$\frac{1}{2\rho_s^+} \|\mathcal{P}_{\Omega}(\nabla f(w') - \nabla f(w))\|_2^2 \leq f(w') - f(w) - \langle \nabla f(w), w' - w \rangle.$$

The claim follows by adding the two inequalities with w and w' interchanged. \square

The following corollary provides the lower bound for the gradient.

Corollary 9. Assume the function $F(w)$ satisfies the \mathcal{D} -RSC, then

$$\rho_s^- \|w' - w\|_2^2 \leq \langle w' - w, \nabla F(w') - \nabla F(w) \rangle \quad (41)$$

for all w and w' such that $|\text{supp}_{\mathcal{D}}(w) \cup \text{supp}_{\mathcal{D}}(w')| \leq s$.

Proof. From the \mathcal{D} -RSC assumption, we can write

$$F(w') - F(w) - \langle \nabla F(w), w' - w \rangle \geq \frac{\rho_s^-}{2} \|w' - w\|_2^2.$$

Swapping w and w' , we also have

$$F(w) - F(w') - \langle \nabla F(w'), w - w' \rangle \geq \frac{\rho_s^-}{2} \|w' - w\|_2^2.$$

The result follows by adding the two inequalities. \square

The next corollary provides key estimates for our convergence analysis. Recall that we assume $\{f_i(w)\}_{i=1}^M$ satisfies the \mathcal{D} -RSS and $F(w) = \sum_{i=1}^M f_i(w)$ satisfies the \mathcal{D} -RSC.

Corollary 10. Let i be an index selected with probability $p(i)$ from the set $[M]$. For any fixed sparse vectors w and w' , let Ω be a set such that $\text{supp}_{\mathcal{D}}(w) \cup \text{supp}_{\mathcal{D}}(w') \in \Omega$ and denote $s = |\Omega|$. We have

$$\mathbb{E}_i \left\| w' - w - \frac{\gamma}{Mp(i)} \mathcal{P}_{\Omega} (\nabla f_i(w') - \nabla f_i(w)) \right\|_2 \leq \sqrt{1 - (2 - \gamma\alpha_s)\gamma\rho_s^-} \|w' - w\|_2 \quad (42)$$

where we define $\alpha_s \triangleq \max_i \frac{\rho_s^+(i)}{Mp(i)}$. In addition, we have

$$\mathbb{E}_i \left\| w' - w - \frac{\gamma}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\|_2 \leq \sqrt{1 + \gamma^2\alpha_s\bar{\rho}_s^+ - 2\gamma\rho_s^-} \|w' - w\|_2, \quad (43)$$

where $\bar{\rho}_s^+ \triangleq \frac{1}{M} \sum_i \rho_s^+(i)$.

The difference between the estimates (42) and (43) is the additional term $\|\mathcal{P}_{\Omega^c} (\nabla f_i(w') - \nabla f_i(w))\|_2$ with $\Omega^c = [n] \setminus \Omega$ appearing in (43).

Proof. We will use the co-coercivity property that appeared in inequality (38) in Corollary 8. We

have

$$\begin{aligned}
& \mathbb{E}_i \left\| w' - w - \frac{\gamma}{Mp(i)} \mathcal{P}_\Omega (\nabla f_i(w') - \nabla f_i(w)) \right\|_2^2 \\
&= \|w' - w\|_2^2 + \mathbb{E}_i \frac{\gamma^2}{(Mp(i))^2} \|\mathcal{P}_\Omega (\nabla f_i(w') - \nabla f_i(w))\|_2^2 \\
&\quad - 2\gamma \mathbb{E}_i \left\langle w' - w, \mathcal{P}_\Omega \frac{1}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\rangle \\
&\leq \|w' - w\|_2^2 + \gamma^2 \mathbb{E}_i \frac{\rho_s^+(i)}{(Mp(i))^2} \langle w' - w, \nabla f_i(w') - \nabla f_i(w) \rangle \\
&\quad - 2\gamma \mathbb{E}_i \left\langle w' - w, \frac{1}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\rangle \\
&\leq \|w' - w\|_2^2 + \gamma^2 \max_i \frac{\rho_s^+(i)}{Mp(i)} \mathbb{E}_i \left\langle w' - w, \frac{1}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\rangle \\
&\quad - 2\gamma \mathbb{E}_i \left\langle w' - w, \frac{1}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\rangle \\
&= \|w' - w\|_2^2 - \left(2\gamma - \gamma^2 \max_i \frac{\rho_s^+(i)}{Mp(i)} \right) \mathbb{E}_i \left\langle w' - w, \frac{1}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\rangle \\
&= \|w' - w\|_2^2 - (2\gamma - \gamma^2 \alpha_s) \langle w' - w, \nabla F(w') - \nabla F(w) \rangle \\
&\leq \|w' - w\|_2^2 - (2\gamma - \gamma^2 \alpha_s) \rho_s^- \|w' - w\|_2^2,
\end{aligned}$$

where the first inequality follows from (38) and the last inequality follows from (41). Applying the known result $(\mathbb{E}Z)^2 \leq \mathbb{E}Z^2$ completes the proof of (42).

The proof of (43) is similar to that of (42), except now we are not able to apply the co-coercivity inequality. Expanding the left hand side and applying the definition of \mathcal{D} -RSS together with the inequality (41), we derive

$$\begin{aligned}
& \mathbb{E}_i \left\| w' - w - \frac{\gamma}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\|_2^2 \\
&= \|w' - w\|_2^2 + \mathbb{E}_i \frac{\gamma^2}{(Mp(i))^2} \|\nabla f_i(w') - \nabla f_i(w)\|_2^2 \\
&\quad - 2\gamma \mathbb{E}_i \left\langle w' - w, \frac{1}{Mp(i)} (\nabla f_i(w') - \nabla f_i(w)) \right\rangle \\
&= \|w' - w\|_2^2 + \mathbb{E}_i \frac{\gamma^2}{(Mp(i))^2} \|\nabla f_i(w') - \nabla f_i(w)\|_2^2 \\
&\quad - 2\gamma \langle w' - w, \nabla F(w') - \nabla F_i(w) \rangle \\
&\leq \|w' - w\|_2^2 + \mathbb{E}_i \frac{\gamma^2}{(Mp(i))^2} (\rho_s^+(i))^2 \|w' - w\|_2^2 - 2\gamma \rho_s^- \|w' - w\|_2^2.
\end{aligned}$$

We further have

$$\mathbb{E}_i \frac{\gamma^2}{(Mp(i))^2} (\rho_s^+(i))^2 \leq \gamma^2 \max_i \frac{\rho_s^+(i)}{Mp(i)} \mathbb{E}_i \frac{\rho_s^+(i)}{Mp(i)} = \gamma^2 \alpha_s \sum_i \frac{\rho_s^+(i)}{Mp(i)} p(i) = \gamma^2 \alpha_s \bar{\rho}_s^+,$$

where we recall that $\alpha_s = \max_i \frac{\rho_s^+(i)}{Mp(i)}$ and $\bar{\rho}_s^+ = \frac{1}{M} \sum_i \rho_s^+(i)$. Substitute this result into the above inequality and then use the inequality $(\mathbb{E}Z)^2 \leq \mathbb{E}Z^2$ to complete the proof. \square

8.2 Proof of Theorem 1

Proof of Theorem 1. We notice that $\|w^{t+1} - b^t\|_2 \leq \eta \|b_k^t - b^t\|_2 \leq \eta \|w^* - b^t\|_2$ where b_k^t is the best k -sparse approximation of b^t with respect to \mathcal{D} . Thus,

$$\|w^{t+1} - w^* + w^* - b^t\|_2^2 \leq \eta^2 \|w^* - b^t\|_2^2.$$

Expanding the left hand side of this inequality leads to

$$\begin{aligned} \|w^{t+1} - w^*\|_2^2 &\leq 2 \langle w^{t+1} - w^*, b^t - w^* \rangle + (\eta^2 - 1) \|b^t - w^*\|_2^2 \\ &= 2 \left\langle w^{t+1} - w^*, w^t - w^* - \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^t) \right\rangle \\ &\quad + (\eta^2 - 1) \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^t) \right\|_2^2 \\ &= 2 \left\langle w^{t+1} - w^*, w^t - w^* - \frac{\gamma}{Mp(i_t)} (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\rangle \\ &\quad - 2 \left\langle w^{t+1} - w^*, \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^*) \right\rangle + (\eta^2 - 1) \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^t) \right\|_2^2. \end{aligned}$$

Denote $\Omega = \text{supp}_{\mathcal{D}}(w^{t+1}) \cup \text{supp}_{\mathcal{D}}(w^t) \cup \text{supp}_{\mathcal{D}}(w^*)$ and notice that $|\Omega| \leq 3k$, we get

$$\begin{aligned} \|w^{t+1} - w^*\|_2^2 &\leq 2 \left\langle w^{t+1} - w^*, w^t - w^* - \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\rangle \\ &\quad - 2 \left\langle w^{t+1} - w^*, \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} \nabla f_{i_t}(w^*) \right\rangle + (\eta^2 - 1) \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^t) \right\|_2^2 \\ &\leq 2 \|w^{t+1} - w^*\|_2 \\ &\quad \times \underbrace{\left(\left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\|_2 + \left\| \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} \nabla f_{i_t}(w^*) \right\|_2 \right)}_u \\ &\quad + \underbrace{(\eta^2 - 1) \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^t) \right\|_2^2}_v. \end{aligned}$$

Solving this quadratic inequality $x^2 - 2ux - v \leq 0$ with $x = \|w^{t+1} - w^*\|_2$, we get $x \leq u + \sqrt{u^2 + v} \leq 2u + \sqrt{v}$. Substituting the expressions for u and v above, we arrive at

$$\begin{aligned} \|w^{t+1} - w^*\|_2 &\leq 2 \left(\left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\|_2 + \left\| \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} \nabla f_{i_t}(w^*) \right\|_2 \right) \\ &\quad + \sqrt{\eta^2 - 1} \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^t) \right\|_2 \\ &\leq 2 \left(\left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\|_2 + \left\| \frac{\gamma}{Mp(i_t)} \mathcal{P}_{\Omega} \nabla f_{i_t}(w^*) \right\|_2 \right) \\ &\quad + \sqrt{\eta^2 - 1} \left(\left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\|_2 + \left\| \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^*) \right\|_2 \right). \end{aligned}$$

Denote I_t as the set containing all indices i_1, i_2, \dots, i_t randomly selected at or before step t of the algorithm: $I_t = \{i_1, \dots, i_t\}$. It is clear that I_t determines the solutions w^1, \dots, w^{t+1} . We also

denote the conditional expectation $\mathbb{E}_{i_t|I_{t-1}} \|w^{t+1} - w^*\|_2 \triangleq \mathbb{E}_{i_t}(\|w^{t+1} - w^*\|_2 | I_{t-1})$. Now taking the conditional expectation on both sides of the above inequality we obtain

$$\begin{aligned} & \mathbb{E}_{i_t|I_{t-1}} \|w^{t+1} - w^*\|_2 \\ & \leq 2 \left(\mathbb{E}_{i_t|I_{t-1}} \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \mathcal{P}_\Omega (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\|_2 + \mathbb{E}_{i_t|I_{t-1}} \left\| \frac{\gamma}{Mp(i_t)} \mathcal{P}_\Omega \nabla f_{i_t}(w^*) \right\|_2 \right) \\ & \quad + \sqrt{\eta^2 - 1} \left(\mathbb{E}_{i_t|I_{t-1}} \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} (\nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)) \right\|_2 + \mathbb{E}_{i_t|I_{t-1}} \left\| \frac{\gamma}{Mp(i_t)} \nabla f_{i_t}(w^*) \right\|_2 \right). \end{aligned}$$

Conditioning on I_{t-1} , w^t can be seen as a fixed vector. We apply the inequality (42) of Corollary 10 for the first term and (43) for the third term, we get

$$\begin{aligned} & \mathbb{E}_{i_t|I_{t-1}} \|w^{t+1} - w^*\|_2 \\ & \leq 2\sqrt{(1 - (2\gamma - \gamma^2\alpha_{3k})\rho_{3k}^-)} \|w^t - w^*\|_2 + 2\frac{\gamma}{\min_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\mathcal{P}_\Omega \nabla f_{i_t}(w^*)\|_2 \\ & \quad + \sqrt{\eta^2 - 1} \sqrt{1 + \gamma^2\alpha_{3k}\bar{\rho}_{3k}^+ - 2\gamma\rho_{3k}^-} \|w^t - w^*\|_2 + \sqrt{\eta^2 - 1} \frac{\gamma}{\min_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\nabla f_{i_t}(w^*)\|_2 \\ & = \left(2\sqrt{(1 - (2\gamma - \gamma^2\alpha_{3k})\rho_{3k}^-)} + \sqrt{(\eta^2 - 1)(1 + \gamma^2\alpha_{3k}\bar{\rho}_{3k}^+ - 2\gamma\rho_{3k}^-)} \right) \|w^t - w^*\|_2 \\ & \quad + \frac{\gamma}{\min_{i_t} Mp(i_t)} \left(2\mathbb{E}_{i_t} \|\mathcal{P}_\Omega \nabla f_{i_t}(w^*)\|_2 + \sqrt{\eta^2 - 1} \mathbb{E}_{i_t} \|\nabla f_{i_t}(w^*)\|_2 \right) \\ & \leq \kappa \|w^t - w^*\|_2 + \sigma, \end{aligned}$$

where κ and σ_{w^*} are defined in Theorem 1. Taking the expectation on both sides with respect to I_{t-1} yields

$$\mathbb{E}_{I_t} \|w^{t+1} - w^*\|_2 \leq \kappa \mathbb{E}_{I_{t-1}} \|w^t - w^*\|_2 + \sigma.$$

Applying this result recursively over t iterations yields the desired result:

$$\begin{aligned} \mathbb{E}_{I_t} \|w^{t+1} - w^*\|_2 & \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \sum_{j=0}^t \kappa^j \sigma \\ & \leq \kappa^{t+1} \|w^0 - w^*\|_2 + \frac{1}{1 - \kappa} \sigma. \end{aligned}$$

□

8.3 Proof of Theorem 2

The proof of Theorem 2 is a consequence of the following three lemmas. Denote I_t as the set containing all indices i_1, i_2, \dots, i_t randomly selected at or before step t of the algorithm: $I_t = \{i_1, \dots, i_t\}$ and denote the conditional expectation $\mathbb{E}_{i_t|I_{t-1}} \|w^{t+1} - w^*\|_2 \triangleq \mathbb{E}_{i_t}(\|w^{t+1} - w^*\|_2 | I_{t-1})$.

Lemma 1. *The recovery error at the $(t+1)$ -th iteration is upper bounded by*

$$\|w^{t+1} - w^*\|_2 \leq (1 + \eta_2) \|b^t - w^*\|_2.$$

Lemma 2. *Denote $\widehat{\Gamma}$ as the set obtained from the t -th iteration and i as the index selected randomly from $[M]$ with probability $p(i)$. We have,*

$$\mathbb{E}_{I_t} \|b^t - w^*\|_2 \leq \sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} \mathbb{E}_{I_t} \|P_{\widehat{\Gamma}^c}(b^t - w^*)\|_2 + \sigma_1$$

where $\alpha_k = \max_i \frac{\rho_k^+(i)}{Mp(i)}$ and

$$\sigma_1 \triangleq \frac{3}{\rho_{4k}^-} \frac{1}{\min_i Mp(i)} \max_{|\Omega| \leq 3k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2.$$

Lemma 3. Denote $\hat{\Gamma}$ as the set obtained from the t -th iteration. Then,

$$\mathbb{E}_{I_t} \|\mathcal{P}_{\hat{\Gamma}^c}(b^t - w^*)\|_2 \leq \left(\max_i \sqrt{Mp(i)} \sqrt{\frac{\frac{2\eta_1^2-1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-}{\rho_{4k}^-}} + \frac{\sqrt{\eta_1^2-1}}{\eta_1} \right) \|w^t - w^*\|_2 + \sigma_2, \quad (44)$$

where

$$\sigma_2 \triangleq \frac{2 \max_{i \in [M]} p(i)}{\rho_{4k}^- \min_{i \in [M]} p(i)} \max_{|\Omega| \leq 4k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2.$$

We are now able to prove Theorem 2. We have a series of inequalities that follow from the above lemmas,

$$\begin{aligned} & \mathbb{E}_{I_t} \|w^{t+1} - w^*\|_2 \\ & \leq (1 + \eta_2) \mathbb{E}_{I_t} \|b^t - w^*\|_2 \quad (\text{Lemma 1}) \\ & \leq (1 + \eta_2) \sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} \mathbb{E}_{I_t} \|P_{\hat{\Gamma}^c}(b^t - w^*)\|_2 + (1 + \eta_2) \sigma_1 \quad (\text{Lemma 2}) \\ & \leq (1 + \eta_2) \sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} \left(\max_i \sqrt{Mp(i)} \sqrt{\frac{\frac{2\eta_1^2-1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-}{\rho_{4k}^-}} + \frac{\sqrt{\eta_1^2-1}}{\eta_1} \right) \mathbb{E}_{I_{t-1}} \|w^t - w^*\|_2 \\ & + (1 + \eta_2) \left(\sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} \sigma_2 + \sigma_1 \right), \end{aligned} \quad (45)$$

where the last inequality follows from Lemma 3. Replacing the definition of κ in (22) and noticing that σ_{w^*} defined in (23) is greater than the second term of the last equation (it is due to $\max_{|\Omega| \leq 4k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2 \geq \max_{|\Omega| \leq 3k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2$), we arrive at

$$\mathbb{E}_{I_t} \|w^{t+1} - w^*\|_2 \leq \kappa \mathbb{E}_{I_{t-1}} \|w^t - w^*\|_2 + \sigma_{w^*}. \quad (46)$$

Applying this inequality recursively t times will complete the proof.

To the end of this section, we prove three lemmas stated above.

Proof of Lemma 1. Recall that b^t is the vector obtained from the t -th iteration. From the algorithm, we have

$$\|w^{t+1} - b^t\|_2 \leq \eta_2 \|b_k^t - b^t\|_2 \leq \eta_2 \|w^* - b^t\|_2$$

where b_k^t is the best k -sparse approximation of b^t with respect to the set \mathcal{D} . We thus have

$$\begin{aligned} \|w^{t+1} - w^*\|_2 & \leq \|(w^{t+1} - b^t) + (b^t - w^*)\|_2 \\ & \leq \|w^{t+1} - b^t\|_2 + \|b^t - w^*\|_2 \leq (1 + \eta_2) \|b^t - w^*\|_2. \end{aligned} \quad (47)$$

□

Proof of Lemma 2. Denote the set $\mathcal{C}_{\hat{\Gamma}} \triangleq \{w : w = \sum_{j \in \hat{\Gamma}} \alpha_j d_j\}$. It is clear that $\mathcal{C}_{\hat{\Gamma}}$ is a convex set, so the estimation step can be written as

$$b^t = \underset{w}{\operatorname{argmin}} F(w) \quad \text{such that} \quad w \in \mathcal{C}_{\hat{\Gamma}}.$$

Optimization theory states that (Proposition 4.7.1 of [4])

$$\langle \nabla F(b^t), b^t - z \rangle \leq 0 \quad \text{for all } z \in \mathcal{C}_{\hat{\Gamma}}.$$

Put differently, we have

$$\langle \nabla F(b^t), \mathcal{P}_{\hat{\Gamma}}(b^t - z) \rangle \leq 0 \quad \text{for all } z.$$

Denote by i an index selected randomly from $[M]$ with probability $p(i)$ and independent from all the random indices i_t and recall that $\nabla F(b^t) = \mathbb{E}_i \frac{1}{Mp(i)} \nabla f_i(b^t)$. The above inequality can be read as

$$0 \geq \left\langle \mathbb{E}_i \frac{1}{Mp(i)} \nabla f_i(b^t), \mathcal{P}_{\hat{\Gamma}}(b^t - z) \right\rangle = \mathbb{E}_i \left\langle \frac{1}{Mp(i)} \mathcal{P}_{\hat{\Gamma}} \nabla f_i(b^t), \mathcal{P}_{\hat{\Gamma}}(b^t - z) \right\rangle \quad \text{for all } z. \quad (48)$$

We first derive the upper bound of $\|\mathcal{P}_{\hat{\Gamma}}(b^t - w^*)\|_2$. For any $\gamma > 0$, we have

$$\begin{aligned} \|\mathcal{P}_{\hat{\Gamma}}(b^t - w^*)\|_2^2 &= \langle \mathcal{P}_{\hat{\Gamma}}(b^t - w^*), b^t - w^* \rangle \\ &= \left\langle \mathcal{P}_{\hat{\Gamma}}(b^t - w^*), b^t - w^* - \mathbb{E}_i \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}}(\nabla f_i(b^t) - \nabla f_i(w^*)) \right\rangle \\ &\quad + \left\langle \mathcal{P}_{\hat{\Gamma}}(b^t - w^*), \mathbb{E}_i \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}} \nabla f_i(b^t) \right\rangle - \left\langle \mathcal{P}_{\hat{\Gamma}}(b^t - w^*), \mathbb{E}_i \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}} \nabla f_i(w^*) \right\rangle \\ &= \mathbb{E}_i \left\langle \mathcal{P}_{\hat{\Gamma}}(b^t - w^*), b^t - w^* - \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}}(\nabla f_i(b^t) - \nabla f_i(w^*)) \right\rangle \\ &\quad + \mathbb{E}_i \left\langle \mathcal{P}_{\hat{\Gamma}}(b^t - w^*), \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}} \nabla f_i(b^t) \right\rangle - \mathbb{E}_i \left\langle \mathcal{P}_{\hat{\Gamma}}(b^t - w^*), \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}} \nabla f_i(w^*) \right\rangle \\ &\leq \|\mathcal{P}_{\hat{\Gamma}}(b^t - w^*)\|_2 \mathbb{E}_i \left\| b^t - w^* - \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}}(\nabla f_i(b^t) - \nabla f_i(w^*)) \right\|_2 \\ &\quad + \|\mathcal{P}_{\hat{\Gamma}}(b^t - w^*)\|_2 \mathbb{E}_i \frac{\gamma}{Mp(i)} \|\mathcal{P}_{\hat{\Gamma}} \nabla f_i(w^*)\|_2 \end{aligned}$$

where the inequality follows from (48) and the Cauchy-Schwarz inequality. Canceling the common term in both sides, we derive

$$\|\mathcal{P}_{\hat{\Gamma}}(b^t - w^*)\|_2 \leq \mathbb{E}_i \left\| b^t - w^* - \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}}(\nabla f_i(b^t) - \nabla f_i(w^*)) \right\|_2 + \mathbb{E}_i \frac{\gamma}{Mp(i)} \|\mathcal{P}_{\hat{\Gamma}} \nabla f_i(w^*)\|_2.$$

We bound the first term of the right-hand side. For a fixed realization of the random vector b^t , we apply Corollary 10 to obtain

$$\mathbb{E}_i \left\| b^t - w^* - \frac{\gamma}{Mp(i)} \mathcal{P}_{\hat{\Gamma}}(\nabla f_i(b^t) - \nabla f_i(w^*)) \right\|_2 \leq \sqrt{(1 - (2\gamma - \gamma^2 \alpha_{4k}) \rho_{4k}^-)} \|b^t - w^*\|_2. \quad (49)$$

Applying this result to the above inequality and taking the expectation with respect to i_t yields

$$\|\mathcal{P}_{\hat{\Gamma}}(b^t - w^*)\|_2 \leq \sqrt{(1 - (2\gamma - \gamma^2 \alpha_{4k}) \rho_{4k}^-)} \|b^t - w^*\|_2 + \frac{\gamma}{\min_i Mp(i)} \mathbb{E}_i \|\mathcal{P}_{\hat{\Gamma}} \nabla f_i(w^*)\|_2, \quad (50)$$

We now apply this inequality to get

$$\begin{aligned}
\|b^t - w^*\|_2^2 &= \|\mathcal{P}_{\widehat{\Gamma}}(b^t - w^*)\|_2^2 + \|\mathcal{P}_{\widehat{\Gamma}^c}(b^t - w^*)\|_2^2 \\
&\leq \left(\sqrt{(1 - (2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-)} \|b^t - w^*\|_2 + \frac{\gamma}{\min_i Mp(i)} \mathbb{E}_i \|\mathcal{P}_{\widehat{\Gamma}} \nabla f_i(w^*)\|_2 \right)^2 \\
&\quad + \|\mathcal{P}_{\widehat{\Gamma}^c}(b^t - w^*)\|_2^2.
\end{aligned} \tag{51}$$

Solving the quadratic polynomial $ax^2 - 2bx - c \leq 0$ with $x = \|b^t - w^*\|_2$, $a = (2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-$, $b = \sqrt{(1 - (2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-)} \frac{\gamma}{\min_i Mp(i)} \mathbb{E}_i \|\mathcal{P}_{\widehat{\Gamma}} \nabla f_i(w^*)\|_2$, and $c = (\frac{\gamma}{\min_i Mp(i)} \mathbb{E}_i \|\mathcal{P}_{\widehat{\Gamma}} \nabla f_i(w^*)\|_2)^2 + \|\mathcal{P}_{\widehat{\Gamma}^c}(b^t - w^*)\|_2^2$, we get

$$\|b^t - w^*\|_2 \leq \frac{b + \sqrt{b^2 + ac}}{a} \leq \sqrt{\frac{c}{a}} + \frac{2b}{a}.$$

Replacing these quantities a , b , and c yields

$$\begin{aligned}
\|b^t - w^*\|_2 &\leq \frac{1}{\sqrt{(2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-}} \|\mathcal{P}_{\widehat{\Gamma}^c}(b^t - w^*)\|_2 \\
&\quad + \left(\frac{1}{\sqrt{(2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-}} + \frac{2\sqrt{(1 - (2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-)}}{(2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-} \right) \frac{\gamma}{\min_i Mp(i)} \mathbb{E}_i \|\mathcal{P}_{\widehat{\Gamma}} \nabla f_i(w^*)\|_2 \\
&\leq \frac{1}{\sqrt{(2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-}} \|\mathcal{P}_{\widehat{\Gamma}^c}(b^t - w^*)\|_2 + \frac{3}{(2\gamma - \gamma^2\alpha_{4k})\rho_{4k}^-} \frac{\gamma}{\min_i Mp(i)} \mathbb{E}_i \|\mathcal{P}_{\widehat{\Gamma}} \nabla f_i(w^*)\|_2.
\end{aligned}$$

Optimizing γ that maximizes $(2\gamma - \gamma^2\alpha_{4k})$, we get $\gamma = \frac{1}{\alpha_{4k}}$. Plugging this value into the above inequality and taking the expectation with respect to I_t (notice that the random variable b^t is determined by random indices i_0, \dots, i_t)

$$\mathbb{E}_{I_t} \|b^t - w^*\|_2 \leq \sqrt{\frac{\alpha_{4k}^-}{\rho_{4k}^-}} \mathbb{E}_{I_t} \|\mathcal{P}_{\widehat{\Gamma}^c}(b^t - w^*)\|_2 + \frac{3}{\rho_{4k}^-} \frac{1}{\min_i Mp(i)} \mathbb{E}_{i, I_t} \|\mathcal{P}_{\widehat{\Gamma}} \nabla f_i(w^*)\|_2.$$

The proof is completed. \square

Proof of Lemma 3. Since b^t and w^t are in $\text{span}(\mathcal{D}_{\widehat{\Gamma}})$, we have $\mathcal{P}_{\widehat{\Gamma}^c} b^t = 0$ and $\mathcal{P}_{\widehat{\Gamma}^c} w^t = 0$. Therefore,

$$\|\mathcal{P}_{\widehat{\Gamma}^c}(b^t - w^*)\|_2 = \|\mathcal{P}_{\widehat{\Gamma}^c}(w^t - w^*)\|_2 \leq \|\mathcal{P}_{\Gamma^c}(w^t - w^*)\|_2 = \|\Delta - \mathcal{P}_{\Gamma}\Delta\|_2, \tag{52}$$

where we denote $\Delta \triangleq w^* - w^t$. The goal is to estimate $\|\Delta - \mathcal{P}_{\Gamma}\Delta\|_2$. Let $R \triangleq \text{supp}_{\mathcal{D}}(\Delta)$ and apply the \mathcal{D} -RSC, we have

$$\begin{aligned}
F(w^*) - F(w^t) - \frac{\rho_{4k}^-}{2} \|w^* - w^t\|_2^2 &\geq \langle \nabla F(w^t), w^* - w^t \rangle \\
&= \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), \Delta \right\rangle \\
&= \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \mathcal{P}_R \nabla f_{i_t}(w^t), \Delta \right\rangle \\
&\geq -\mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_R \nabla f_{i_t}(w^t) \right\|_2 \|\Delta\|_2.
\end{aligned} \tag{53}$$

The right-hand side can be lower bounded by applying inequality (11), which yields $\|\mathcal{P}_R \nabla f_{i_t}(w^t)\|_2 \leq \|\mathcal{P}_\Gamma \nabla f_{i_t}(w^t)\|_2 + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\mathcal{P}_{\Gamma^c} \nabla f_{i_t}(w^t)\|_2$. We now apply this observation to the above inequality. Denote $z \triangleq -\frac{\mathcal{P}_\Gamma \nabla f_{i_t}(w^t)}{\|\mathcal{P}_\Gamma \nabla f_{i_t}(w^t)\|_2} \|\Delta\|_2$ and $x \triangleq \frac{\mathcal{P}_{\Gamma^c} \nabla f_{i_t}(w^t)}{\|\mathcal{P}_{\Gamma^c} \nabla f_{i_t}(w^t)\|_2} \|\Delta\|_2$, we have

$$\begin{aligned}
& -\mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_R \nabla f_{i_t}(w^t) \right\|_2 \|\Delta\|_2 \\
& \geq -\mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_\Gamma \nabla f_{i_t}(w^t) \right\|_2 \|\Delta\|_2 - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_{\Gamma^c} \nabla f_{i_t}(w^t) \right\|_2 \|\Delta\|_2 \\
& = \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \mathcal{P}_\Gamma \nabla f_{i_t}(w^t), z \right\rangle - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \mathcal{P}_{\Gamma^c} \nabla f_{i_t}(w^t), x \right\rangle \\
& = \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), z \right\rangle - \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x \right\rangle \\
& = \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), z - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x \right\rangle,
\end{aligned} \tag{54}$$

where the second equality follows from $\text{supp}_{\mathcal{D}}(z) = \Gamma$ and $\langle \mathcal{P}_\Gamma r, z \rangle = \langle r, \mathcal{P}_\Gamma z \rangle = \langle r, z \rangle$. Denote $y \triangleq z - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x$ and combine (53) and (54) to arrive at

$$F(w^*) - F(w^t) - \frac{\rho_{4k}^-}{2} \|\Delta\|_2^2 \geq \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), y \right\rangle. \tag{55}$$

We now use the \mathcal{D} -RSS property to lower bound the right-hand side of the above inequality. Recall that from the definition of \mathcal{D} -RSS, we can show that

$$\langle \nabla f_{i_t}(w^t), y \rangle \geq f_{i_t}(w^t + y) - f_{i_t}(w^t) - \frac{\rho_{4k}^+(i_t)}{2} \|y\|_2^2.$$

Multiply both sides with $\frac{1}{Mp(i_t)}$ and take the expectation with respect to the index i_t and recall that $\mathbb{E}_{i_t} \frac{1}{Mp(i_t)} f_{i_t}(w^t) = F(w^t)$, we have

$$\mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), y \right\rangle \geq \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} f_{i_t}(w^t + y) - F(w^t) - \frac{1}{2} \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2.$$

Combining with inequality (55) and removing the common terms yields

$$\begin{aligned}
\frac{1}{2} \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2 - \frac{\rho_{4k}^-}{2} \|\Delta\|_2^2 & \geq \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} f_{i_t}(w^t + y) - F(w^*) \\
& = \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} (f_{i_t}(w^t + y) - f_{i_t}(w^*))
\end{aligned}$$

where the equality follows from $F(w^*) = \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} f_{i_t}(w^*)$. Applying the \mathcal{D} -RSC one more time to

the right-hand side and then taking the expectation, we get

$$\begin{aligned}
& \frac{1}{2} \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2 - \frac{\rho_{4k}^-}{2} \|\Delta\|_2^2 \\
& \geq \frac{\rho_{4k}^-}{2} \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} \|w^t + y - w^*\|_2^2 + \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^*), w^t + y - w^* \right\rangle \\
& = \frac{\rho_{4k}^-}{2} \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} \|\Delta - y\|_2^2 + \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \mathcal{P}_{\Gamma \cup R} \nabla f_{i_t}(w^*), y - \Delta \right\rangle \\
& \geq \frac{\rho_{4k}^-}{2} \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} \|\Delta - y\|_2^2 - \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} \|\mathcal{P}_{\Gamma \cup R} \nabla f_{i_t}(w^*)\|_2 \|\Delta - y\|_2 \\
& \geq \frac{\rho_{4k}^-}{2 \max_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\Delta - y\|_2^2 - \frac{\max_{i_t} \|\mathcal{P}_{\Gamma \cup R} \nabla f_{i_t}(w^*)\|_2}{\min_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\Delta - y\|_2 \\
& \geq \frac{\rho_{4k}^-}{2 \max_{i_t} Mp(i_t)} (\mathbb{E}_{i_t} \|\Delta - y\|_2)^2 - \frac{\max_{i_t} \|\mathcal{P}_{\Gamma \cup R} \nabla f_{i_t}(w^*)\|_2}{\min_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\Delta - y\|_2.
\end{aligned} \tag{56}$$

Solving the quadratic inequality $au^2 - 2bu - c \leq 0$ with $u = \mathbb{E}_{i_t} \|\Delta - y\|_2$, $a = \frac{\rho_{4k}^-}{\max_{i_t} Mp(i_t)}$, $b = \frac{\max_{i_t} \|\mathcal{P}_{\Gamma \cup R} \nabla f_{i_t}(w^*)\|_2}{\min_{i_t} Mp(i_t)}$, and $c = \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2 - \rho_{4k}^- \|\Delta\|_2^2$, we obtain

$$\mathbb{E}_{i_t} \|\Delta - y\|_2 \leq \sqrt{\frac{c}{a}} + \frac{2b}{a}. \tag{57}$$

Now plugging the definition of y we can obtain the lower bound of the left-hand side. We have,

$$\begin{aligned}
\|\Delta - y\|_2 & = \left\| \Delta - z + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x \right\|_2 \\
& \geq \|\Delta - z\|_2 - \left\| \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x \right\|_2 \\
& = \left\| \Delta + \frac{\|\Delta\|_2}{\|\mathcal{P}_\Gamma \nabla f_{i_t}(w^t)\|_2} \mathcal{P}_\Gamma \nabla f_{i_t}(w^t) \right\|_2 - \left\| \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \frac{\nabla f_{i_t}(w^t)}{\|\nabla f_{i_t}(w^t)\|_2} \|\Delta\|_2 \right\|_2 \\
& \geq \|\Delta - \mathcal{P}_\Gamma \Delta\|_2 - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\Delta\|_2,
\end{aligned}$$

where the first inequality follows from the triangular argument; the last inequality follows from the observation that for any vector v , $\|\Delta - \mathcal{P}_\Gamma v\|_2 \geq \|\Delta - \mathcal{P}_\Gamma \Delta\|_2$. Here, $v = -\frac{\|\Delta\|_2}{\|\mathcal{P}_\Gamma \nabla f_{i_t}(w^t)\|_2} \nabla f_{i_t}(w^t)$.

Therefore,

$$\mathbb{E}_{i_t} \|\Delta - y\|_2 \geq \mathbb{E}_{i_t} \|\Delta - \mathcal{P}_\Gamma \Delta\|_2 - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\Delta\|_2.$$

Plugging this inequality into (57), we get

$$\mathbb{E}_{i_t} \|\Delta - \mathcal{P}_\Gamma \Delta\|_2 \leq \sqrt{\frac{c}{a}} + \frac{2b}{a} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\Delta\|_2. \tag{58}$$

The last step is to substitute values of a , b , and c defined above into this inequality. From the definition of y together with the observation that x is orthogonal with z , we have

$$\begin{aligned}\|y\|_2^2 &= \|z\|_2^2 + \frac{\eta_1^2 - 1}{\eta_1^2} \|x\|_2^2 \\ &= \left\| \frac{\mathcal{P}_\Gamma \nabla f_{i_t}(w^t)}{\|\mathcal{P}_\Gamma \nabla f_{i_t}(w^t)\|_2} \|\Delta\|_2 \right\|_2^2 + \frac{\eta_1^2 - 1}{\eta_1^2} \left\| \frac{\mathcal{P}_{\Gamma^c} \nabla f_{i_t}(w^t)}{\|\mathcal{P}_{\Gamma^c} \nabla f_{i_t}(w^t)\|_2} \|\Delta\|_2 \right\|_2^2 \\ &= \|\Delta\|_2^2 + \frac{\eta_1^2 - 1}{\eta_1^2} \|\Delta\|_2^2 = \frac{2\eta_1^2 - 1}{\eta_1^2} \|\Delta\|_2^2.\end{aligned}$$

Thus, the quantity c defined above is bounded by

$$\begin{aligned}c &= \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \frac{2\eta_1^2 - 1}{\eta_1^2} \|\Delta\|_2^2 - \rho_{4k}^- \|\Delta\|_2^2 \\ &\leq \max_{i_t} \rho_{4k}^+(i_t) \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} \frac{2\eta_1^2 - 1}{\eta_1^2} \|\Delta\|_2^2 - \rho_{4k}^- \|\Delta\|_2^2 \\ &= \left(\frac{2\eta_1^2 - 1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^- \right) \|\Delta\|_2^2.\end{aligned}$$

Now combine this inequality with (58) and plug values of a and b , we obtain

$$\begin{aligned}\mathbb{E}_{i_t} \|\Delta - \mathcal{P}_\Gamma \Delta\|_2 &\leq \sqrt{\frac{c}{a}} + \frac{2b}{a} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\Delta\|_2 \\ &\leq \max_{i_t} \sqrt{Mp(i_t)} \sqrt{\frac{2\eta_1^2 - 1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-} \|\Delta\|_2 \\ &\quad + \frac{2 \max_{i_t} p(i_t)}{\rho_{4k}^- \min_{i_t} p(i_t)} \max_{|\Omega| \leq 4k, i_t \in [M]} \|\mathcal{P}_\Omega \nabla f_{i_t}(w^*)\|_2 + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\Delta\|_2 \\ &= \left(\max_{i_t} \sqrt{Mp(i_t)} \sqrt{\frac{2\eta_1^2 - 1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \right) \|\Delta\|_2 \\ &\quad + \frac{2 \max_{i_t} p(i_t)}{\rho_{4k}^- \min_{i_t} p(i_t)} \max_{|\Omega| \leq 4k, i_t \in [M]} \|\mathcal{P}_\Omega \nabla f_{i_t}(w^*)\|_2\end{aligned}$$

The proof follows by combining this inequality with (52). \square

8.4 Proof of Theorem 3

Proof. At the t -th iteration, denote $g_{i_t}(w) \triangleq \nabla f_{i_t}(w) + e^t$. The proof of Theorem 3 is essentially the same as that of Theorem 1 with $\nabla f_{i_t}(w^t)$ and $\nabla f_{i_t}(w^*)$ replaced by $g_{i_t}(w^t)$ and $g_{i_t}(w^*)$, respectively. Following the same proof as in Theorem 1, we arrive at

$$\begin{aligned}&\mathbb{E}_{i_t|I_{t-1}} \|w^{t+1} - w^*\|_2 \\ &\leq 2 \left(\mathbb{E}_{i_t|I_{t-1}} \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} \mathcal{P}_\Omega (g_{i_t}(w^t) - g_{i_t}(w^*)) \right\|_2 + \mathbb{E}_{i_t|I_{t-1}} \left\| \frac{\gamma}{Mp(i_t)} \mathcal{P}_\Omega g_{i_t}(w^*) \right\|_2 \right) \\ &\quad + \sqrt{\eta^2 - 1} \left(\mathbb{E}_{i_t|I_{t-1}} \left\| w^t - w^* - \frac{\gamma}{Mp(i_t)} (g_{i_t}(w^t) - g_{i_t}(w^*)) \right\|_2 + \mathbb{E}_{i_t|I_{t-1}} \left\| \frac{\gamma}{Mp(i_t)} g_{i_t}(w^*) \right\|_2 \right).\end{aligned}$$

Notice that $g_{i_t}(w^t) - g_{i_t}(w^*) = \nabla f_{i_t}(w^t) - \nabla f_{i_t}(w^*)$, we can apply the inequality (42) of Corollary 10 for the first term and (43) for the third term of the summation to obtain

$$\begin{aligned}
& \mathbb{E}_{i_t|I_{t-1}} \|w^{t+1} - w^*\|_2 \\
& \leq 2\sqrt{(1 - (2\gamma - \gamma^2\alpha_{3k})\rho_{3k}^-)} \|w^t - w^*\|_2 + 2\frac{\gamma}{\min_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\mathcal{P}_\Omega \nabla g_{i_t}(w^*)\|_2 \\
& \quad + \sqrt{\eta^2 - 1} \sqrt{1 + \gamma^2\alpha_{3k}\bar{\rho}_{3k}^+ - 2\gamma\rho_{3k}^-} \|w^t - w^*\|_2 + \sqrt{\eta^2 - 1} \frac{\gamma}{\min_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\nabla g_{i_t}(w^*)\|_2 \\
& = \left(2\sqrt{(1 - (2\gamma - \gamma^2\alpha_{3k})\rho_{3k}^-)} + \sqrt{(\eta^2 - 1)(1 + \gamma^2\alpha_{3k}\bar{\rho}_{3k}^+ - 2\gamma\rho_{3k}^-)} \right) \|w^t - w^*\|_2 \\
& \quad + \frac{\gamma}{\min_{i_t} Mp(i_t)} \left(2\mathbb{E}_{i_t} \|\mathcal{P}_\Omega \nabla g_{i_t}(w^*)\|_2 + \sqrt{\eta^2 - 1} \mathbb{E}_{i_t} \|\nabla g_{i_t}(w^*)\|_2 \right) \\
& \leq \kappa \|w^t - w^*\|_2 + (\sigma_{w^*} + \sigma_{e^t}),
\end{aligned}$$

where κ and σ_{w^*} are defined in (15) and (16) and

$$\sigma_{e^t} \triangleq \frac{\gamma}{\min_i Mp(i)} \left(2 \max_{|\Omega| \leq 3k} \|\mathcal{P}_\Omega e^t\|_2 + \sqrt{\eta^2 - 1} \|e^t\|_2 \right).$$

Taking the expectation on both sides with respect to I_{t-1} yields

$$\mathbb{E}_{I_t} \|w^{t+1} - w^*\|_2 \leq \kappa \mathbb{E}_{I_{t-1}} \|w^t - w^*\|_2 + (\sigma_{w^*} + \sigma_e),$$

where $\sigma_e = \max_{j \in [t]} \sigma_{e^j}$. Applying this result recursively completes the proof. \square

8.5 Proof of Theorem 4

Proof. The analysis of Theorem 4 follows closely to that of Theorem 2. In particular, we will apply Lemmas 1, 2, and the following lemma to obtain the proof.

Lemma 4. Denote $\hat{\Gamma}$ as the set obtained from the t -th iteration. Then,

$$\mathbb{E}_{i_t} \|\mathcal{P}_{\hat{\Gamma}^c}(b^t - w^*)\|_2 \leq \left(\max_i \sqrt{Mp(i)} \sqrt{\frac{2\eta_1^2 - 1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \right) \|w^t - w^*\|_2 + \sigma_2, \quad (59)$$

where

$$\sigma_2 \triangleq \frac{2 \max_{i_t} p(i_t)}{\rho_{4k}^- \min_{i_t} p(i_t)} \left(\max_{|\Omega| \leq 4k, i_t \in [M]} \|\mathcal{P}_\Omega \nabla f_{i_t}(w^*)\|_2 + \max_t \|e^t\|_2 \right).$$

Given this lemma, the proof is exactly the same as that of Theorem 2. Now, we proceed to prove Lemma 4. Denote $\Delta = w^* - w^t$ and $g_{i_t}(w) \triangleq \nabla f_{i_t}(w) + e^t$. Similar to the analysis of Lemma 3, we start by applying the \mathcal{D} -RSC,

$$\begin{aligned}
& F(w^*) - F(w^t) - \frac{\rho_{4k}^-}{2} \|w^* - w^t\|_2^2 \\
& \geq \langle \nabla F(w^t), w^* - w^t \rangle \\
& = \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), \Delta \right\rangle \\
& = \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \mathcal{P}_R(g_{i_t}(w^t)) - e^t, \Delta \right\rangle \\
& \geq -\mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_R g_{i_t}(w^t) \right\|_2 \|\Delta\|_2 - \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \mathcal{P}_R e^t, \Delta \right\rangle.
\end{aligned} \quad (60)$$

Again, applying inequality (11) allows us to write $\|\mathcal{P}_R g_{i_t}(w^t)\|_2 \leq \|\mathcal{P}_\Gamma g_{i_t}(w^t)\|_2 + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\mathcal{P}_{\Gamma^c} g_{i_t}(w^t)\|_2$. We now apply this observation to the above inequality. Denote $z \triangleq -\frac{\mathcal{P}_\Gamma g_{i_t}(w^t)}{\|\mathcal{P}_\Gamma g_{i_t}(w^t)\|_2} \|\Delta\|_2$ and $x \triangleq \frac{\mathcal{P}_{\Gamma^c} g_{i_t}(w^t)}{\|\mathcal{P}_{\Gamma^c} g_{i_t}(w^t)\|_2} \|\Delta\|_2$ and follow the same procedure in formula (54) with $\nabla f_{i_t}(w^t)$ replaced by $g_{i_t}(w^t)$, we arrive at

$$\begin{aligned}
& -\mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_R g_{i_t}(w^t) \right\|_2 \|\Delta\|_2 \\
& \geq -\mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_\Gamma g_{i_t}(w^t) \right\|_2 \|\Delta\|_2 - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \mathbb{E}_{i_t} \left\| \frac{1}{Mp(i_t)} \mathcal{P}_{\Gamma^c} g_{i_t}(w^t) \right\|_2 \|\Delta\|_2 \\
& = \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} g_{i_t}(w^t), z - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x \right\rangle \\
& = \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), z - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x \right\rangle + \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} e^t, z - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x \right\rangle.
\end{aligned} \tag{61}$$

Denote $y \triangleq z - \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} x$ and combine (60) and (61), we get

$$\begin{aligned}
F(w^*) - F(w^t) - \frac{\rho_{4k}^-}{2} \|\Delta\|_2^2 & \geq \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), y \right\rangle - \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} e^t, \Delta - y \right\rangle \\
& \geq \mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), y \right\rangle - \frac{1}{\min_{i_t} Mp(i_t)} \|e^t\|_2 \mathbb{E}_{i_t} \|\Delta - y\|_2,
\end{aligned} \tag{62}$$

where the last argument follows from Cauchy-Schwarz inequality. We now use the \mathcal{D} -RSS property to lower bound the right-hand side of the above inequality. Recall that from the definition of \mathcal{D} -RSS, we can show that

$$\langle \nabla f_{i_t}(w^t), y \rangle \geq f_{i_t}(w^t + y) - f_{i_t}(w^t) - \frac{\rho_{4k}^+(i_t)}{2} \|y\|_2^2.$$

Multiply both sides with $\frac{1}{Mp(i_t)}$ and take the expectation on both sides with respect to the index i_t and recall that $\mathbb{E}_{i_t} \frac{1}{Mp(i_t)} f_{i_t}(w^t) = F(w^t)$, we have

$$\mathbb{E}_{i_t} \left\langle \frac{1}{Mp(i_t)} \nabla f_{i_t}(w^t), y \right\rangle \geq \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} f_{i_t}(w^t + y) - F(w^t) - \frac{1}{2} \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2.$$

Combining with inequality (62) and removing the common terms yields

$$\begin{aligned}
& \frac{1}{2} \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2 - \frac{\rho_{4k}^-}{2} \|\Delta\|_2^2 \\
& \geq \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} f_{i_t}(w^t + y) - F(w^*) - \frac{1}{\min_{i_t} Mp(i_t)} \|e^t\|_2 \mathbb{E}_{i_t} \|\Delta - y\|_2 \\
& = \mathbb{E}_{i_t} \frac{1}{Mp(i_t)} (f_{i_t}(w^t + y) - f_{i_t}(w^*)) - \frac{1}{\min_{i_t} Mp(i_t)} \|e^t\|_2 \mathbb{E}_{i_t} \|\Delta - y\|_2.
\end{aligned}$$

Apply the \mathcal{D} -RSC one more time to the right-hand side and follow a similar procedure as (56) with

an additional term involving the gradient noise, we get

$$\begin{aligned}
& \frac{1}{2} \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2 - \frac{\rho_{4k}^-}{2} \|\Delta\|_2^2 \\
& \geq \frac{\rho_{4k}^-}{2 \max_{i_t} Mp(i_t)} (\mathbb{E}_{i_t} \|\Delta - y\|_2)^2 - \frac{\max_{i_t} \|\mathcal{P}_{\Gamma \cup R} \nabla f_{i_t}(w^*)\|_2}{\min_{i_t} Mp(i_t)} \mathbb{E}_{i_t} \|\Delta - y\|_2 \quad (63) \\
& \quad - \frac{1}{\min_{i_t} Mp(i_t)} \|e^t\|_2 \mathbb{E}_{i_t} \|\Delta - y\|_2.
\end{aligned}$$

Solving the quadratic inequality $au^2 - 2bu - c \leq 0$ where $u = \mathbb{E}_{i_t} \|\Delta - y\|_2$, $a = \frac{\rho_{4k}^-}{\max_{i_t} Mp(i_t)}$, $b = \frac{\max_{i_t} \|\mathcal{P}_{\Gamma \cup R} \nabla f_{i_t}(w^*)\|_2}{\min_{i_t} Mp(i_t)} + \frac{1}{\min_{i_t} Mp(i_t)} \|e^t\|_2$, and $c = \mathbb{E}_{i_t} \frac{\rho_{4k}^+(i_t)}{Mp(i_t)} \|y\|_2^2 - \rho_{4k}^- \|\Delta\|_2^2$, we obtain

$$\mathbb{E}_{i_t} \|\Delta - y\|_2 \leq \sqrt{\frac{c}{a}} + \frac{2b}{a}. \quad (64)$$

Following the same steps after inequality (57), we arrive at

$$\begin{aligned}
\mathbb{E}_{i_t} \|\Delta - \mathcal{P}_\Gamma \Delta\|_2 & \leq \sqrt{\frac{c}{a}} + \frac{2b}{a} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\Delta\|_2 \\
& = \max_i \sqrt{Mp(i)} \sqrt{\frac{\frac{2\eta_1^2 - 1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-}{\rho_{4k}^-}} \|\Delta\|_2 \\
& \quad + \frac{2 \max_i p(i)}{\rho_{4k}^- \min_i p(i)} \left(\max_{|\Omega| \leq 4k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2 + \|e^t\|_2 \right) + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \|\Delta\|_2 \\
& = \left(\max_i \sqrt{Mp(i)} \sqrt{\frac{\frac{2\eta_1^2 - 1}{\eta_1^2} \rho_{4k}^+ - \rho_{4k}^-}{\rho_{4k}^-}} + \frac{\sqrt{\eta_1^2 - 1}}{\eta_1} \right) \|\Delta\|_2 \\
& \quad + \frac{2 \max_i p(i)}{\rho_{4k}^- \min_i p(i)} \left(\max_{|\Omega| \leq 4k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2 + \|e^t\|_2 \right).
\end{aligned}$$

The proof follows by combining this inequality with (52). \square

8.6 Proof of Theorem 5

Proof. The proof follows exactly the same steps as that of Theorem 2 in Section 8.3. The only difference is Lemma 2, which is now replaced by the following lemma.

Lemma 5. Denote $\widehat{\Gamma}$ as the set obtained from the t -th iteration and i as the index selected randomly from $[M]$ with probability $p(i)$. We have,

$$\mathbb{E}_{I_t} \|b^t - w^*\|_2 \leq \sqrt{\frac{\alpha_{4k}}{\rho_{4k}^-}} \mathbb{E}_{I_t} \|P_{\widehat{\Gamma}^c}(b^t - w^*)\|_2 + \sigma_1 + \epsilon^t,$$

where $\alpha_k = \max_i \frac{\rho_k^+(i)}{Mp(i)}$ and $\sigma_1 \triangleq \frac{3}{\rho_{4k}^-} \frac{1}{\min_i Mp(i)} \max_{|\Omega| \leq 3k, i \in [M]} \|\mathcal{P}_\Omega \nabla f_i(w^*)\|_2$.

From the triangular inequality,

$$\mathbb{E}_{i_t} \|b^t - w^*\|_2 \leq \mathbb{E}_{i_t} \|b_{\text{opt}}^t - w^*\|_2 + \mathbb{E}_{i_t} \|b^t - b_{\text{opt}}^t\|_2 \leq \mathbb{E}_{i_t} \|b_{\text{opt}}^t - w^*\|_2 + \epsilon^t.$$

Applying Lemma 2 to get upper bound for $\mathbb{E}_{i_t} \|b_{\text{opt}}^t - w^*\|_2$ will complete the proof of this lemma. \square

References

- [1] A. Agarwal, S. Negahban, and M. J. Wainwright. Fast global convergence of gradient methods for high-dimensional statistical recovery. *Ann. Statist.*, 40(5):2452–2482, 2012.
- [2] F. Bach and E. Moulines. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. *Ad. Neural Inf. Process. Sys. (NIPS)*, 2011.
- [3] S. Bahmani, B. Raj, and P. T. Boufounos. Greedy sparsity-constrained optimization. *Journ. Mach. Learn. Res.*, 14:807–841, Mar. 2013.
- [4] D. P. Bertsekas, A. Nedic, and A. E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, 2003.
- [5] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comp. Harm. Anal.*, 27(3):265–274, 2009.
- [6] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [7] E. J. Candès. Compressive sampling. In *Proceed. Int. Congress Math.*, volume 3, pages 1433–1452. Madrid, Spain, 2006.
- [8] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory*, 52(2):5406–5425, Feb. 2006.
- [9] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223, Aug. 2006.
- [10] E. J. Candès and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Trans. Inf. Theory*, 57(4):2342–2359, Apr. 2011.
- [11] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inf. Theory*, 51(12):4203–4215, Dec. 2005.
- [12] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. Willsky. The convex geometry of linear inverse problems. *Found. Comp. Math.*, 12(6):805–849, Dec. 2012.
- [13] Compressed sensing webpage. <http://www.dsp.ece.rice.edu/cs/>.
- [14] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.*, 57:1413–1457, 2004.
- [15] M. Davenport, D. Needell, and M. Wakin. Signal space CoSaMP for sparse recovery with redundant dictionaries. *IEEE Trans. Inf. Theory*, 59(10):6820 – 6829, 2012.

- [16] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, April 2006.
- [17] P. Drineas, M. W. Mahoney, S. Muthukrishnan, and T. Sarlós. Faster least squares approximation. *Numerische Mathematik*, 117(2):219–249, Feb. 2011.
- [18] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright. Gradient projection for sparse reconstruction: Application to Compressed Sensing and other inverse problems. *IEEE J. Sel. Topics Sign. Process.: Spec. Issue Convex Opt. Methods. Sign. Process.*, 1(4):586–598, 2007.
- [19] S. Foucart. Sparse recovery algorithms: sufficient conditions in terms of restricted isometry constants. *Approx. Theory*, 13:65–77, 2012.
- [20] R. Giryes and M. Elad. RIP-based near-oracle performance guarantees for SP, CoSaMP, and IHT. *IEEE Trans. Sign. Process.*, 60(3):1465–1468, 2012.
- [21] R. Giryes, S. Nam, M. Elad, R. Gribonval, and M. E. J. Davies. Greedy-like algorithms for the cosparsity analysis model. *Linear Algebra Appl.*, 441:22–60, 2014.
- [22] R. Giryes and D. Needell. Greedy signal space methods for incoherence and beyond. Submitted, 2013.
- [23] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [24] P. Jain, R. Mega, and I. Dhillon. Guaranteed rank minimization via Singular Value Projection. In *Ad. Neural Infor. Proc. Sys. (NIPS)*, pages 937–945, Vancouver, BC, Canada, Dec. 2010.
- [25] A. Jalali, C. C. Johnson, and P. Ravikumar. On learning discrete graphical models using greedy methods. In *Ad. Neural Infor. Proc. Sys. (NIPS)*, Granada, Spain, Dec. 2011.
- [26] S. Kaczmarz. Angenäherte auflösung von systemen linearer gleichungen. *Bull. Int. Acad. Polon. Sci. Lett. Ser. A*, pages 335–357, 1937.
- [27] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky. An interior-point method for large-scale L1-regularized least squares. *IEEE J. Sel. Topics Sign. Process.*, 1(4):606–617, 2007.
- [28] D. Needell, N. Srebro, and R. Ward. Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm. Submitted, 2013.
- [29] D. Needell and J. A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Appl. Comp. Harm. Anal.*, 26:301–321, 2008.
- [30] D. Needell and R. Vershynin. Signal recovery from inaccurate and incomplete measurements via regularized orthogonal matching pursuit. *IEEE J. Sel. Topics Sign. Process.*, 4(2):310–316, Apr. 2010.
- [31] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. In *Ad. Neural Inf. Process. Sys. (NIPS)*, Vancouver, BC, Canada, Dec. 2009.
- [32] Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM J. Optimiz.*, 22(2):341–362, 2012.

- [33] N. H. Nguyen, S. Chin, and T. D. Tran. A unified iterative greedy algorithm for sparsity-constrained optimization. 2013. Submitted.
- [34] R. Richtárik and M. Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Math. Program.*, 144(2):1–38, 2014.
- [35] S. Shalev-Shwartz, A. Gonen, and O. Shamir. Large-scale convex minimization with a low-rank constraint. In *Proc. 29th Inter. Conf. Mach. Learn. (ICML)*, pages 329–336, Bellevue, WA, USA, June 2011.
- [36] S. Shalev-Shwartz, N. Srebro, , and T. Zhang. Trading accuracy for sparsity in optimization problems with sparsity constraints. *SIAM Journ. Opt.*, 20(6):2807–2832, 2010.
- [37] James C Spall. *Introduction to stochastic search and optimization: estimation, simulation, and control*, volume 65. John Wiley & Sons, 2005.
- [38] T. Strohmer and R. Vershynin. A randomized Kaczmarz algorithm with exponential convergence. *J. Fourier Anal. Appl.*, 15(2):262–278, 2009.
- [39] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996.
- [40] J. A. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. Inf. Theory*, 50(10):2231–2242, 2004.
- [41] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. Roy. Statist. Soc. Ser. B*, 68(1):49–67, 2006.
- [42] X.-T. Yuan, P. Li, and T. Zhang. Gradient hard thresholding pursuit for sparsity-constrained optimization. In *Proc. 32nd Inter. Conf. Mach. Learn. (ICML)*, Beijing, China, Dec. 2014.
- [43] X.-T. Yuan and S. Yan. Forward basis selection for sparse approximation over dictionary. In *Proc. 15th Inter. Conf. Art. Intel. Statist. (AISTATS)*, pages 1377–1388, La Palma, Canary Islands, Apr. 2012.
- [44] T. Zhang. Adaptive forward-backward greedy algorithm for learning sparse representations. *IEEE Trans. Inf. Theory*, 57(7):4689–4708, Jul. 2011.
- [45] Tong Zhang. Sparse recovery with orthogonal matching pursuit under RIP. *IEEE Trans. Inf. Theory*, 57(9):6215–6221, 2011.