

The Stochastic Proximal Distance Algorithm

by

Haoyu Jiang

Department of Statistical Science
Duke University

Date: _____

Approved:

Jason Xu, Supervisor

Simon Mak

Yuansi Chen

A thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science
in the Department of Statistical Science
in the Graduate School of
Duke University

2023

ABSTRACT

The Stochastic Proximal Distance Algorithm

by

Haoyu Jiang

Department of Statistical Science
Duke University

Date: _____

Approved:

Jason Xu, Supervisor

Simon Mak

Yuansi Chen

An abstract of a thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science
in the Department of Statistical Science
in the Graduate School of
Duke University

2023

Copyright © 2023 by Haoyu Jiang
All rights reserved

Abstract

Stochastic versions of proximal methods have gained much attention in statistics and machine learning. These algorithms tend to admit simple, scalable forms, and enjoy numerical stability via implicit updates. In this work, we propose and analyze a stochastic version of the recently proposed proximal distance algorithm, a class of iterative optimization methods that recover a desired constrained estimation problem as a penalty parameter $\rho \rightarrow \infty$. By uncovering connections to related stochastic proximal methods and interpreting the penalty parameter as the learning rate, we justify heuristics used in practical manifestations of the proximal distance method, establishing their convergence guarantees for the first time. Moreover, we extend recent theoretical devices to establish finite error bounds and a complete characterization of convergence rates regimes. We validate our analysis via a thorough empirical study, also showing that unsurprisingly, the proposed method outpaces batch versions on popular learning tasks.

Contents

Abstract	iv
List of Figures	vii
1 Introduction	1
2 Background and Motivation	4
3 A Stochastic Proximal Distance Algorithm	7
4 Theoretical Analysis	12
5 Empirical Study	17
5.1 Synthetic Data	17
5.2 Real Data	21
6 Conclusion	26
A Proofs of Theoretical Results	27
A.1 Proof of Proposition 1	27
A.2 Proof of Proposition 2	27
A.3 Proof of Theorem 1	30
A.4 Proof of Theorem 2	33
A.5 Proof of Theorem 3	37
B Derivation of Proximal Distance Updates	40
B.1 Sparsity/Unit Ball Constrained Linear Regression	40
B.2 Sparsity/Unit Ball Constrained Logistic Regression	41
C Data Generation and Additional Empirical Results	43

List of Figures

5.1	Finite error at different γ in different settings under the unit ball constraint (top panel) and sparsity set constraint (bottom panel)	18
5.2	Runtime as sample size and batch size vary.	20
5.3	Box plots of runtime, iteration numbers, accuracy and AUC for all the algorithms.	23
5.4	ROC curves for all the algorithms. The left plot contains the results for batch size 200 and the right one is for for batch size 500.	24
5.5	Results of the stochastic proximal distance algorithm and the projected SGD with random initial guesses.	25
C.1	Finite error in terms of the parameter at different γ under the unit ball constraint (top panel) and sparsity constraint (bottom panel) . .	44
C.2	Finite error at different ρ_1 in different settings under the unit ball constraint	45
C.3	Box plots of error in various settings under the unit ball constraint. .	46
C.4	Box plots of error in various settings under the sparsity constraint. For reference to give a sense of scale, $\ \theta_*\ \approx 12$ here.	47

Chapter 1

Introduction

Optimization of an objective function subject to constraints or regularization is ubiquitous in statistics and machine learning. Consider the canonical problem setting

$$\min_{\theta} F(\theta) = \frac{1}{n} \sum_{i=1}^n f(\theta; z_i) \quad \text{s. t. } \theta \in C \quad (1.1)$$

where $\{z_i\}_{i=1}^n$ is the observed data, θ is the model parameter, f is typically a measure of fit, and C is a set imposing desired structure on the solution. Toward solving constrained and regularized objectives, there are many variants of gradient methods popular in the literature. We survey a number of these in Section 2, which often employing projection or proximal steps to handle constraints and non-smooth penalties. Here we focus attention on a recent method by [LK15] that is appealing in providing a broad framework to handle many constrained problems gracefully. Called the *proximal distance algorithm*, their approach transforms (1.1) to the unconstrained problem

$$\min_{\theta} F(\theta) + \frac{\rho}{2} \text{dist}(\theta, C)^2, \quad (1.2)$$

where $\text{dist}(\theta, C) = \inf_{\theta' \in C} \|\theta - \theta'\|_2$ denotes the Euclidean distance between θ and the constraint set C . Using the idea of distance majorization [CZL14], this reformulation can be solved via iterative algorithms from the perspective of Majorization-Minimization, or MM [HL04, Mai15] as long as the projection operators are practical to evaluate. This is the case for many common constraints including sparsity, rank, and shape restrictions; a broad variety of examples are considered in [XCL17, KZL19, LPZL22].

This proximal distance method can be viewed as a smooth version of the classical penalty method due to [Cou43], but under a squared penalty as in (1.2), the unconstrained solutions recover the constrained solution only in the limit as $\rho \rightarrow \infty$ [BB73]. Early attempts to set ρ at a large value encountered slow progress, as the reciprocal $1/\rho$ plays the

role of a step-size in the algorithm, discussed further below. This is supported by recent analyses by [KZL19, LPZL22] revealing a convergence rate in convex cases of $O(\rho k^{-1})$, where k is the iteration counter. In light of this observation, researchers have suggested gradually increasing ρ within the iterative algorithm. [KZL19] suggests setting an initial penalty value $\rho = \rho_0$, rate $\gamma > 1$ and frequency m , and then updating $\rho_k = \rho_0 \gamma^{\lfloor \frac{k}{m} \rfloor}$, which increases roughly exponentially. While this is consistent with the intuition of approaching the constrained solution, the traditional MM geometry used to derive the algorithm—along with its corresponding guarantees— does not apply to heuristics that entail a sequence of *changing* objective functions (1.2) indexed by the increasing sequence ρ_k . These heuristics for increasing ρ play an important role in practice so that iterates eventually respect the constraint when the penalty becomes large, but allow for large enough data-driven steps in the earlier stages of the algorithm.

Fortunately, there is a large literature on related stochastic proximal methods, which are increasingly prominent in statistical and machine learning problems involving large data sets [THA20, LLW22]. In this paper, we draw new connections between the proximal distance algorithm and these studies, leveraging techniques for studying their convergence toward rigorously resolving the open questions regarding the convergence of the proximal distance algorithm under various ρ_k schedules. Our point of departure is to propose a stochastic version of the proximal distance algorithm. By evaluating the loss at only a subsample of the data, we break from the original geometry of majorization from which the method was derived. However, we take a different view by drawing analogies to implicit gradient methods. Noting the relation between the penalty parameter and the step size or *learning rate* in proximal/implicit algorithms, we establish connections to implicit SGD [TAR14a, TA17, LLW22], the proximal Robbins-Monro algorithm [THA20], and incremental proximal algorithms [Ber11]. We provide new theoretical analyses that reveal convergence rates in several regimes under various polynomial learning rate schedules under weaker assumptions than previous studies.

Our contributions bridge a theoretical gap in justifying heuristics commonly used in

practical implementations of the proximal distance algorithm. While the findings are largely theoretical in nature, the resulting stochastic algorithm naturally enables large-scale data to be analyzed that would be intractable under the full batch version of the method. The remainder of this manuscript is organized as follows: we begin with an overview and necessary background in Section 2. Next, Section 3 proposes our stochastic version of the proximal distance algorithm, and discusses connections and differences to related stochastic proximal methods. In Section 4, we present our theoretical analysis of the algorithm, establishing finite error bounds, revealing convergence rates and providing a rigorous justification for mechanisms to increase ρ . In Section 5, we conduct empirical study to validate our analysis and investigate behavior in practice for both convex and non-convex settings, while demonstrating computational gains over the full version. We conclude and discuss future directions in Section 6.

Chapter 2

Background and Motivation

Proximal Algorithms Many powerful iterative algorithms for optimization involve the proximal operator [BC⁺11, PB⁺14], defined as

$$\operatorname{prox}_f(y) = \arg \min_x f(x) + \frac{1}{2} \|x - y\|^2.$$

A fundamental example is the *proximal point algorithm* [Roc76, PB⁺14], which minimizes an objective F with successive proximal operations:

$$\theta_{k+1} = \operatorname{prox}_{\alpha F}(\theta_k) = \arg \min_{\theta} F(\theta) + \frac{1}{2\alpha} \|\theta - \theta_k\|^2.$$

Intuitively, this seeks to shrink toward the current iterate θ_k when minimizing F , where a parameter α determines the strength of the shrinkage. Methods such as *proximal gradient descent* [PB⁺14, LL15] consider composite objectives, alternating between proximal operations and gradient descent steps.

Of particular interest here is the proximal distance algorithm [LK15]. It extends the penalty method of [Cou43] which relaxes a constrained problem (1.1) to an unconstrained problem of the form

$$\min_{\theta} F(\theta) + \rho q(\theta),$$

where $q(\theta) \geq 0$ is a penalty that vanishes on the constraint set: $q(\theta) = 0$ for all $\theta \in C$. Solutions to this unconstrained problem θ_*^ρ are guaranteed to converge to the constrained solution as $\rho \rightarrow \infty$ [BB73]. The proximal distance algorithm combines this idea with distance majorization [CZL14], taking the penalty q to be the squared distance between θ and the constraint set C . A key algorithmic component is the set projection: note the squared distance between θ and C can also be expressed as

$$q(\theta) = \operatorname{dist}(\theta, C)^2 = \|\theta - P_C(\theta)\|^2,$$

where $P_C(\theta) = \arg \min_{\theta' \in C} \|\theta' - \theta\|_2$ denotes the projection of θ on C . [LK15] utilize this to convert the constrained problem (1.1) to an unconstrained form:

$$\min_{\theta} F(\theta) + \frac{\rho}{2} \|\theta - P_C(\theta)\|^2. \quad (2.1)$$

Though (2.1) is often not directly solvable, the proximal distance algorithm makes use of majorization-minimization (MM) to efficiently solve a sequence of simpler subproblems. In more detail, distance majorization [CZL14] implies that the following *surrogate function*

$$F(\theta) + \frac{\rho}{2} \|\theta - P_C(\theta_{k-1})\|^2 \quad (2.2)$$

majorizes the expression in (2.1). The principle of MM suggests then minimizing the surrogate (2.2):

$$\theta_k = \arg \min_{\theta} F(\theta) + \frac{\rho}{2} \|\theta - P_C(\theta_{k-1})\|^2 = \text{prox}_{\rho^{-1}F}[P_C(\theta_{k-1})]. \quad (2.3)$$

The resulting MM iteration consists of alternatively defining surrogates and minimizing them according to these updates. By ensuring a descent property, it is easy to show that the iterate sequence θ_k converges to a local minimum θ_*^ρ of (2.1). In convex settings, θ_*^ρ is the global minimizer.

In practice, convergence can be slow when fixing a large value for ρ , a pitfall past approaches seek to remedy by employing an increasing sequence of penalty parameters $\{\rho_k\}$. When ρ is small, the objective is dominated by the loss F which drives the MM updates; as ρ grows larger, the penalty plays a stronger role, guiding $\{\theta_k\}$ toward the constraint set. Heuristically increasing ρ at an effective rate plays a key role in the practical performance of the method, which we will investigate more closely below.

Stochastic Proximal Methods [RM51] proposed a stochastic approximation method for finding the roots of functions that lays the foundation for many stochastic optimization methods used in large scale statistical and machine learning problems. A well-known instance is stochastic gradient descent (SGD) [NJLS09, MB11, BM17], a method for smooth

unconstrained optimization that combines the Robbins-Monro algorithm with gradient descent in using subsamples toward approximate steps. It updates

$$\theta_{k+1} = \theta_k - \alpha_k \nabla f(\theta_k, z_{\xi_k}),$$

where z_{ξ_k} denotes a randomly drawn observation from the dataset and α_k is the sequence of learning rates. For convex cases, θ_k converges in probability to the solution whenever

$$\sum_{k=1}^{\infty} \alpha_k = \infty, \quad \sum_{k=1}^{\infty} \alpha_k^2 < \infty.$$

While SGD enables many learning tasks in the large-scale and online settings, it can be sensitive to initial learning rate, with small learning rates effecting slow convergence and large learning rates risking divergence [MB11, RB14]. Stochastic versions of proximal methods have been studied to improve numerical instability. The *stochastic proximal point algorithm* [Bia16, AD19], also called *implicit SGD* [TAR14a, TA17, LLW22], applies stochastic approximation to the proximal point algorithm. Subsampling one data point at each iteration,

$$\theta_{k+1} = \arg \min_{\theta} f(\theta; z_{\xi_k}) + \frac{1}{2\alpha_k} \|\theta - \theta_k\|^2 = \theta_k - \alpha_k \nabla f(\theta_{k+1}; z_{\xi_k}), \quad (2.4)$$

where the second equation is obtained by equating the gradient to zero. Note θ_{k+1} appears on both sides; for this reason, the method is referred to as *implicit*. By virtue of the proximal operation, this implicit iteration is more stable than (explicit) SGD with respect to the choice of initial step size in theory and practice [RB14, TAR14a]. [THA20] also consider an idealized implicit SGD in the setting where ξ_k itself is not observable. They incorporate a nested SGD as a subroutine and show that the method still inherits desirable stability properties. Similar *stochastic proximal gradient algorithms* [Nit14] subsample in the gradient descent step, and further employ Monte Carlo methods to approximate the gradient when it is analytically intractable [AFM17]. Finally, the incremental proximal method of [Ber11] is related to stochastic proximal gradient, but performs an incremental version of both the gradient descent step and the proximal step, and focuses on constrained problems.

Chapter 3

A Stochastic Proximal Distance

Algorithm

Motivated by the recent success of stochastic proximal methods in machine learning problems on massive datasets, it is natural to now consider a stochastic version of the proximal distance algorithm. Doing so allows the method to scale to problem sizes that were previously infeasible, and to apply to online data settings. Scaling the method is not the focus of our study, however; our analysis yields more surprising results on the convergence rates properties and properties of various penalty parameter schedules. This sheds light on the precise role of ρ_k not known in previous studies—even in the full batch setting, convergence guarantees were lacking.

For intuition, we first consider the case that F is convex. Taking the gradient of the surrogate function (2.2) and equating to 0 yields the stationarity equation

$$\nabla F(\theta) + \rho_k[\theta - P_C(\theta_{k-1})] = 0,$$

which furnishes the next iterate θ_k as the minimizer:

$$\theta_k = P_C(\theta_{k-1}) - \frac{1}{\rho_k} \nabla F(\theta_k). \tag{3.1}$$

We see from (3.1) that the update resembles an *implicit gradient* step but with an additional projection inside. Note that $1/\rho_k$ plays the role of the step size or *learning rate* in this analogy. Bridging to the intuition behind methods in Section 2, this leads to a natural stochastic variant of the proximal distance algorithm by approximating the implicit gradient using only a subset of the data. Doing so yields the following *stochastic proximal distance*

Algorithm 1 Stochastic proximal distance algorithm

input: data $\{z_i\}_{i=1}^n$, initial penalty ρ_1 , rate γ , batch size b , initial guess θ_0 , iteration cap K_{\max} , tolerance ϵ

output: $\hat{\theta} \approx \arg \min_{\theta \in C} \frac{1}{n} \sum_{i=1}^n f(\theta; z_i)$

Initialize $k = 0$

repeat

$k \leftarrow k + 1, \rho \leftarrow \rho_1 \times k^\gamma;$

$\mathcal{I} \leftarrow b$ samples taken from $\{1, \dots, n\}$ w/o replacement;

$\theta_k \leftarrow \text{prox}_{\rho^{-1}b^{-1} \sum_{i \in \mathcal{I}} f(\cdot; z_i)}[P_C(\theta_{k-1})]$

until $|F[P_C(\theta_k)] - F[P_C(\theta_{k-1})]| < \epsilon$ or $k = K_{\max}$

$\hat{\theta} \leftarrow P_C(\theta_k)$

iteration:

$$\begin{aligned} \theta_k &= P_C(\theta_{k-1}) - \frac{1}{\rho_k} \nabla f(\theta_k; z_{\xi_k}) \\ &= \arg \min_{\theta} f(\theta; z_{\xi_k}) + \frac{\rho_k}{2} \|\theta - P_C(\theta_{k-1})\|^2 \\ &= \text{prox}_{\rho_k^{-1} f(\cdot; z_{\xi_k})}[P_C(\theta_{k-1})], \end{aligned} \tag{3.2}$$

where the index ξ_k is drawn uniformly from $\{1, 2, \dots, n\}$. Here we have considered the subsample to be a single point for exposition, though the idea applies immediately to minibatches of size b by sampling a set of indices \mathcal{I}_k without replacement with $|\mathcal{I}_k| = b$. In this case, $f(\theta; z_{\xi_k})$ in (3.2) is replaced by $\frac{1}{b} \sum_{i \in \mathcal{I}_k} f(\theta; z_i)$; Algorithm 1 summarizes the procedure in pseudocode.

It is worth noting that although the connection to implicit gradient descent provides a useful perspective to derive an analogy, our algorithm does *not* require differentiability in practice. Algorithm 1 requires only that the projection and the proximal mapping are computable. This is the case for many common constraints (see Chapter 1), and a convex loss component suffices to guarantee the existence of the proximal mapping of $P_C(\theta_{k-1})$. After convergence, we advocate one final projection step to ensure constraints are enforced exactly, as denoted in the last line of the pseudocode.

Intuition and Convergence The proximal distance algorithm is originally derived from the perspective of MM. From this view, the stochastic proximal distance algorithm (3.2) substitutes the loss at one data point to approximate the average loss across all the data when building the surrogate function. Under this approximation, however, majorization of the objective no longer holds. Thus, the overall geometry and descent property from the MM algorithm it modifies no longer carry through. Our intuition tells us that while the loss evaluated on a subsample fluctuates, it has mean equal to the objective and should tend to decrease the loss, though we do not have a precise theory of MM in expectation here.

We may also view the stochastic proximal distance algorithm (3.2) from the lens of gradient descent. At each iteration, we are approximating a true direction of steepest descent given the data, $\nabla F(\theta_k^+)$, by a noisy direction $\nabla f(\theta_k; z_{\xi_k})$ based on only a subsample from the data. Here θ_k^+ denotes the update from the full non-stochastic or *batch* update: that is, $\theta_k^+ = \text{prox}_{\rho_k^{-1}F}[P_C(\theta_{k-1})]$. It is worth pointing out that unlike explicit gradient methods, the gradient $\nabla f(\theta_k; z_{\xi_k})$ is no longer an unbiased estimate of the true direction (this is also true for implicit SGD). Denote $\theta_k^i = \text{prox}_{\rho_k^{-1}f(\cdot; z_i)}[P_C(\theta_{k-1})]$ for all $i = 1, \dots, n$: then denoting the natural filtration $\mathcal{F}_{k-1} = \sigma(\xi_1, \dots, \xi_{k-1})$,

$$E(\nabla f(\theta_k; z_{\xi_k}) | \mathcal{F}_{k-1}) = \frac{1}{n} \sum_{i=1}^n \nabla f(\theta_k^i; z_i) \neq \nabla F(\theta_k^+).$$

While we expect these to yield a descent direction on average, the right hand side emphasizes that the noisy direction is not necessarily unbiased as is the case in implicit update schemes. This motivates a rigorous analysis of the convergence properties, presented in Section 4.

Following the intuition in previous work on proximal distance algorithms, the penalty parameter sequence ρ_k must increase to approach the constrained solution. This agrees with viewing our method from the lens of stochastic gradient methods, where a diminishing learning rate is necessary to decrease the variance of the noisy direction to ensure convergence. We see then that here ρ_k plays a dual role as a step-size algorithmically, and as a parameter to modulate the transition from measure of fit to constraint term. This key observation, which we expand upon immediately below, leads us to import theoretical tools

used to analyze the convergence of stochastic approximation algorithms to better understand the sequence of proximal distance objectives, for which only heuristic justifications are previously available [XL21, LPZL22].

Connection to Incremental Proximal Methods Though the prior discussion does not rigorously establish convergence, we find a surprising connection between the stochastic algorithm and incremental proximal methods. Indeed, despite deriving from an entirely different perspective, we found Algorithm 1 shares an *identical update rule* with the incremental proximal algorithm of [Ber11]. There, the authors consider the problem

$$\min_{\theta \in C} \sum_{i=1}^n f(\theta; z_i) + g(\theta; z_i),$$

and derive updates consisting of a combination of an incremental proximal step and an incremental gradient step:

$$\begin{aligned} \eta_k &= P_C[\theta_k - \alpha_k \nabla f(\eta_k; z_{\xi_k})], \\ \theta_{k+1} &= P_C[\eta_k - \alpha_k \nabla g(\eta_k; z_{\xi_k})]. \end{aligned} \tag{3.3}$$

The authors show that the order of these steps is exchangeable, and moreover the projection in the first step is optional and can be omitted. If we set $g = 0$ and remove the optional projection, the first step in (3.3) becomes an unconstrained proximal map, and the second gradient step reduces to a projection step of its result. Upon inspection, this is equivalent to the update (3.2) defining our stochastic proximal distance algorithm. Bridging these methodologies will allow us to leverage techniques from that line of work toward convergence analysis in the following section.

Contrast to Prior Work We have seen that Algorithm 1 is similar in spirit to implicit SGD. In particular, the inverse penalty $1/\rho_k$ plays the same role as the learning rate α_k in the SGD rule (2.4). This observation is useful in the following section toward characterizing convergence and finite error bounds, incorporating properties of the constraint set C into existing theoretical frameworks for gradient-based methods. The additional projection

in (3.2) encourages iterates to stay near constraints throughout, in a sense performing a projected implicit step.

[KZL19] pointed out that the proximal distance algorithm reduces to the proximal gradient algorithm in the convex case. Consider the unconstrained problem

$$\min_{\theta} F(\theta) + q_{\rho}(\theta) \text{ such that } q_{\rho}(\theta) = \frac{\rho}{2} \|\theta - P_C(\theta)\|^2 :$$

since $\nabla q_{\rho}(\theta) = \rho[\theta - P_C(\theta)]$ in the convex case [Tod03, Lan16], the proximal distance algorithm for this problem with fixed ρ can equivalently be written as:

$$\theta_k = \text{prox}_{\rho^{-1}F}[P_C(\theta_{k-1})] = \text{prox}_{\rho^{-1}F}[\theta_{k-1} - \rho^{-1}\nabla q_{\rho}(\theta_{k-1})],$$

i.e. gradient descent on q_{ρ} followed by a proximal operation on F with step-size ρ^{-1} .

There are two key differences in considering the *stochastic versions* of the proximal distance and proximal gradient algorithms. First, with an increasing ρ schedule, the component q_{ρ} —and thus objective sequence—is also changing, which is unusual as proximal gradient methods address a particular fixed objective. Second, while both algorithms modify their non-stochastic counterparts via subsampling, the randomness is introduced *implicitly* in the stochastic proximal distance algorithm, and explicitly in the stochastic proximal gradient algorithm (even though the update scheme of the latter has an implicit nature). Partial intuition is conveyed by observing that the proximal mapping in the proximal distance method *mirrors* the role it usually plays in related algorithms. That is, typically a gradient step is performed on the loss F , with a proximal map then applied to the (often non-smooth) regularizer q . Instead, proximal distance updates *reverse these roles*. In doing so, both the gradient step and the stochasticity enter the iteration implicitly, implying better behavior from a stability standpoint. This in part leads to the desirable properties we prove below.

Chapter 4

Theoretical Analysis

In this section, we establish convergence guarantees and finite error bounds for our proposed algorithm in the convex case. Specifically, analysis is conducted on the projected sequence $P_C(\theta_k)$. This is more straightforward and meaningful in non-asymptotic analysis, because $P_C(\theta_k)$ satisfies the constraint while θ_k does not. We present the assumptions and main results here with some discussion, with complete proof details for each provided in the Appendix.

In our analysis, we focus on polynomial schedules of increasing ρ_k . The first assumption guarantees that the Robbins-Monro conditions are satisfied within this class:

Assumption 1. $\rho_k = \rho_1 \cdot k^\gamma$, $0.5 < \gamma \leq 1$.

One immediately sees that the “learning rate” $\frac{1}{\rho_k}$ satisfies $\sum_{k=1}^{+\infty} \frac{1}{\rho_k} = +\infty$ and $\sum_{k=1}^{+\infty} \frac{1}{\rho_k^2} < +\infty$, which is common in proving convergence in stochastic optimization methods.

Next, we make assumptions to place our theoretical analysis in the convex setting.

Assumption 2. *The constraint set C is closed and convex, the objective F is continuous, strictly convex and coercive.*

Assumption 3. *The loss components $f(\theta; z_i)$ are differentiable and convex for any z_i .*

In particular, this ensures the solution is well-posed:

Proposition 1. *Under Assumption 2, the constrained solution $\theta_* = \arg \min_{\theta \in C} F(\theta)$ exists and is unique.*

The next assumption enables us to analyze the stochastic proximal update in an implicit gradient manner:

Assumption 4. *The gradient of the loss component $\nabla f(\theta, z_i)$ is L -Lipschitz for any z_i .*

This assumption is common in studies of implicit SGD [THA20, LLW22], and holds for many common models such as generalized linear models (GLMs).

Now, we begin by establishing the following proposition, useful for the proof of the main theorems. It also indicates that the stochastic proximal distance algorithm is stable in some sense, because the two quantities below can both explode in SGD if the step size is too large.

Proposition 2. *Under Assumptions 1, 2, 3 and 4, denote $s := \sum_{k=1}^{+\infty} \frac{1}{k^{2\gamma}}$ and $G^2 := E\{\|\nabla f(\theta_*; z_{\xi_k})\|^2\} = \frac{1}{n} \sum_{i=1}^n \|\nabla f(\theta_*; z_i)\|^2$, then the iterative sequence and its gradient are bounded as follows.*

- (a). $E\{\|P_C(\theta_k) - \theta_*\|^2\} \leq \exp(\frac{4L^2s}{\rho_1^2}) \cdot \{\frac{4G^2s}{\rho_1^2} + \|P_C(\theta_0) - \theta_*\|^2\} := r^2$
- (b). $E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2\} \leq 2(G^2 + r^2L^2) := c^2$

Our first theorem establishes convergence to the constrained solution with an appropriate polynomial rate schedule for ρ_k . The result is analogous to Proposition 9 of [Ber11], but our proof does *not* assume bounded gradients.

Theorem 1. *Under Assumptions 1, 2, 3 and 4, $P_C(\theta_k)$ converges to $\theta_* = \arg \min_{\theta \in C} F(\theta)$ almost surely.*

Next, we impose an “identifiability condition” in that there is enough curvature near the global solution so that it is sufficiently separated from other potential parameter values.

Assumption 5. *For all $\theta \in C$, the loss function F satisfies*

$$\nabla F(\theta)^T(\theta - \theta_*) \geq F(\theta) - F(\theta_*) + \frac{\mu}{2}\|\theta - \theta_*\|^2.$$

Note that Assumption 5 is implied by strong convexity, which is typical in finite error bound analysis in similar analyses of related methods [TA17]. That is, our condition is strictly weaker than those imposed in such studies—we only require the inequality to hold at θ_* rather than everywhere on the domain. Next, we characterize the convergence behavior of the iterate sequence in various regimes of the rate parameters.

Theorem 2. Under Assumptions 1, 2, 3, 4 and 5, denote $\delta_k := E[\|P_C(\theta_k) - \theta_*\|^2]$, and let c, r, G, s as defined in Proposition 2. Then it holds that

$$\delta_k \leq \frac{2m(1 + \frac{\mu}{\rho_1})}{\mu\rho_1} k^{-\gamma} + (1 + \frac{\mu}{\rho_1})^{-\phi_\gamma(k)} (\delta_0 + A)$$

where $m = 5c^2 + 2L^2r^2 + 2LGr$, $A = \frac{ms}{\rho_1^2}(1 + \frac{\mu}{\rho_1})^{k_0}$, $k_0 > 0$ is a constant, and

$$\phi_\gamma(k) = \begin{cases} k^{1-\gamma} & \text{if } \gamma \in (0.5, 1) \\ \log k & \text{if } \gamma = 1. \end{cases}$$

In particular, this indicates that the convergence rate for $\|P_C(\theta_k) - \theta_*\|^2$ is $\mathcal{O}(k^{-\gamma})$ whenever $\gamma \in (0.5, 1)$. Otherwise, if $\gamma = 1$, it becomes $\mathcal{O}(k^{-\min\{1, \log(1 + \frac{\mu}{\rho_1})\}})$. Another interesting takeaway is that the forgetting rate of the initial error δ_0 is $(1 + \frac{\mu}{\rho_1})^{-\phi_\gamma(k)}$, regardless of the initial penalty ρ_1 . This is one rigorous characterization of the stability properties inherited by our stochastic proximal distance algorithm as an implicit scheme. In contrast, for explicit stochastic gradient algorithms such as SGD, it is well known that the initial error can be amplified arbitrarily if the initial learning rate is improper [MB11, KLRT15]. We next establish a similar result characterizing global convergence of the objective sequence.

Theorem 3. Let Assumptions 1, 2, 3, 4 and 5, hold, and denote $\zeta_k := E\{F[P_C(\theta_k)] - F(\theta_*)\}$. Then,

$$\begin{aligned} \zeta_k \leq & G \sqrt{\frac{2m(1 + \frac{\mu}{\rho_1})}{\mu\rho_1} k^{-\frac{\gamma}{2}} + G(1 + \frac{\mu}{\rho_1})^{-\frac{\phi_\gamma(k)}{2}} \sqrt{\delta_0 + A}} \\ & + Bk^{-\gamma} + L \cdot (1 + \frac{\mu}{\rho_1})^{-\phi_\gamma(k)} (\delta_0 + A) \end{aligned}$$

where A, ϕ_γ and δ_0 are as defined in Theorem 2, and $B = \frac{\mu c^2 + 2Lm(1 + \frac{\mu}{\rho_1})}{\mu\rho_1}$.

This result indicates that the convergence rate of the objectives is $\mathcal{O}(k^{-\frac{\gamma}{2}})$ if $\gamma \in (0.5, 1)$ and becomes $\mathcal{O}(k^{-\min[\frac{1}{2}, \frac{1}{2}\log(1 + \frac{\mu}{\rho_1})]})$ for $\gamma = 1$. It follows that the fastest worst-case rate it can achieve is $\mathcal{O}(k^{-\frac{1}{2}})$. Moreover, we see that the convergence rate is consistent with that of the sequence of iterates $E\|P_C(\theta_k) - \theta_*\|^2$ in the sense that both rates are monotonically increasing with respect to γ . This differs from a similar analysis of the idealized proximal

Robbins-Monro method in [THA20], where the authors show that the convergence rate of $E\|\theta_k - \theta_*\|^2$ is also $\mathcal{O}(k^{-\gamma})$, but is at odds with the convergence rate of $E[F(\theta_k) - F(\theta_*)]$ which is non-monotonic in γ . Indeed, the best rate they obtain is of order $\mathcal{O}(k^{-\frac{1}{3}})$, and is achieved at $\gamma = \frac{2}{3}$ in their analysis.

We briefly remark that all proofs of the above results apply to the mini-batch (and full batch) case with only slight modifications. For the mini batch with batch size b case, A subset $\mathcal{I}_k \subset \{1, \dots, n\}$ of size b is drawn without replacement at each iteration, and $f(\theta; z_{\xi_k})$ is replaced by $\frac{1}{b} \sum_{i \in \mathcal{I}_k} f(\theta; z_i)$; the results hold immediately since the latter satisfies the same conditions as $f(\theta; z_{\xi_k})$.

Novelty of Theoretical Devices It is worth emphasizing several ways our analyses improve upon prior techniques. First, to establish almost sure convergence in Theorem 1, Assumption 4 imposes only that the *gradient* of the objective is Lipschitz. Our condition is significantly less restrictive than prior work guaranteeing almost sure convergence, and importantly can be verified on many common model classes such as GLMs. In contrast, the technique by [Ber11] requires that the gradient of each component function has bounded norm almost surely. From our strictly weaker condition, we may derive Proposition 2 (b), which substitutes for their quite strong condition a similar role in the proof. Importantly, a merit of implicit schemes lies in their advantage to avoid the instability from potentially explosive behavior in explicit gradient schemes. Assuming bounded gradient components as in prior work precludes this possibility, missing this “interesting” gap by restricting to a regime where explicit schemes also succeed.

Next, our finite error bounds in Theorem 2 do not require that the objective is twice differentiable and Lipschitz as in previous studies of implicit SGD [TA17]. More importantly, these past analyses posit that the loss is globally both Lipschitz and strongly convex, which cannot simultaneously hold true [AD19, LLW22]. While we are inspired by their analyses, our proof techniques for analyzing the finite error bounds of the objective function make a departure from recent studies of proximal algorithms by [THA20, LLW22]. They derive a recursion relating the objective error at θ_{k-1} and θ_k , and employ a clever technique

to analyze this recursion through constructing another sequence for bounding the error, resulting in a case-based analysis depending on the value of γ . In our setting, it is not obvious how to leverage these recursive arguments, as the projection operator complicates analysis. Instead, we study the relation between ζ_k and δ_k , an approach that does not require discussing different cases of γ values. As a result, our error bounds not only yield superior rates of convergence, but are arguably more natural in this sense.

Chapter 5

Empirical Study

5.1 Synthetic Data

As the breadth of proximal distance algorithms has been demonstrated in past works [XCL17, LPZL22], we focus on assessing whether the theoretical results align with performance in practice on synthetic experiments.

First we consider linear and logistic regression in low and high dimensional settings. Specifically, we consider a simple convex example involving the unit ball constraint $C_u = \{\theta \in R^p : \|\theta\|_2 \leq 1\}$, as well as the non-convex exact sparsity set $C_s = \{\theta \in R^p : \|\theta\|_0 \leq s\}$, where p is the dimension of covariates and the sparsity level is s . In both cases, projection is very efficient: $P_{C_u}(\theta)$ is obtained by simply scaling θ if its norm is larger than 1, and $P_{C_s}(\theta)$ is computed by setting all but the largest (in absolute value) s elements to 0. We will find that while our guarantees are sufficient for convergence, the algorithm often succeeds when not all conditions are met, exploring behavior for γ outside of the guaranteed range and considering *nonconvex* constraint such as C_s . We describe each simulation study below, with complete details on data generation and deriving the proximal updates in the Appendix.

Constrained Linear Regression At each iteration of Algorithm 1, denote the subsampled data as $\tilde{X} \in R^{b \times p}$ and corresponding observations $\tilde{y} \in R^b$, where b is the batch size. In the linear setting, the proximal update for θ has closed form:

$$\begin{aligned}\theta_k &= (b\rho_k I_p + \tilde{X}^T \tilde{X})^{-1} [b\rho_k P_C(\theta_{k-1}) + \tilde{X}^T \tilde{y}] \\ &= [I_p - \tilde{X}^T (b\rho_k I_b + \tilde{X} \tilde{X}^T)^{-1} \tilde{X}] [P_C(\theta_{k-1}) + b^{-1} \rho_k^{-1} \tilde{X}^T \tilde{y}].\end{aligned}$$

The second equation follows by invoking the Woodbury formula, and significantly reduces computations when $b \ll p$. We consider various combinations of (n, p) , and keep the batch

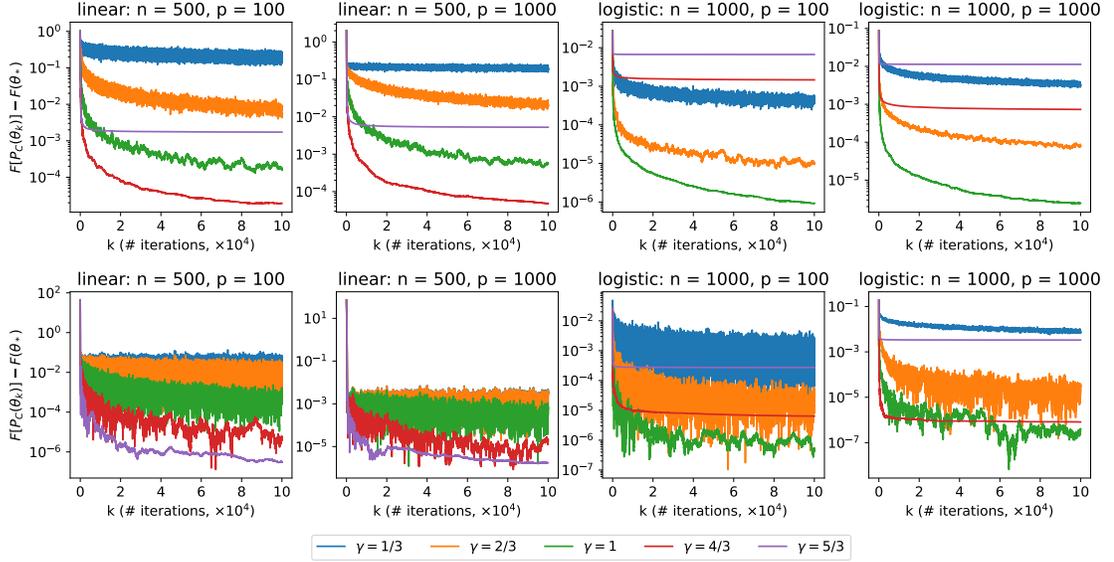


Figure (5.1) Finite error at different γ in different settings under the unit ball constraint (top panel) and sparsity set constraint (bottom panel)

size to be 5% of the sample size.

Constrained Logistic Regression Unlike the linear case above, there is no closed form solution for the proximal operation in constrained logistic regression. We employ a Newton iteration with update rule as follows:

For $i = 1, \dots, S$ Newton steps

$$\beta_i = \beta_{i-1} - \eta \cdot [\rho_k I + \frac{1}{b} \tilde{X}^T \tilde{W} \tilde{X}]^{-1} \{ -\frac{1}{b} \tilde{X}^T (\tilde{y} - \tilde{p}) + \rho_k [\beta_{i-1} - P_C(\theta_{k-1})] \}.$$

Here \tilde{p} is a $b \times 1$ vector, \tilde{W} is a $b \times b$ diagonal matrix; both depend on the parameter β_{i-1} . The stepsize η is selected by Armijo backtracking, and the inverse matrix above again can be rewritten to reduce complexity when $b \ll p$:

$$\rho_k^{-1} I_p - \rho_k^{-1} \tilde{X}^T (b \rho_k I_b + \tilde{W} \tilde{X} \tilde{X}^T)^{-1} \tilde{W} \tilde{X}.$$

We take subsamples of size 20% of the total dataset.

Validation of Theoretical Results We begin by validating the finite error properties in Theorems 2 and 3 under various penalty rates γ from a fixed initial penalty. Though our

proofs assume γ lies in $(\frac{1}{2}, 1]$, we explore values beyond this range as well in $\{\frac{1}{3}, \frac{2}{3}, 1, \frac{4}{3}, \frac{5}{3}\}$ since Assumption 1 is sufficient but not necessary. We set the same seed across different γ settings to control the stochasticity due to subsampling in our comparison; we set $\rho_1 = 10^{-1}$ in the unit ball example and $\rho_1 = 10^{-3}$ in the sparsity constraint case.

Figure 5.1 summarizes the results in terms of objective convergence $F[P_C(\theta_k)] - F(\theta_*)$. The analogous plots for the optimization variable $\|P_C(\theta_k) - \theta_*\|^2$ are displayed in Figure C.1 in the Appendix and convey the same trends; this is to be expected as Theorems 2 and 3 yield the same rate regimes for the iterate sequence and objective sequence. Moreover, in all settings, the algorithms succeed in making steady progress under values in the guaranteed range (albeit slowly for $\gamma = 1/3$). In line with our convergence theory, the ranking of their rates of progress is monotonic in $\gamma \in \{\frac{1}{3}, \frac{2}{3}, 1\}$. Interestingly, this trend continues beyond $\gamma > 1$ in some settings. For example, in the linear case under unit ball constraint, as γ increases to $\frac{4}{3}$, the algorithm continues to make even faster progress. However, the method appears to approach an asymptote prematurely for $\gamma = \frac{5}{3}$, a pattern that emerges starting from $\gamma = \frac{4}{3}$ in the logistic case. Similar patterns are observed for both models under the non-convex sparsity constraint. This suggests that the sufficient conditions in our general analysis are not necessary ones. In particular, it may be possible to derive sharper characterizations of convergence rates when additional “nice” structure of the objectives hold on a case by case basis, such as in our linear examples. We also note that the paths in the sparse regression example appear choppier than those subject to the unit ball, which may be due to the non-convexity of the sparsity set.

These trends agree with the results suggested by our theoretical analysis. Noting that increasing ρ is advantageous for the first term to decrease but disadvantageous for the second term in the bound given by Theorem 2, we additionally provide a thorough experimental study of the behavior and interaction between parameters. These additional empirical results appear in the Appendix to respect space limitations, and convey a more complete understanding of the qualitative behavior of the algorithm as ρ_1, γ vary.

Computational Advantages Finally, we present a runtime comparison to show that unsurprisingly, the stochastic version of the proximal distance algorithm outpaces its full batch counterpart. Note that in practice, one would not evaluate the objective values in Figure 5.1 as they require the entire dataset; this comparison focuses on naturally imposing a stopping rule based on relative difference in the parameter sequence. Rather than push the limits of how far one can scale the algorithm, which depends heavily on implementation and computing environment, we run a focused study on a single machine with in dimension $p = 1000$ and sample size n varying from 5×10^4 to 5×10^5 . We illustrate performance under sampling batches of size $b = 100, 500, 1000, n$.

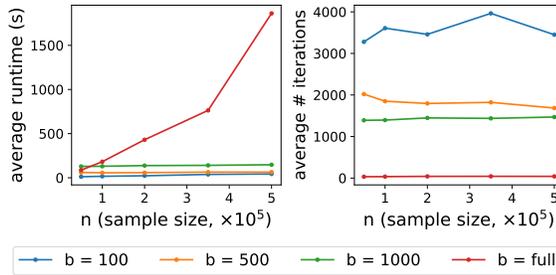


Figure (5.2) Runtime as sample size and batch size vary.

Results in terms of wall-clock time as well as iterations until convergence are displayed in Figure 5.2. Though the stochastic algorithm requires more iterations until convergence, a simple back-of-the-envelope calculation given by multiplying the batch size and iteration count reveals that it consistently succeeds having only considered a fraction of the complete dataset—for instance, 100×3500 requires processing much less data than even one iteration of the full batch algorithm. This discrepancy becomes more pronounced as the problem scale increases. Unsurprisingly, then, the stochastic method scales much better as the size of the entire dataset n grows in terms of real runtime.

5.2 Real Data

In this section, we assess our proposed stochastic proximal distance algorithm in solving a sparsity constrained logistic regression on a real-world data. The data we use is the Oral Toxicity data set from the UCI Machine Learning Repository ¹. It contains 1024 binary features (molecular fingerprints) and 1 binary response variable (toxic or not toxic), and there are 8992 observations (chemicals) in total.

Benchmark Algorithms We compare our proposed stochastic proximal distance algorithm with its batch version and two other projection based algorithms. One of them is the *projected gradient descent* method [JTK14] (also called iterative hard thresholding), which is the projection version of the gradient descent method for constrained optimization problems. In solving problem (1.1), it updates:

$$\theta_k = P_C[\theta_{k-1} - \alpha \nabla F(\theta_{k-1})],$$

where α is the step size, and this is an explicit gradient descent step followed by a projection operation. The other is the *projected SGD* [NJLS09], which is the projection version of the SGD. Similar to SGD, it approximates the gradient by a subset of the data and uses decreasing step sizes α_k . It updates:

$$\theta_k = P_C[\theta_{k-1} - \alpha_k \nabla f(\theta_{k-1}; z_{\xi_k})],$$

where z_{ξ_k} is randomly drawn from the data set and it can also be a mini batch.

Data Processing and Model Tuning We randomly split 80% and 20% of the data as the training set and testing set. It is worth noting that the data is imbalanced: the proportions of the positive (toxic) group and negative (not toxic) group in the training set are 8.29% and 91.71%, respectively. We tried common strategies (such as random oversampling and SMOTE) to address the imbalance, but they hardly work, so we finally used the original data.

¹<https://archive.ics.uci.edu/ml/datasets/QSAR+oral+toxicity>

For our stochastic proximal distance algorithm, we choose the increasing schedule of the penalty ρ_k to be $\rho_k = \rho_1 \cdot k$, and we use 10-fold cross validation to tune the sparsity parameter s and the initial penalty ρ_1 , it turns out that the optimal parameter values in our grid are $s = 100$ and $\rho_1 = 1 \times 10^{-2}$. For the projected gradient descent and projected SGD methods, we use the same sparsity level $s = 100$ and also tune their step sizes, where the decreasing schedule of the step sizes in the projected SGD is $\alpha_k = \alpha_1/k$ [NJLS09].

Results We assess the four algorithms in terms of computational speed and predictive performance. We report their runtime, number of iterations, accuracy, AUC and plot their ROC curves. Specifically, we tried two different batch sizes (200 and 500) for the stochastic proximal distance algorithm and the projected SGD are stochastic algorithms. Also, because these two algorithms are stochastic algorithms, we run each of them 30 times under each setting of batch size, and we run the two batch version algorithms once. The results of the runtime, iteration numbers, accuracy and AUC are shown in Figure 5.3, and the ROC curves are presented in Figure 5.4. There, “Sto-Prox” refers to the stochastic proximal distance algorithm and “Prox” is its batch version, “Proj-SGD” is the projected SGD algorithm and “Proj-GD” is its batch version, the number 200 and 500 refers to the batch size in the stochastic algorithms.

There are several interesting messages from the results. First, it can be seen that the batch version proximal distance algorithm is much slower than other algorithms, although it achieves the highest accuracy and AUC. This is a common drawback for proximal algorithms when the proximal mapping is hard to evaluate. Remarkably, our proposed stochastic proximal distance algorithms significantly reduces the runtime while maintaining satisfactory accuracy and AUC. The projected SGD method also scales well on this data set, but its accuracy on the test set equals 91.9% in all 30 experiments, which is exactly the proportion of the negative group in the testing set. This means that it is equivalent to a dumb classifier, which always classify observations in to the negative group, in terms of the accuracy. For the projected gradient descent method, it achieves high accuracy and AUC with a relatively short runtime. However, it requires much more iterations than the (stochastic) proximal

distance algorithm. Although it is not significantly slower than the stochastic proximal distance algorithm on this data, it's probable that it will have runtime much longer than the stochastic proximal distance algorithms on a larger data set.

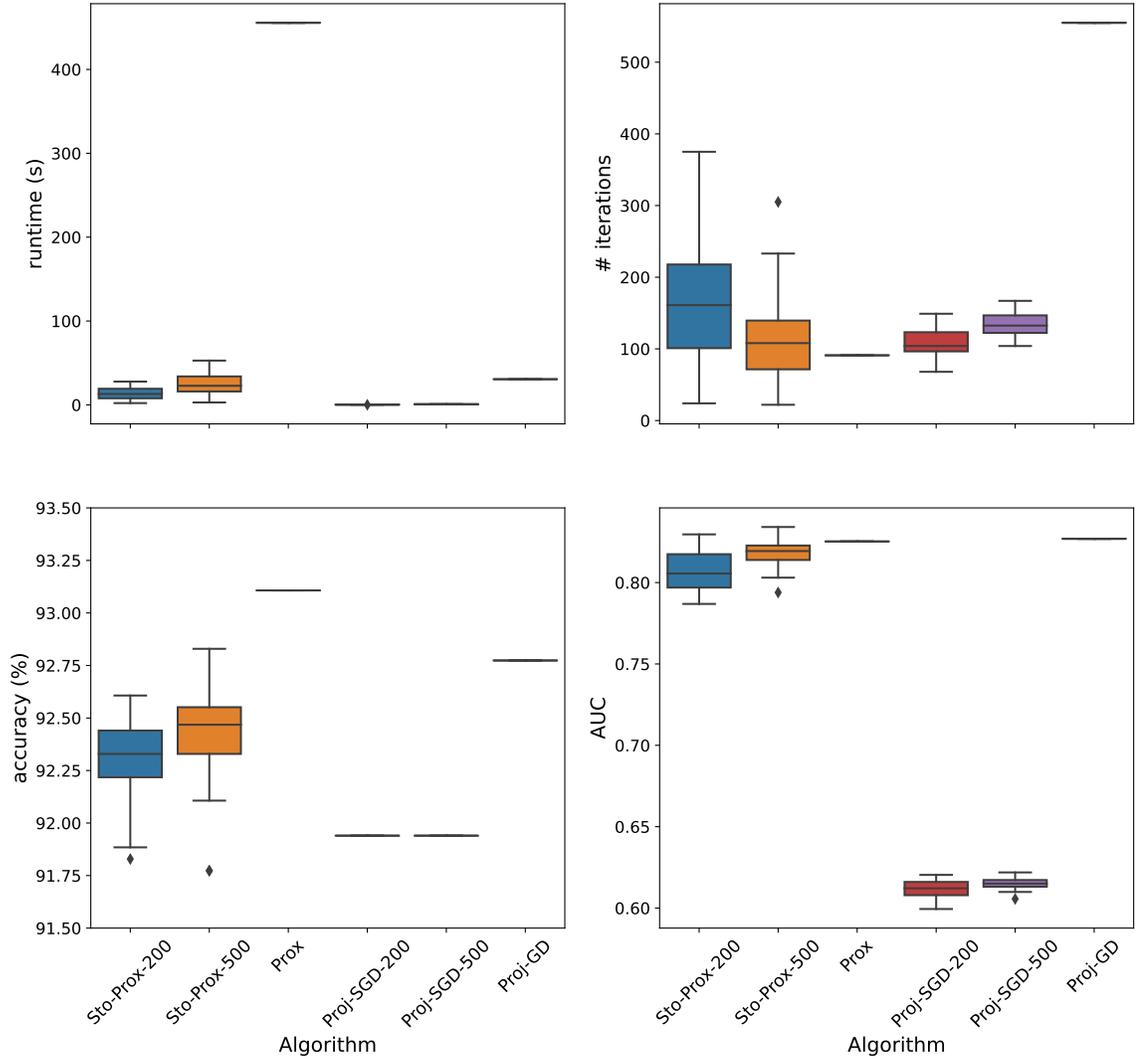


Figure (5.3) Box plots of runtime, iteration numbers, accuracy and AUC for all the algorithms.

In the above experiments, we used a zero vector as our initial guess for all the algorithms, because this setting had shown to perform well in sparse estimation problems in previous practices of the proximal distance algorithm. We also did additional experiments to explore the effects of initial guesses on the stochastic proximal distance algorithm and the projected

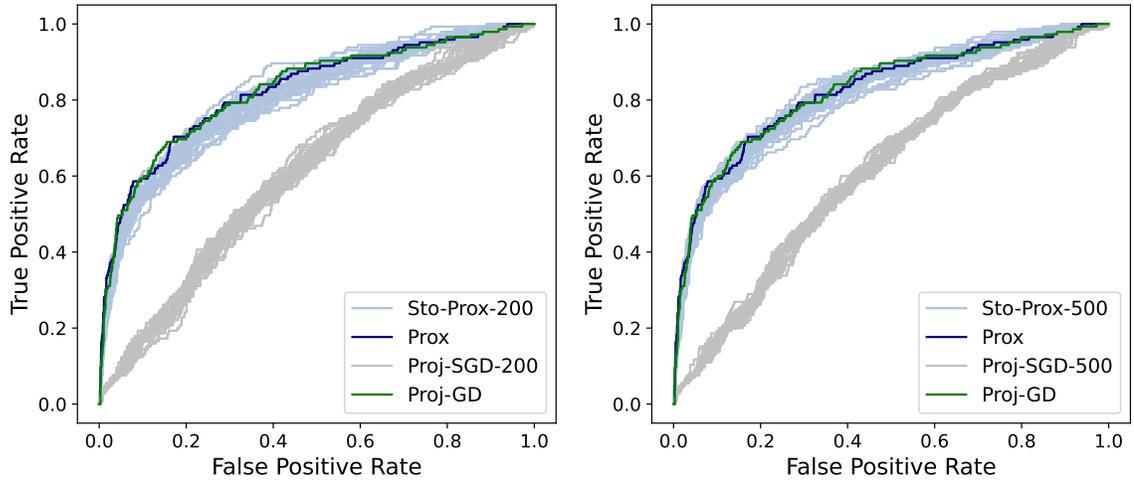


Figure (5.4) ROC curves for all the algorithms. The left plot contains the results for batch size 200 and the right one is for for batch size 500.

SGD. We run them 30 times using batch size 200, with a random initial guess θ_0 generated each time, where each element of θ_0 was independently drawn from $\text{Normal}(0, \sigma^2)$. We tried different values of σ and the results are summarized in Figure 5.5, which manifest that our proposed stochastic proximal distance algorithm is more robust than the projected SGD in terms of the choice of the initial guess.

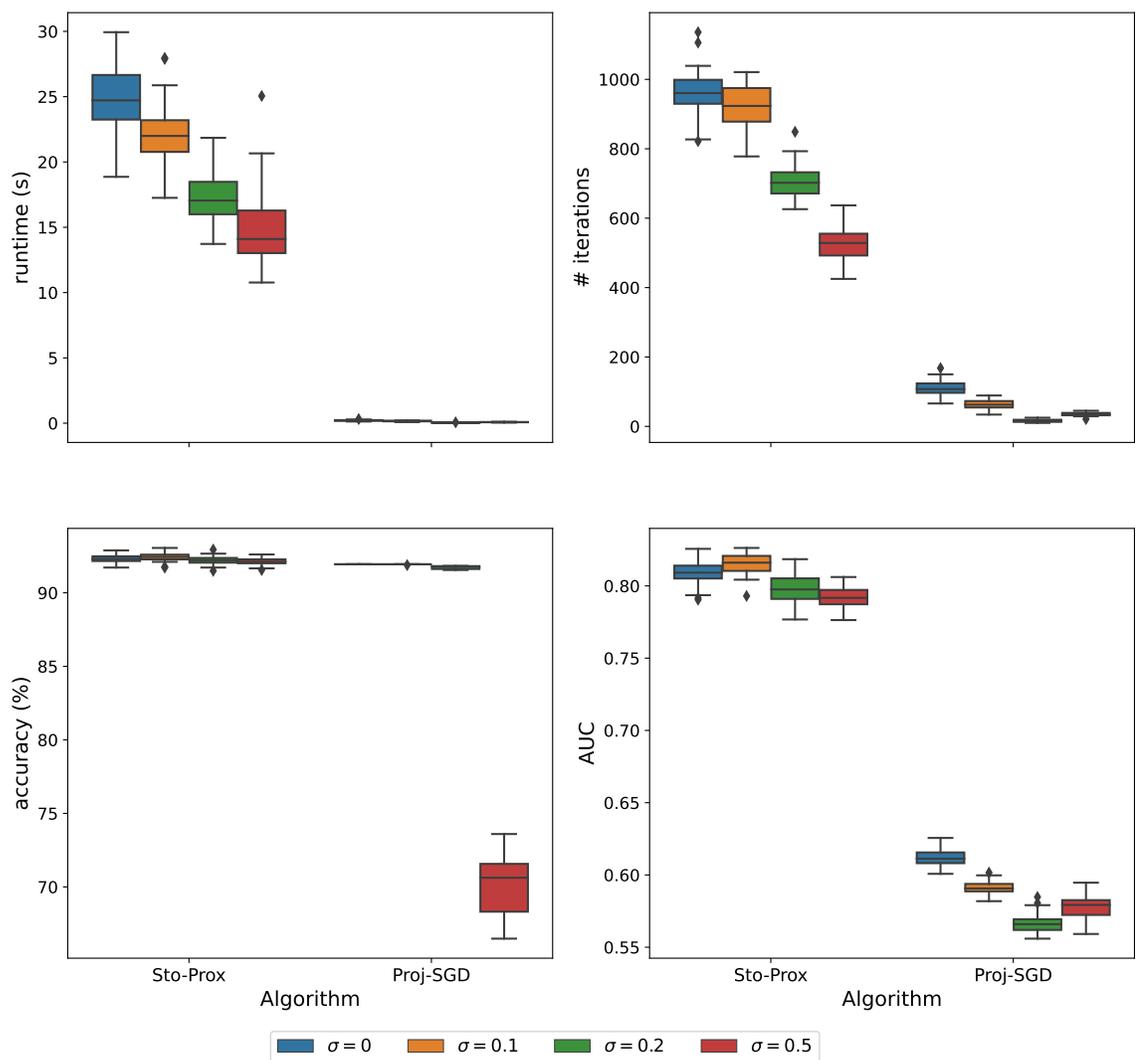


Figure (5.5) Results of the stochastic proximal distance algorithm and the projected SGD with random initial guesses.

Chapter 6

Conclusion

This paper proposes and analyzes a stochastic version of the recent proximal distance algorithm. The algorithm allows for a wide range of constraints to be considered in a projection-based framework, and our contributions now extend these merits to large-scale settings, and imbue them with new theoretical insights. In particular, we establish new convergence guarantees when the parameter ρ —and in turn the sequence of objectives—change over iterations, leveraging tools from the analysis of stochastic and incremental schemes when the MM geometry no longer holds.

We provide finite error analyses in the case of polynomial rate schedules, and our contributions open the door to future directions for exponential and other rate schedules [KZL19]. In particular, our experiments suggest that the conditions for convergence are sufficient, but may enjoy sharper analyses when assuming specific structures of objectives, or perhaps by investigating the interplay between mini-batch size and convergence rate. Finally, we focused our theoretical treatment to the convex setting, but even the simple sparse example makes use of projection onto a non-convex constraint set with empirical success. One of the advantages of proximal distance methods is its success in many non-convex settings [LPZL22]. Indeed, theoretical results on convergence are available in the fixed ρ setting and batch case; see [KZL19]. A challenging but natural direction to further investigate is deriving finite error bounds and characterizing convergence rates in such non-convex problems.

Appendix A

Proofs of Theoretical Results

A.1 Proof of Proposition 1

Proof. For some $\theta_0 \in C$, define $C' := C \cap \{\theta : F(\theta) \leq F(\theta_0)\}$. Because F is continuous, strictly convex and coercive, C' is a compact subset of C . Then, there exists a $\theta_* \in C'$ such that $F(\theta_*) = \min_{\theta \in C'} F(\theta)$. Because C' is constructed as the sublevel sets from θ_0 contained in C , it follows that $\theta_* \in C$ is the minimizer of F in C .

Next to prove its uniqueness, suppose by contradiction that there can exist two distinct solutions $\theta_1, \theta_2 \in C$, s.t. $F(\theta_1) = F(\theta_2) = \min_{\theta \in C} F(\theta)$. Since C is convex, for any $\lambda \in (0, 1)$, we have that $\theta_3 = \lambda\theta_1 + (1 - \lambda)\theta_2 \in C$. On the other hand, by strict convexity of F ,

$$F(\theta_3) < \lambda F(\theta_1) + (1 - \lambda)F(\theta_2) = \min_{\theta \in C} F(\theta).$$

This is a contradiction; it thus follows that the solution $\theta_* = \operatorname{argmin}_{\theta \in C} F(\theta)$ must be unique. \square

A.2 Proof of Proposition 2

We first introduce several lemmas.

Lemma 1. *For a proximal mapping $x_k = x_{k-1} - \frac{1}{\rho_k} \nabla g(x_k)$, where the function g is convex and differentiable, we have:*

$$(a). \quad \|\nabla g(x_k)\| \leq \|\nabla g(x_{k-1})\|$$

$$(b). \quad \|x_k - y\|^2 \leq \|x_{k-1} - y\|^2 - \frac{2}{\rho_k} [g(x_k) - g(y)] - \|x_k - x_{k-1}\|^2 \text{ for any } y$$

Here result (a) is collected from Lemma 3.1 of [LLW22] and result (b) is a variant of Proposition 1 in [Ber11]. This lemma indicates that a proximal update based on g is

guaranteed to make progress with respect to g : (a) says that the gradient of the updated parameter will have smaller norm than before; (b) implies that the updated parameter will get closer to the minimal point if we choose y to be a minimum point of g .

Proof. Result (b) can be seen upon expansion:

$$\begin{aligned}
\|x_{k-1} - y\|^2 &= \|x_{k-1} - x_k + x_k - y\|^2 \\
&= \|x_{k-1} - x_k\|^2 + \|x_k - y\|^2 + 2(x_{k-1} - x_k)^T(x_k - y) \\
&= \|x_{k-1} - x_k\|^2 + \|x_k - y\|^2 + \frac{2}{\rho_k} \nabla g(x_k)^T(x_k - y) \\
&\geq \|x_{k-1} - x_k\|^2 + \|x_k - y\|^2 + \frac{2}{\rho_k} [g(x_k) - g(y)] \quad (\text{by convexity of } g)
\end{aligned}$$

Upon rearranging, the result follows:

$$\|x_k - y\|^2 \leq \|x_{k-1} - y\|^2 - \frac{2}{\rho_k} [g(x_k) - g(y)] - \|x_{k-1} - x_k\|^2$$

□

Lemma 2. *If a non-negative sequence $\{y_0, y_1, \dots\}$ satisfies that $y_k \leq (1 + a_k)y_{k-1} + b_k$, where $a_k, b_k \geq 0$, $\sum_{i=1}^{+\infty} a_i < +\infty$ and $\sum_{i=1}^{+\infty} b_i < +\infty$. Then $\forall k$,*

$$y_k \leq \exp\left(\sum_{i=1}^{+\infty} a_i\right) \cdot \sum_{i=1}^{+\infty} b_i + \exp\left(\sum_{i=1}^{+\infty} a_i\right) \cdot y_0.$$

Proof. By induction,

$$\begin{aligned}
y_k &\leq b_k + (1 + a_k)y_{k-1} \\
&\leq b_k + (1 + a_k)b_{k-1} + (1 + a_k)(1 + a_{k-1})y_{k-2} \\
&\leq b_k + (1 + a_k)b_{k-1} + (1 + a_k)(1 + a_{k-1})b_{k-2} + (1 + a_k)(1 + a_{k-1})(1 + a_{k-2})y_{k-3} \\
&\dots \\
&\leq b_k + (1 + a_k)b_{k-1} + (1 + a_k)(1 + a_{k-1})b_{k-2} + \dots + (1 + a_k)(1 + a_{k-1})\dots(1 + a_2)b_1 \\
&\quad + (1 + a_k)(1 + a_{k-1})\dots(1 + a_1)y_0 \\
&\leq C_k \cdot [b_k + b_{k-1} + \dots + b_1] + C_k \cdot y_0,
\end{aligned}$$

where $C_k = (1 + a_k)(1 + a_{k-1})\dots(1 + a_1)$. We have $C_k \leq \exp(\sum_{i=1}^{+\infty} a_i)$ because

$$\log C_k = \sum_{i=1}^k \log(1 + a_i) \leq \sum_{i=1}^k a_i \leq \sum_{i=1}^{+\infty} a_i.$$

Therefore,

$$\begin{aligned}
y_k &\leq \exp\left(\sum_{i=1}^{+\infty} a_i\right) \cdot [b_k + b_{k-1} + \dots + b_1] + \exp\left(\sum_{i=1}^{+\infty} a_i\right) \cdot y_0 \\
&\leq \exp\left(\sum_{i=1}^{+\infty} a_i\right) \cdot \sum_{i=1}^{+\infty} b_i + \exp\left(\sum_{i=1}^{+\infty} a_i\right) \cdot y_0
\end{aligned}$$

□

Now we are in the position to prove Proposition 2.

Proof of Proposition 2

Proof. We first prove result (a), beginning by seeking a recursion of the squared difference between the projected iterate and θ_* :

$$\begin{aligned}
&\|P_C(\theta_k) - \theta_*\|^2 \\
&\leq \|\theta_k - \theta_*\|^2 \quad (\text{by non-expansive property}) \\
&\leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{f(\theta_k; z_{\xi_k}) - f(\theta_*; z_{\xi_k})\} - \|\theta_k - P_C(\theta_{k-1})\|^2 \quad (\text{by Lemma 1 (b)}) \\
&\leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} + \frac{2}{\rho_k} \{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_k; z_{\xi_k})\} \\
&\leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} \\
&\quad + \frac{2}{\rho_k} \nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T [P_C(\theta_{k-1}) - \theta_k] \quad (\text{by convexity}) \\
&= \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} + \frac{2}{\rho_k^2} \nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T \nabla f(\theta_k; z_{\xi_k}) \\
&\leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} + \frac{2}{\rho_k^2} \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2,
\end{aligned} \tag{A.1}$$

where the last inequality is due to Cauchy–Schwarz inequality and Lemma 1 (a). Also, due to triangular inequality and Assumption 4, we have

$$\begin{aligned}
\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\| - \|\nabla f(\theta_*; z_{\xi_k})\| &\leq \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}] - \nabla f(\theta_*; z_{\xi_k})\| \\
&\leq L \cdot \|P_C(\theta_{k-1}) - \theta_*\| \\
&\Rightarrow \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\| \leq \|\nabla f(\theta_*; z_{\xi_k})\| + L \|P_C(\theta_{k-1}) - \theta_*\| \\
&\Rightarrow \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \leq 2(\|\nabla f(\theta_*; z_{\xi_k})\|^2 + L^2 \|P_C(\theta_{k-1}) - \theta_*\|^2),
\end{aligned} \tag{A.2}$$

where the last inequality is because $(a + b)^2 \leq 2(a^2 + b^2)$. Plugging (A.2) into (A.1) leads to:

$$\begin{aligned} \|P_C(\theta_k) - \theta_*\|^2 &\leq \left(1 + \frac{4L^2}{\rho_k^2}\right) \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} \\ &\quad + \frac{4}{\rho_k^2} \|\nabla f(\theta_*; z_{\xi_k})\|^2, \end{aligned}$$

Taking expectation on both sides leads to:

$$\begin{aligned} &E\{\|P_C(\theta_k) - \theta_*\|^2\} \\ &\leq \left(1 + \frac{4L^2}{\rho_k^2}\right) E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} - \frac{2}{\rho_k} E\{F[P_C(\theta_{k-1})] - F(\theta_*)\} + \frac{4}{\rho_k^2} E\{\|\nabla f(\theta_*; z_{\xi_k})\|^2\} \\ &\leq \left(1 + \frac{4L^2}{\rho_k^2}\right) E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{4}{\rho_k^2} E\{\|\nabla f(\theta_*; z_{\xi_k})\|^2\} \\ &= \left(1 + \frac{4L^2}{\rho_1^2 \cdot k^{2\gamma}}\right) E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{4G^2}{\rho_1^2 \cdot k^{2\gamma}}, \end{aligned}$$

where the second term on the RHS of the first line is obtained by taking conditional expectation given \mathcal{F}_{k-1} and then averaging over \mathcal{F}_{k-1} , the penultimate step is because that θ_* is the constrained solution, and the last inequality is obtained by plugging in $\rho_k = \rho_1 \cdot k^\gamma$ and $G^2 := E\{\|\nabla f(\theta_*; z_{\xi_k})\|^2\} = \frac{1}{n} \sum_{i=1}^n \|\nabla f(\theta_*; z_i)\|^2$. Now we can invoke Lemma 2 and obtain conclusion (a). To prove conclusion (b), we only need to take expectation on both sides of (A.2):

$$E\{\|\nabla f[P_C(\theta_k); z_{\xi_k}]\|^2\} \leq 2(G^2 + r^2 L^2) := c^2.$$

□

Based on the above Lemmas and Propositions, next we will prove the main Theorems.

A.3 Proof of Theorem 1

We prove a similar result to Proposition 9 of [Ber11] under strictly weaker assumptions. The difference is that we do not require that $\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|$ is bounded almost surely (see Assumption 4 (b) in [Ber11]). Instead, we assume only that the gradient of each component

function is L -Lipschitz, which does not require the gradients to be bounded. Under this assumption, we proved Proposition 2 that plays a similar role as the Assumption 4 (b) in [Ber11].

We first recall the Supermartingale Convergence Theorem.

Lemma 3. (*Supermartingale Convergence Theorem [RS71, Ber11]*) *Let Y_k , Z_k and W_k , $k = 0, 1, \dots$, be three sequences of random variables and let \mathcal{F}_k , $k = 0, 1, \dots$, be sets of random variables such that $\mathcal{F}_k \subset \mathcal{F}_{k+1}$ for all k . Suppose that:*

(1) *The random variables Y_k , Z_k and W_k are nonnegative, and are functions of the random variables in \mathcal{F}_k .*

(2) *For each k , $E\{Y_{k+1}|\mathcal{F}_k\} \leq Y_k - Z_k + W_k$.*

(3) *There holds, with probability 1, $\sum_{k=0}^{\infty} W_k < \infty$.*

Then we have $\sum_{k=0}^{\infty} Z_k < \infty$, and the sequence Y_k converges to a nonnegative random variable Y , with probability 1.

Proof of Theorem 1

Proof. We begin by bounding the squared difference between the projected iterate and θ^* , and we can start from Eq. (A.1):

$$\begin{aligned} \|P_C(\theta_k) - \theta_*\|^2 &\leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} \\ &\quad + \frac{2}{\rho_k^2} \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \end{aligned}$$

Taking expectations conditional on the filtration \mathcal{F}_{k-1} on both sides leads to:

$$\begin{aligned} E\{\|P_C(\theta_k) - \theta_*\|^2 \mid \mathcal{F}_{k-1}\} &\leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k} \{F[P_C(\theta_{k-1})] - F(\theta_*)\} \\ &\quad + \frac{2}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\} \end{aligned} \tag{A.3}$$

The last term on the RHS satisfies condition (3) in Lemma 3:

$$\sum_{k=1}^{\infty} \frac{2}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\} < \infty$$

with probability 1, because for any $M > 0$,

$$\begin{aligned}
& P\left(\sum_{k=1}^{\infty} \frac{2}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\} \geq M\right) \\
& \leq \frac{E\left\{\sum_{k=1}^{\infty} \frac{2}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\}\right\}}{M} \quad (\text{by Markov's inequality}) \\
& = \frac{\sum_{k=1}^{\infty} \frac{2}{\rho_k^2} E\{E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\}\}}{M} \\
& = \frac{\sum_{k=1}^{\infty} \frac{2}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2\}}{M} \\
& \leq \frac{\sum_{k=1}^{\infty} \frac{2c^2}{\rho_k^2}}{M} \quad (\text{by Proposition 2 (b), and the numerator is finite by Assumption 1}) \\
& \Rightarrow 1 \geq P\left(\sum_{k=1}^{\infty} \frac{2}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\} < M\right) \geq 1 - \frac{\sum_{k=1}^{\infty} \frac{2c^2}{\rho_k^2}}{M} \rightarrow 1, \text{ as } M \rightarrow +\infty \\
& \Rightarrow P\left(\sum_{k=1}^{\infty} \frac{2}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\} < +\infty\right) = 1
\end{aligned}$$

Now, all the conditions in Lemma 3 are met and we can invoke the supermartingale convergence theorem with (A.3). In particular, this implies that with probability 1, $\|P_C(\theta_k) - \theta_*\|^2$ converges to some nonnegative random variable Y , and

$$\sum_{k=1}^{\infty} \frac{2}{\rho_k} \{F[P_C(\theta_{k-1})] - F(\theta_*)\} < \infty. \quad (\text{A.4})$$

From here onward, one can establish the result by following an identical argument in Proposition 9 of [Ber11]. We instead provide a simpler, more direct proof for completeness.

To this end, we first show that with probability 1, there exists a subsequence indexed by $\{i_k\}_{k=1}^{+\infty}$ such that $F[P_C(\theta_{i_k-1})] - F(\theta_*) \rightarrow 0$. Assume by contradiction that with positive probability, no subsequences of $F[P_C(\theta_{k-1})] - F(\theta_*)$ converge to 0. This implies that we can find $\delta > 0$ such that only finitely many k satisfy $F[P_C(\theta_{k-1})] - F(\theta_*) < \delta$. Indeed, otherwise for all $\delta > 0$, the condition $F[P_C(\theta_{k-1})] - F(\theta_*) < \delta$ is satisfied for infinitely many k , so that we could have chosen such a convergent subsequence.

However, if $F[P_C(\theta_{k-1})] - F(\theta_*) < \delta$ only holds for finitely many k , then there are infinitely many terms in the sum that contribute positive mass. Since $\sum_{k=1}^{\infty} \frac{1}{\rho_k}$ is divergent

too,

$$\sum_{k=1}^{\infty} \frac{2}{\rho_k} \{F[P_C(\theta_{k-1})] - F(\theta_*)\} = \infty,$$

a contradiction of (A.4).

Therefore, it must hold with probability 1 that there exists a subsequence $F[P_C(\theta_{i_k-1})] - F(\theta_*) \rightarrow 0$. As a result, with probability 1, the same limit Y given by the supermartingale theorem is shared by the subsequence indexed $\{i_k\}_{k=1}^{+\infty}$:

$$\|P_C(\theta_{i_k-1}) - \theta_*\|^2 \rightarrow Y.$$

It remains to transfer a similar result onto the iterate sequence. Let $\omega \in \Omega$, where we denote the probability space $(\Omega, \mathcal{F}, \mathcal{P})$. Since F is *strictly* convex and θ_* is the (constrained) minimizer, we have

$$\lim_{k \rightarrow \infty} \|P_C(\theta_{i_k-1}(\omega)) - \theta_*\|^2 > 0 \Rightarrow \lim_{k \rightarrow \infty} F[P_C(\theta_{i_k-1}(\omega))] - F(\theta_*) > 0.$$

It follows that

$$P(\lim_{k \rightarrow \infty} \|P_C(\theta_{i_k-1}) - \theta_*\|^2 > 0) \leq P(\lim_{k \rightarrow \infty} F[P_C(\theta_{i_k-1})] - F(\theta_*) > 0) = 0.$$

That is, with probability 1 the limit $Y = 0$, and as a result the entire iterate sequence $P_C(\theta_k) \rightarrow \theta_*$ almost surely.

□

A.4 Proof of Theorem 2

We first introduce a useful lemma.

Lemma 4. (Combining Lemma 2 and Corollary 1 in [THA20]) For sequences $a_n = a_1 n^{-\alpha}$ and $b_n = b_1 n^{-\beta}$, where $\alpha > \beta$, and $a_1, b_1, \beta > 0$ and $1 < \alpha < 1 + \beta$. Define,

$$\delta_n \triangleq \frac{1}{a_n} (a_{n-1}/b_{n-1} - a_n/b_n).$$

If a positive sequence $y_n > 0$ satisfies the recursive inequality

$$y_n \leq \frac{1}{1 + b_n} y_{n-1} + a_n.$$

Then,

$$y_n \leq 2 \frac{a_1(1 + b_1)}{b_1} n^{-\alpha + \beta} + \exp(-\log(1 + b_1)\phi_\beta(n))[y_0 + (1 + b_1)^{n_0} A],$$

where $A = \sum_{i=1}^{\infty} a_i$, $\phi_\beta(n) = n^{1-\beta}$ if $\beta \in (0.5, 1)$ and $\phi_\beta(n) = \log n$ if $\beta = 1$, and n_0 is a positive integer satisfying that $\delta_{n_0} < 1$.

Now we begin to prove Theorem 2.

Proof of Theorem 2

Proof. Our objective is to seek a recursion for $E\{\|P_C(\theta_k) - \theta_*\|^2\}$. We first borrow a nice idea from recent work by [LLW22], which entails rewriting our implicit scheme as an explicit update plus a remainder term R_k . This allows for a cleaner analysis, showing that the terms involving R_k go to zero sufficiently fast. To this end, we begin by rewriting the updating rule as

$$\begin{aligned} \theta_k &= P_C(\theta_{k-1}) - \frac{1}{\rho_k} \nabla f(\theta_k; z_{\xi_k}) \\ &= P_C(\theta_{k-1}) - \frac{1}{\rho_k} \nabla f[P_C(\theta_{k-1}); z_{\xi_k}] + R_k, \end{aligned}$$

where

$$R_k = \frac{1}{\rho_k} \{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}] - \nabla f(\theta_k; z_{\xi_k})\}.$$

The motivation for doing this rewrite is the fact mentioned in the main paper that

$$E\{\nabla f(\theta_k; z_{\xi_k}) \mid \mathcal{F}_{k-1}\} \neq \nabla F(\theta_k),$$

but

$$E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}] \mid \mathcal{F}_{k-1}\} = \nabla F[P_C(\theta_{k-1})].$$

This latter equation enables us to use the fact that F is strongly convex at θ_* (Assumption

5): we have

$$\begin{aligned}
& E\{\|P_C(\theta_k) - \theta_*\|^2\} \\
& \leq E\{\|\theta_k - \theta_*\|^2\} \quad (\text{by non-expansive property}) \\
& = E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{1}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2\} + E\{\|R_k\|^2\} \\
& \quad - \frac{2}{\rho_k} E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T [P_C(\theta_{k-1}) - \theta_*]\} \\
& \quad - \frac{2}{\rho_k} E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T R_k\} + 2E\{[P_C(\theta_{k-1}) - \theta_*]^T R_k\}.
\end{aligned} \tag{A.5}$$

Then we proceed to bound each term on the RHS of (A.5). The second term can be bounded by Proposition 2 (b):

$$\frac{1}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2\} \leq \frac{c^2}{\rho_k^2}$$

The third term:

$$\begin{aligned}
& E\{\|R_k\|^2\} \\
& = \frac{1}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}] - \nabla f(\theta_k; z_{\xi_k})\|^2\} \\
& = \frac{1}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 + \|\nabla f(\theta_k; z_{\xi_k})\|^2 - 2\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T \nabla f(\theta_k; z_{\xi_k})\} \\
& \leq \frac{4}{\rho_k^2} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2\} \quad (\text{by Cauchy-Schwarz inequality and Lemma 1 (a)}) \\
& \leq \frac{4c^2}{\rho_k^2} \quad (\text{by Proposition 2 (b)})
\end{aligned}$$

The fourth term:

$$\begin{aligned}
& E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T [P_C(\theta_{k-1}) - \theta_*]\} \\
& = E\{E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T [P_C(\theta_{k-1}) - \theta_*] \mid \mathcal{F}_{k-1}\}\} \\
& = E\{\nabla F[P_C(\theta_{k-1})]^T [P_C(\theta_{k-1}) - \theta_*]\} \\
& \geq E\{F[P_C(\theta_{k-1})] - F(\theta_*) + \frac{\mu}{2}\|P_C(\theta_{k-1}) - \theta_*\|^2\} \quad (\text{by Assumption 5}) \\
& \geq \frac{\mu}{2} E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} \quad (\text{since } \theta_* \text{ is the constrained solution}) \\
& \Rightarrow -\frac{2}{\rho_k} E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T [P_C(\theta_{k-1}) - \theta_*]\} \leq -\frac{\mu}{\rho_k} E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\}
\end{aligned}$$

The fifth term:

$$\begin{aligned}
& E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T R_k\} \\
&= \frac{1}{\rho_k} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 - \nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T \nabla f(\theta_k; z_{\xi_k})\} \\
&\geq \frac{1}{\rho_k} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 - \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\| \cdot \|\nabla f(\theta_k; z_{\xi_k})\|\} \\
&\geq 0 \quad (\text{by Lemma 1 (a)}) \\
&\Rightarrow -\frac{2}{\rho_k} E\{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]^T R_k\} \leq 0
\end{aligned}$$

Finally, to bound the sixth term:

$$\begin{aligned}
& 2E\{[P_C(\theta_{k-1}) - \theta_*]^T R_k\} \\
&= \frac{2}{\rho_k} E\{[P_C(\theta_{k-1}) - \theta_*]^T \{\nabla f[P_C(\theta_{k-1}); z_{\xi_k}] - \nabla f(\theta_k; z_{\xi_k})\}\} \\
&\leq \frac{2}{\rho_k} E\{\|P_C(\theta_{k-1}) - \theta_*\| \cdot \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}] - \nabla f(\theta_k; z_{\xi_k})\|\} \\
&\leq \frac{2L}{\rho_k} E\{\|P_C(\theta_{k-1}) - \theta_*\| \cdot \|P_C(\theta_{k-1}) - \theta_k\|\} \quad (\text{by Assumption 4}) \\
&= \frac{2L}{\rho_k^2} E\{\|P_C(\theta_{k-1}) - \theta_*\| \cdot \|\nabla f(\theta_k; z_{\xi_k})\|\} \\
&\leq \frac{2L}{\rho_k^2} E\{\|P_C(\theta_{k-1}) - \theta_*\| \cdot \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|\} \quad (\text{by Lemma 1 (a)}).
\end{aligned}$$

From the second line of Eq. (A.2), we have

$$\begin{aligned}
& \|P_C(\theta_{k-1}) - \theta_*\| \cdot \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\| \\
&\leq L \cdot \|P_C(\theta_{k-1}) - \theta_*\|^2 + \|\nabla f(\theta_*; z_{\xi_k})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\|.
\end{aligned}$$

Then, the sixth term

$$\begin{aligned}
& 2E\{[P_C(\theta_{k-1}) - \theta_*]^T R_k\} \\
&\leq \frac{2L^2}{\rho_k^2} E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{2L}{\rho_k^2} E\{\|\nabla f(\theta_*; z_{\xi_k})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\|\} \\
&= \frac{2L^2}{\rho_k^2} E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{2L}{\rho_k^2} E\{\|\nabla f(\theta_*; z_{\xi_k})\|\} \cdot E\{\|P_C(\theta_{k-1}) - \theta_*\|\} \\
&\leq \frac{2L^2 r^2 + 2LGr}{\rho_k^2},
\end{aligned}$$

where the last inequality is according to Proposition 2 (a) and Jensen's inequality, and the penultimate step is by independence.

Organizing these bounds together, we obtain the following recursion:

$$\begin{aligned}
& E\{\|P_C(\theta_k) - \theta_*\|^2\} \\
& \leq (1 - \frac{\mu}{\rho_k})E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{1}{\rho_k^2}(5c^2 + 2L^2r^2 + 2LGr) \\
& \leq \frac{1}{1 + \frac{\mu}{\rho_k}}E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{1}{\rho_k^2}(5c^2 + 2L^2r^2 + 2LGr) \\
& = \frac{1}{1 + \frac{\mu}{\rho_1 k^\gamma}}E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + \frac{1}{\rho_1^2 k^{2\gamma}}(5c^2 + 2L^2r^2 + 2LGr),
\end{aligned}$$

where the last step is due to plugging in $\rho_k = \rho_1 k^\gamma$, and the penultimate step is because that $(1 - \frac{\mu}{\rho_k})(1 + \frac{\mu}{\rho_k}) < 1$. Now we can invoke Lemma 4 and the result follows. \square

A.5 Proof of Theorem 3

Proof. We start from Eq. (A.1):

$$\begin{aligned}
\|P_C(\theta_k) - \theta_*\|^2 & \leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \frac{2}{\rho_k}\{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} \\
& \quad + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2
\end{aligned}$$

Upon rearranging and further manipulation, we obtain

$$\begin{aligned}
& \frac{2}{\rho_k}\{f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k})\} \\
& \leq \|P_C(\theta_{k-1}) - \theta_*\|^2 - \|P_C(\theta_k) - \theta_*\|^2 + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \\
& = \|P_C(\theta_{k-1}) - \theta_*\|^2 - \|P_C(\theta_k) - P_C(\theta_{k-1}) + P_C(\theta_{k-1}) - \theta_*\|^2 + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \\
& = -\|P_C(\theta_k) - P_C(\theta_{k-1})\|^2 - 2[P_C(\theta_k) - P_C(\theta_{k-1})]^T[P_C(\theta_{k-1}) - \theta_*] \\
& \quad + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \\
& \leq 2\|P_C(\theta_k) - P_C(\theta_{k-1})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\| + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \\
& \leq 2\|\theta_k - P_C(\theta_{k-1})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\| + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \quad (\text{by nonexpansiveness}) \\
& = \frac{2}{\rho_k}\|\nabla f(\theta_k; z_{\xi_k})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\| + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \\
& \leq \frac{2}{\rho_k}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\| \cdot \|P_C(\theta_{k-1}) - \theta_*\| + \frac{2}{\rho_k^2}\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \quad (\text{by Lemma 1 (a)})
\end{aligned}$$

Rearranging again and cancelling constants reveals

$$\begin{aligned}
& f[P_C(\theta_{k-1}); z_{\xi_k}] - f(\theta_*; z_{\xi_k}) \\
& \leq \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\| \cdot \|P_C(\theta_{k-1}) - \theta_*\| + \frac{1}{\rho_k} \cdot \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \\
& \leq L \cdot \|P_C(\theta_{k-1}) - \theta_*\|^2 + \|\nabla f(\theta_*; z_{\xi_k})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\| + \frac{1}{\rho_k} \cdot \|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2,
\end{aligned}$$

where the last inequality is due to the second line of Eq. (A.2). Taking conditional expectations given \mathcal{F}_{k-1} on both sides leads to

$$\begin{aligned}
& F[P_C(\theta_{k-1})] - F(\theta_*) \\
& \leq L \cdot \|P_C(\theta_{k-1}) - \theta_*\|^2 + E\{\|\nabla f(\theta_*; z_{\xi_k})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\| \mid \mathcal{F}_{k-1}\} \\
& \quad + \frac{1}{\rho_k} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2 \mid \mathcal{F}_{k-1}\},
\end{aligned}$$

and taking expectations over \mathcal{F}_{k-1} :

$$\begin{aligned}
& E\{F[P_C(\theta_{k-1})] - F(\theta_*)\} \\
& \leq L \cdot E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + E\{\|\nabla f(\theta_*; z_{\xi_k})\| \cdot \|P_C(\theta_{k-1}) - \theta_*\|\} \\
& \quad + \frac{1}{\rho_k} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2\} \\
& = L \cdot E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + E\{\|\nabla f(\theta_*; z_{\xi_k})\|\} \cdot E\{\|P_C(\theta_{k-1}) - \theta_*\|\} \\
& \quad + \frac{1}{\rho_k} E\{\|\nabla f[P_C(\theta_{k-1}); z_{\xi_k}]\|^2\} \quad (\text{two independent terms}) \\
& \leq L \cdot E\{\|P_C(\theta_{k-1}) - \theta_*\|^2\} + G \cdot E\{\|P_C(\theta_{k-1}) - \theta_*\|\} + \frac{c^2}{\rho_k},
\end{aligned}$$

where the last inequality is due to Proposition 2 and Jensen's inequality. Replacing the subscript $k-1$ by k will lead to the result:

$$\begin{aligned}
\zeta_k & = E\{F[P_C(\theta_k)] - F(\theta_*)\} \\
& \leq L \cdot E\{\|P_C(\theta_k) - \theta_*\|^2\} + G \cdot E\{\|P_C(\theta_k) - \theta_*\|\} + \frac{c^2}{\rho_{k+1}} \\
& \leq L\delta_k + G\sqrt{\delta_k} + \frac{c^2}{\rho_1} (k+1)^{-\gamma} \quad (\text{by Theorem 2 and Jensen's inequality}) \\
& \leq G\sqrt{\frac{2m(1 + \frac{\mu}{\rho_1})}{\mu\rho_1} k^{-\frac{\gamma}{2}}} + G(1 + \frac{\mu}{\rho_1})^{-\frac{\phi_\gamma(k)}{2}} \sqrt{\delta_0 + A} \\
& \quad + \frac{\mu c^2 + 2Lm(1 + \frac{\mu}{\rho_1})}{\mu\rho_1} k^{-\gamma} + L \cdot (1 + \frac{\mu}{\rho_1})^{-\phi_\gamma(k)} (\delta_0 + A).
\end{aligned}$$

In the last step, we combined the results in Theorem 2 and the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ to bound $\sqrt{\delta_k}$, we also rewrite $\frac{c^2}{\rho_{k+1}} = \frac{c^2}{\rho_1}(k+1)^{-\gamma} \leq \frac{c^2}{\rho_1}k^{-\gamma}$ to make the expression more concise without loss of rates. □

Appendix B

Derivation of Proximal Distance Updates

B.1 Sparsity/Unit Ball Constrained Linear Regression

Denote the whole data set as $X = (x_1, \dots, x_n)^T$ and $y = (y_1, \dots, y_n)^T$. Maximizing the log-likelihood under the sparsity/unit ball constraint leads to the objective:

$$\min_{\theta \in C} F(\theta) = \frac{1}{2n} \sum_{i=1}^n (y_i - x_i^T \theta)^2,$$

then the gradient of $F(\theta)$ is:

$$\begin{aligned} \nabla F(\theta) &= \frac{1}{n} \sum_{i=1}^n -(y_i - x_i^T \theta) x_i \\ &= -\frac{1}{n} (X^T y - X^T X \theta). \end{aligned}$$

For the stochastic proximal distance algorithm, the gradient at a subset $\{\tilde{x}_i, \tilde{y}_i\}_1^b$ is:

$$\frac{1}{b} \sum_{i=1}^b -(\tilde{y}_i - \tilde{x}_i^T \theta) \tilde{x}_i = -\frac{1}{b} (\tilde{X}^T \tilde{y} - \tilde{X}^T \tilde{X} \theta),$$

then the proximal update is:

$$\theta_k = P_C(\theta_{k-1}) + \frac{1}{\rho_k} \frac{1}{b} (\tilde{X}^T \tilde{y} - \tilde{X}^T \tilde{X} \theta_k),$$

and it has a closed form solution:

$$\begin{aligned} \theta_k &= (b\rho_k I + \tilde{X}^T \tilde{X})^{-1} [b\rho_k P_C(\theta_{k-1}) + \tilde{X}^T \tilde{y}] \\ &= [I_p - \tilde{X}^T (b\rho_k I_b + \tilde{X} \tilde{X}^T)^{-1} \tilde{X}] [P_C(\theta_{k-1}) + (b\rho_k)^{-1} \tilde{X}^T \tilde{y}]. \end{aligned}$$

The second equation follows by invoking the Woodbury formula, and significantly reduces computations in calculating the inverse when $b \ll p$.

B.2 Sparsity/Unit Ball Constrained Logistic Regression

For the whole data set $X = (x_1, \dots, x_n)^T$ and $y = (y_1, \dots, y_n)^T$, maximizing the log-likelihood under the sparsity/unit ball constraint leads to the objective:

$$\min_{\theta \in C} F(\theta) = -\frac{1}{n} \sum_{i=1}^n \{y_i(x_i^T \theta) - \log(1 + e^{x_i^T \theta})\}.$$

For the stochastic proximal distance algorithm, denote the loss on the subset $\{\tilde{x}_i, \tilde{y}_i\}_1^b$ at iteration k as

$$g_k(\theta) = -\frac{1}{b} \sum_{i=1}^b \{\tilde{y}_i(\tilde{x}_i^T \theta) - \log(1 + e^{\tilde{x}_i^T \theta})\},$$

then the proximal update is:

$$\theta_k = \arg \min_{\theta} g_k(\theta) + \frac{\rho_k}{2} \|\theta - P_C(\theta_{k-1})\|^2. \quad (\text{B.1})$$

Since a closed form solution is not available for (B.1), we follow [XCL17] and employ a Newton iteration to solve this proximal operation, the iterative algorithm is:

$$\beta_0 = \theta_{k-1}$$

For $i = 1, \dots, S$ Newton steps:

$$\beta_i = \beta_{i-1} - \eta * [\rho_k I + \nabla^2 g_k(\beta_{i-1})]^{-1} \{\nabla g_k(\beta_{i-1}) + \rho_k[\beta_{i-1} - P_C(\theta_{k-1})]\}$$

$$\theta_k = \beta_S.$$

Here η is selected by Armijo backtracking; the gradient and second derivative of g_k are:

$$\begin{aligned} \nabla g_k(\theta) &= -\frac{1}{b} \sum_{i=1}^b \left\{ \tilde{x}_i \left[\tilde{y}_i - \frac{e^{\tilde{x}_i^T \theta}}{1 + e^{\tilde{x}_i^T \theta}} \right] \right\} \\ &= -\frac{1}{b} \tilde{X}^T (\tilde{y} - \tilde{p}), \\ \nabla^2 g_k(\theta) &= -\frac{1}{b} \sum_{i=1}^b \left\{ -\tilde{x}_i \frac{e^{\tilde{x}_i^T \theta}}{(1 + e^{\tilde{x}_i^T \theta})^2} \tilde{x}_i^T \right\} \\ &= \frac{1}{b} \tilde{X}^T \tilde{W} \tilde{X}, \end{aligned}$$

where

$$\tilde{X} = \begin{pmatrix} \tilde{x}_1 \\ \dots \\ \tilde{x}_b \end{pmatrix}, \tilde{y} = \begin{pmatrix} \tilde{y}_1 \\ \dots \\ \tilde{y}_b \end{pmatrix}, \tilde{p} = \begin{pmatrix} \frac{e^{\tilde{x}_1^T \theta}}{1+e^{\tilde{x}_1^T \theta}} \\ \dots \\ \frac{e^{\tilde{x}_b^T \theta}}{1+e^{\tilde{x}_b^T \theta}} \end{pmatrix}, \tilde{W} = \begin{pmatrix} \frac{e^{\tilde{x}_1^T \theta}}{(1+e^{\tilde{x}_1^T \theta})^2} & & \\ & \ddots & \\ & & \frac{e^{\tilde{x}_b^T \theta}}{(1+e^{\tilde{x}_b^T \theta})^2} \end{pmatrix}.$$

Again, when the batch size $b \ll p$, we can save computation by the Woodbury formula:

$$\begin{aligned} (\rho_k I_p + \nabla^2 g_k(\theta))^{-1} &= (\rho_k I_p + \frac{1}{b} \tilde{X}^T \tilde{W} \tilde{X})^{-1} \\ &= \rho_k^{-1} I_p - \rho_k^{-1} \tilde{X}^T (b \rho_k I_b + \tilde{W} \tilde{X} \tilde{X}^T)^{-1} \tilde{W} \tilde{X}. \end{aligned}$$

Furthermore, according to [TAR14b], the proximal operation (B.1) can be simplified to a one-dimensional fixed-point problem when the batch size $b = 1$.

Appendix C

Data Generation and Additional Empirical Results

In both the sparse linear regression and the sparse logistic regression, we generate n independent data $\{y_i, x_i\}_{i=1}^n$, where $y_i \in R$ ($y_i \in \{0, 1\}$) is the response, $x_i = (x_{i1}, \dots, x_{ip})^T \in R^p$ is the predictor, and y_i depends on x_i through θ_{true} , the true regression coefficients. Under the unit ball constraint, $\|\theta_{\text{true}}\|_2 = 2$ and θ_{true} is specified by generating each element from the uniform distribution in $(-7, -4) \cup (4, 7)$ and then scaling the vector. Under the sparsity constraint, $\|\theta_{\text{true}}\|_0 = 5 \ll p$ and θ_{true} is specified by generating 5 i.i.d. random variables from the uniform distribution in $(-7, -4) \cup (4, 7)$ and randomly assigning them to 5 elements of $\mathbf{0}_p$.

Linear Case The data is generated by $x_{ij} \stackrel{i.i.d.}{\sim} N(0, 1)$, and $y_i|x_i \sim N(x_i^T \theta_{\text{true}}, 1)$.

Logistic Case The data is generated by $x_{ij} \stackrel{i.i.d.}{\sim} N(0, 0.3)$, and

$$y_i|x_i \sim \text{Bernoulli}\left(\frac{1}{1 + \exp\{-x_i^T \theta_{\text{true}}\}}\right)$$

Qualitative Behavior and Interaction Between ρ_1, γ We further conducted experiments to explore the effect of the initial penalty ρ_1 and its interaction with γ . Indeed, to illustrate the tradeoff, consider the case $\gamma = 1$, where the convergence rate depends on the relationship between $\log(1 + \frac{\mu}{\rho_1})$ and 1. Analogously to the previous experiments, we fixed $\gamma = 1$ and consider different ρ_1 , the analogous results of Figure 5.1 are shown in Figure C.2. These results show that increasing ρ_1 tends to accelerate the convergence up to a point, but adversely affects convergence if it becomes too large. This behavior is reminiscent of step size selection in general for gradient-based methods.

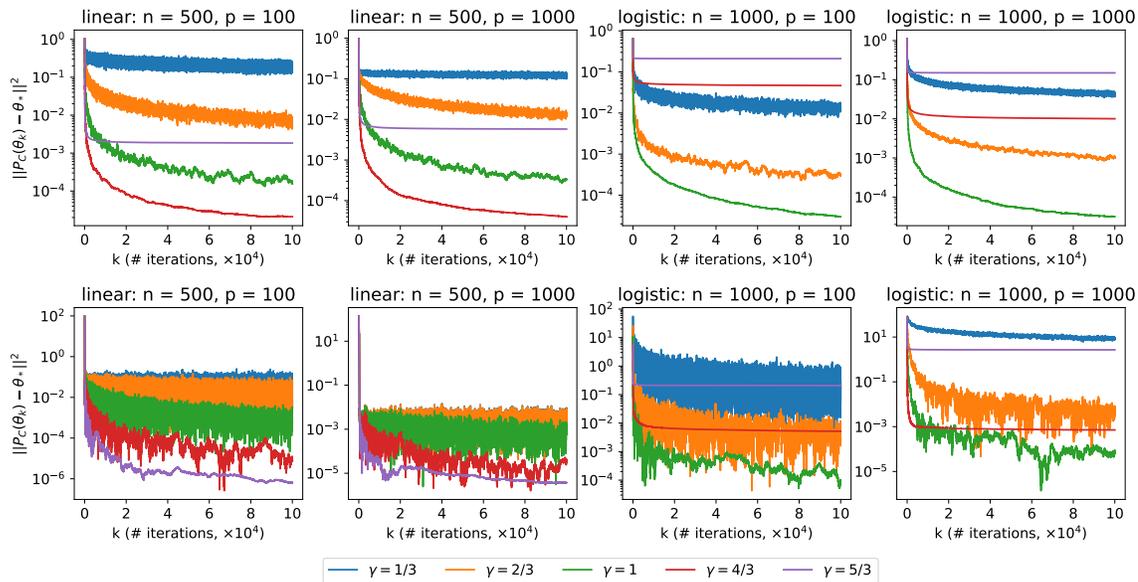


Figure (C.1) Finite error in terms of the parameter at different γ under the unit ball constraint (top panel) and sparsity constraint (bottom panel)

The next set of experiments explores changes in the qualitative behavior of the algorithms over many settings of ρ_1 and γ . We run the method for 10,000 iterations and report results averaged over 50 restarts in each setting. Box plots of log-transformed errors across all the settings appear in Figure C.3 and C.4 in the Appendix. As we increase the initial size ρ_1 , the performance improves up to a point, after which the algorithm begins to fail. This behavior is in line with the intuition behind increasing the penalties: as ρ_k increases, the focus of the algorithm gradually shifts from learning from the data to obeying the constraint. This works best when the transition occurs gradually, a heuristic related to γ which our theoretical results make rigorous. When ρ_1 is too large, the learning is too slow from the beginning, so that the algorithm quickly falls within the constraint but at a point that does not fit the data well. We also see that around this transition (for instance $\rho_1 = 10, \gamma = 5/3$ in the unit ball constrained low dimensional linear regression), cases where γ lies outside the guaranteed range may fail to converge instead of speed up progress—this can be understood as switching from the data-driven to constraint-driven setting too abruptly or prematurely. These trade-offs motivate future theoretical work on

characterizing convergence beyond the sufficient conditions we study, and also give practical guidance in line with intuitions behind the design of the algorithm.

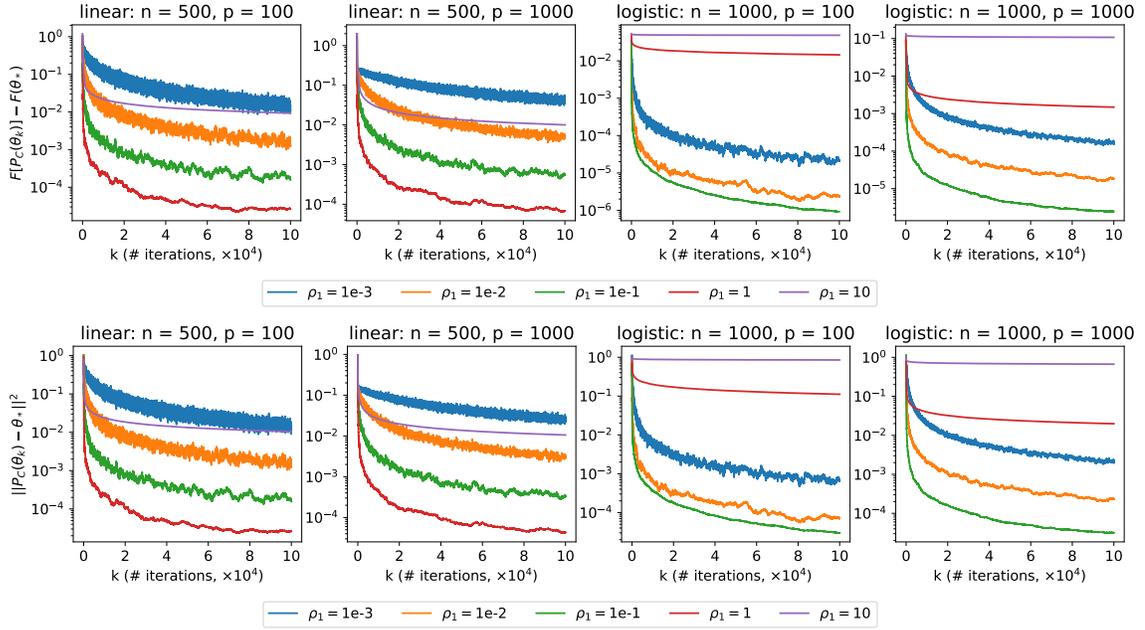


Figure (C.2) Finite error at different ρ_1 in different settings under the unit ball constraint

Figure C.3 presents the box plots of error at different combinations of $\gamma \in \{\frac{1}{3}, \frac{2}{3}, 1, \frac{4}{3}, \frac{5}{3}\}$ and $\rho_1 \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$ under unit ball constrained linear/logistic regression of various dimensions, with 50 repeated experiments at each setting. Figure C.4 presents the results under the sparsity constraint.

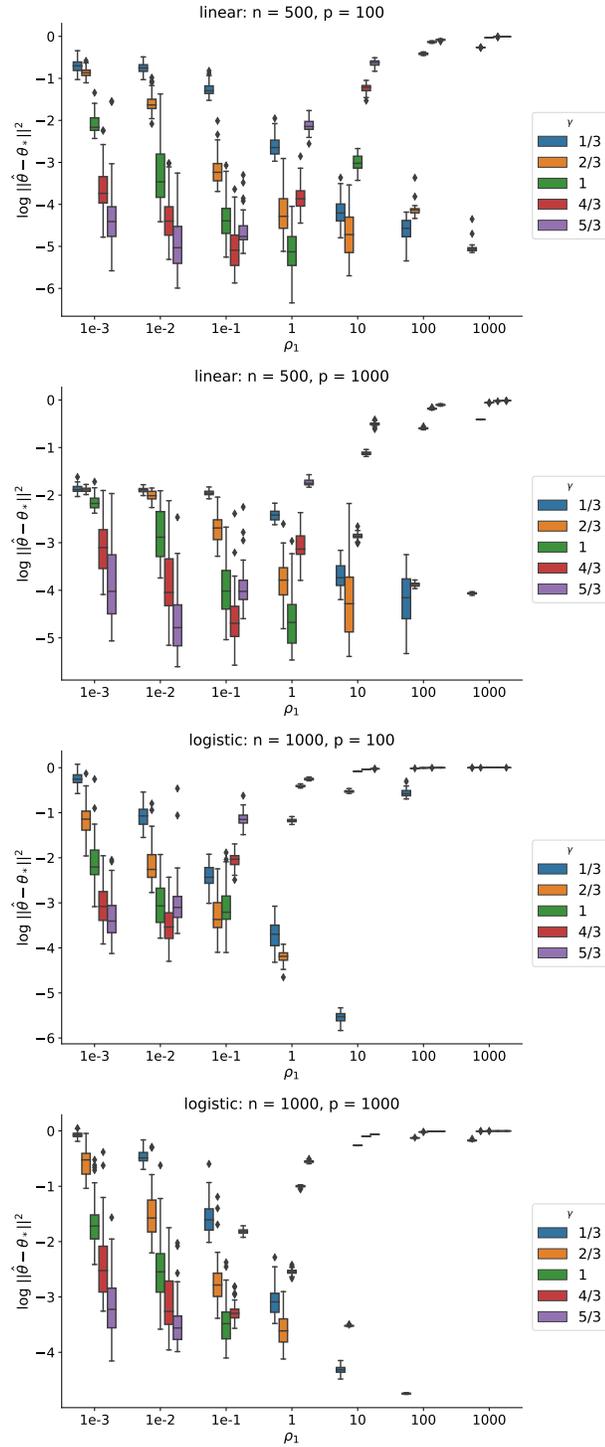


Figure (C.3) Box plots of error in various settings under the unit ball constraint.

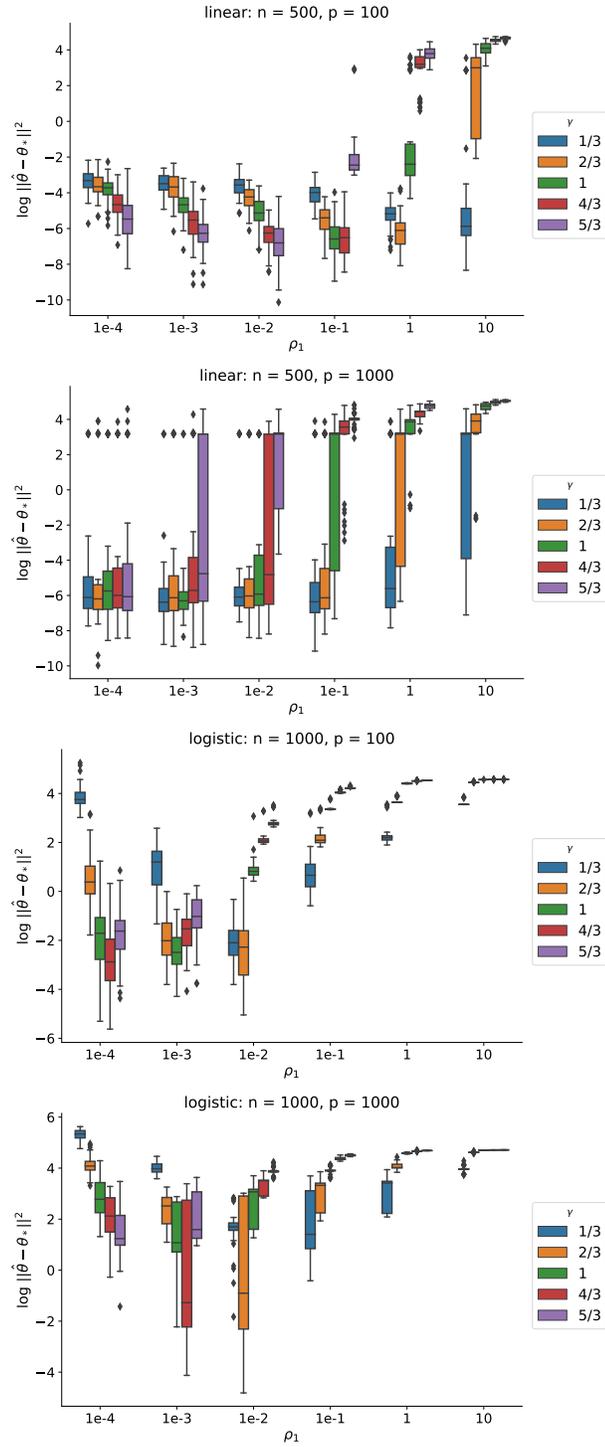


Figure (C.4) Box plots of error in various settings under the sparsity constraint. For reference to give a sense of scale, $\|\theta_*\| \approx 12$ here.

Bibliography

- [AD19] Hilal Asi and John C Duchi. Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019.
- [AFM14] Yves F Atchade, Gersende Fort, and Eric Moulines. On stochastic proximal gradient algorithms. *arXiv preprint arXiv:1402.2365*, 23, 2014.
- [AFM17] Yves F Atchadé, Gersende Fort, and Eric Moulines. On perturbed proximal gradient algorithms. *The Journal of Machine Learning Research*, 18(1):310–342, 2017.
- [BB73] Edward J Beltrami and Robert H Bryant. An algorithmic approach to nonlinear analysis and optimization. 1973.
- [BC⁺11] Heinz H Bauschke, Patrick L Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- [Ber11] Dimitri P Bertsekas. Incremental proximal methods for large scale convex optimization. *Mathematical programming*, 129(2):163–195, 2011.
- [Ber14] Dimitri P Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.
- [Bia16] Pascal Bianchi. Ergodic convergence of a stochastic proximal point algorithm. *SIAM Journal on Optimization*, 26(4):2235–2260, 2016.
- [BM17] Alberto Bietti and Julien Mairal. Stochastic optimization with variance reduction for infinite datasets with finite sum structure. *Advances in Neural Information Processing Systems*, 30, 2017.
- [BPC⁺11] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction

- method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.
- [Cou43] Richard Courant. Variational methods for the solution of problems of equilibrium and vibrations. *Bulletin of the American Mathematical Society*, 49:1–23, 1943.
- [CZL14] Eric C Chi, Hua Zhou, and Kenneth Lange. Distance majorization and its applications. *Mathematical programming*, 146(1):409–436, 2014.
- [HL04] David R Hunter and Kenneth Lange. A tutorial on mm algorithms. *The American Statistician*, 58(1):30–37, 2004.
- [JTK14] Prateek Jain, Ambuj Tewari, and Purushottam Kar. On iterative hard thresholding methods for high-dimensional m-estimation. *Advances in neural information processing systems*, 27, 2014.
- [KLRT15] Jakub Konečný, Jie Liu, Peter Richtárik, and Martin Takáč. Mini-batch semi-stochastic gradient descent in the proximal setting. *IEEE Journal of Selected Topics in Signal Processing*, 10(2):242–255, 2015.
- [KZL19] Kevin L Keys, Hua Zhou, and Kenneth Lange. Proximal distance algorithms: Theory and practice. *The Journal of Machine Learning Research*, 20(1):2384–2421, 2019.
- [Lan16] Kenneth Lange. *MM optimization algorithms*. SIAM, 2016.
- [LK15] Kenneth Lange and Kevin L Keys. The proximal distance algorithm. *arXiv preprint arXiv:1507.07598*, 2015.
- [LL15] Huan Li and Zhouchen Lin. Accelerated proximal gradient methods for non-convex programming. *Advances in neural information processing systems*, 28, 2015.

- [LLW22] Yoonhyung Lee, Sungdong Lee, and Joong-Ho Won. Statistical inference with implicit SGD: proximal Robbins-Monro vs. Polyak-Ruppert, 2022.
- [LPZL22] Alfonso Landeros, Oscar Hernan Madrid Padilla, Hua Zhou, and Kenneth Lange. Extensions to the proximal distance method of constrained optimization. *Journal of Machine Learning Research*, 23(182):1–45, 2022.
- [Mai15] Julien Mairal. Incremental majorization-minimization optimization with application to large-scale machine learning. *SIAM Journal on Optimization*, 25(2):829–855, 2015.
- [MB11] Eric Moulines and Francis Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. *Advances in neural information processing systems*, 24, 2011.
- [Nit14] Atsushi Nitanda. Stochastic proximal gradient descent with acceleration techniques. *Advances in Neural Information Processing Systems*, 27, 2014.
- [NJLS09] Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.
- [PB⁺14] Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and trends[®] in Optimization*, 1(3):127–239, 2014.
- [PSW15] Nicholas G Polson, James G Scott, and Brandon T Willard. Proximal algorithms in statistics and machine learning. *Statistical Science*, 30(4):559–581, 2015.
- [RB14] Ernest K Ryu and Stephen Boyd. Stochastic proximal iteration: a non-asymptotic improvement upon stochastic gradient descent. *Author website, early draft*, 2014.
- [RM51] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

- [Roc76] R Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- [RS71] Herbert Robbins and David Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics*, pages 233–257. Elsevier, 1971.
- [TA17] Panos Toulis and Edoardo M. Airoldi. Asymptotic and finite-sample properties of estimators based on stochastic gradients. *The Annals of Statistics*, 45(4):1694–1727, 2017.
- [TAR14a] Panagiotis Toulis, Edoardo Airoldi, and Jason Rennie. Statistical analysis of stochastic gradient methods for generalized linear models. In *International Conference on Machine Learning*, pages 667–675. PMLR, 2014.
- [TAR14b] Panagiotis Toulis, Edoardo Airoldi, and Jason Rennie. Statistical analysis of stochastic gradient methods for generalized linear models. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 667–675, Beijing, China, 22–24 Jun 2014. PMLR.
- [THA20] Panos Toulis, Thibaut Horel, and Edoardo Airoldi. The proximal Robbins–Monro method. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83, 12 2020.
- [Tod03] Michael J Todd. Convex analysis and nonlinear optimization: Theory and examples. jonathan m. borwein and adrian s. lewis, springer, new york, 2000, 2003.
- [XCL17] Jason Xu, Eric Chi, and Kenneth Lange. Generalized linear model regression under distance-to-set penalties. *Advances in Neural Information Processing Systems*, 30, 2017.

- [XL19] Jason Xu and Kenneth Lange. Power k-means clustering. In *International conference on machine learning*, pages 6921–6931. PMLR, 2019.
- [XL21] Jason Xu and Kenneth Lange. A proximal distance algorithm for likelihood-based sparse covariance estimation. *arXiv preprint arXiv:2109.04497*, 2021.
- [XL22] Jason Xu and Kenneth Lange. A proximal distance algorithm for likelihood-based sparse covariance estimation. *Biometrika*, 02 2022. asac011.