arXiv:2002.07290v2 [math.OC] 2 Jul 2020

Stochastic Gauss-Newton Algorithms for Nonconvex Compositional Optimization

Quoc Tran-Dinh [1] **Nhan H. Pham** [1] **Lam M. Nguyen** [2]

Abstract

We develop two new stochastic Gauss-Newton algorithms for solving a class of non-convex stochastic compositional optimization problems frequently arising in practice. We consider both the expectation and finite-sum settings under standard assumptions, and use both classical stochastic and SARAH estimators for approximating function values and Jacobians. In the expectation case, we establish $\mathcal{O}\left(\varepsilon^{-2}\right)$ iteration-complexity to achieve a stationary point in expectation and estimate the total number of stochastic oracle calls for both function value and its Jacobian, where ε is a desired accuracy. In the finite sum case, we also estimate $\mathcal{O}\left(\varepsilon^{-2}\right)$ iteration-complexity and the total oracle calls with high probability. To our best knowledge, this is the first time such global stochastic oracle complexity is established for stochastic Gauss-Newton methods. Finally, we illustrate our theoretical results via two numerical examples on both synthetic and real datasets.

1. Introduction

We consider the following nonconvex stochastic compositional nonconvex optimization problem:

$$\min_{x\in\mathbb{R}^p}\Big\{\Psi(x) := \phi(F(x)) \equiv \phi\Big(\mathbb{E}_{\xi}\left[\mathbf{F}(x,\xi)\right]\Big)\Big\}, \tag{1}$$

where $\mathbf{F}:\mathbb{R}^p\times\Omega\to\mathbb{R}^q$ is a stochastic function defined on a probability space $(\Omega,\mathbb{P})$, $\phi:\mathbb{R}^q\to\mathbb{R}\cup\{+\infty\}$ is a proper, closed, and convex, but not necessarily smooth function, and F is the expectation of $\mathbf{F}$ w.r.t. to ξ.

As a special case, if Ω is finite, i.e. $\Omega := \{\xi_1,\cdots,\xi_n\}$ and $\mathbb{P}(\xi=\xi_i)=\mathbf{p}_i>0$ for $i\in[n]:=\{1,\cdots,n\}$ and $\sum_{i=1}^n\mathbf{p}_i = 1$, then by introducting $F_i(x) := n\mathbf{p}_i\mathbf{F}(x,\xi_i)$, $F(x)$ can be written into a finite-sum $F(x) := \frac{1}{n}\sum_{i=1}^n F_i(x)$, and (1) reduces to

$$\min_{x\in\mathbb{R}^p}\Big\{\Psi(x) := \phi(F(x)) \equiv \phi\Big(\frac{1}{n}\sum_{i=1}^n F_i(x)\Big)\Big\}. \tag{2}$$

This expression can also be viewed as a stochastic average approximation of $F(x) := \mathbb{E}_{\xi}\left[\mathbf{F}(x,\xi)\right]$ in (1). Note that the setting (1) is completely different from $\min_x\{\Psi(x) := \mathbb{E}_{\xi}\left[\phi(\mathbf{F}(x,\xi),\xi)\right]\}$ in Davis & Grimmer (2019); Davis & Drusvyatskiy (2019); Duchi & Ruan (2018).

Problem (1) or its special form (2) covers various applications in different domains (both deterministic and stochastic) such as penalized problems for constrained optimization, parameter estimation, nonlinear least-squares, system identification, statistical learning, dynamic programming, and minimax problems (Drusvyatskiy & Paquette, 2019; Duchi & Ruan, 2018; Lewis & Wright, 2016; Nesterov, 2007; Tran-Dinh & Diehl, 2011; Wang et al., 2017a). Note that both (1) and (2) cover the composite form

$$\min_{x\in\mathbb{R}^p}\Big\{\Psi(x) := \phi(F(x)) + g(x)\Big\}, \tag{3}$$

for a given convex function g if we introduce $\hat{\phi}(\cdot) := \phi(\cdot) + g(\cdot)$ and $\hat{F}(x) := [F(x); x]$ to reformulate it into (1) or (2). This formulation, on the other hand, is an extension of (1). We will also show how to handle (3) in Subsection 4.3.

Our goal in this paper is to develop novel stochastic methods to solve (1) and (2) based on the following assumptions:

Assumption 1.1. The function Ψ of (1) is bounded from below on its domain, i.e. $\Psi^\star := \inf_x\Psi(x) > -\infty$. The function ϕ is M_ϕ-Lipschitz continuous, and F is L_F-average smooth, i.e., there exist $M_\phi, L_F\in(0,+\infty)$ such that

$$\begin{cases} |\phi(u)-\phi(v)| \le M_\phi\|u-v\|, \ \forall u,v\in\mathbb{R}^q, \\ \mathbb{E}_{\xi}\left[\|\mathbf{F}'(x,\xi)-\mathbf{F}'(y,\xi)\|^2\right] \le L_F^2\|x-y\|^2, \ \forall x,y. \end{cases} \tag{4}$$

For the finite-sum case (2), we impose a stronger assumption that $\|F_i'(x)-F_i'(y)\|\le L_F\|x-y\|$ for all $x,y\in\mathbb{R}^p$ and all $i\in[n]$. Here, we use spectral norm for Jacobian.

The first version was online on Feb 17, 2020 on Arxiv. This is the second version (ICML 2020).

[1]Department of Statistics and Operations Research, The University of North Carolina at Chapel Hill, NC, USA. [2]IBM Research, Thomas J. Watson Research Center, NY, USA. Correspondence to: Quoc Tran-Dinh <quoctd@email.unc.edu>.

Proceedings of the 37^{th} International Conference on Machine Learning, Vienna, Austria, PMLR 119, 2020. Copyright 2020 by the author(s).

Assumption 1.2. There exist $\sigma_F, \sigma_D \in [0, +\infty)$ such that the variance of F and F' is uniformly bounded, i.e., $\mathbb{E}_\xi\left[\|\mathbf{F}(x,\xi) - F(x)\|^2\right] \le \sigma_F^2$ and $\mathbb{E}_\xi\left[\|\mathbf{F}'(x,\xi) - F'(x)\|^2\right] \le \sigma_D^2$, respectively. In the finite sum case (2), we again impose stronger conditions $\|F_i(x) - F(x)\| \le \sigma_F$ and $\|F_i'(x) - F'(x)\| \le \sigma_D$ for all $x \in \mathbb{R}^p$ and for all $i \in [n]$.

Assumptions 1.1 and 1.2 are standard and cover a wide class of models in practice as opposed to existing works. The stronger assumptions imposed on (2) allow us to develop adaptive subsampling schemes later.

Related work. Problem (1) or (2) has been widely studied in the literature under both deterministic (including the finite-sum (2) and $n = 1$) and stochastic settings, see, e.g., (Drusvyatskiy & Paquette, 2019; Duchi & Ruan, 2018; Lewis & Wright, 2016; Nesterov, 2007; Tran-Dinh & Diehl, 2011; Wang et al., 2017a). If $q = 1$ and $\phi(u) = u$, then (1) reduces to the standard stochastic optimization model studied in, e.g. Ghadimi & Lan (2016); Pham et al. (2020). In the deterministic setting, the common method to solve (1) is the *Gauss-Newton* (GN) scheme, which is also known as the *prox-linear method*. This method has been studied in several papers, including Drusvyatskiy & Paquette (2019); Duchi & Ruan (2018); Lewis & Wright (2016); Nesterov (2007); Tran-Dinh & Diehl (2011). In such settings, GN only requires Assumption 1.1 to have global convergence guarantees (Drusvyatskiy & Paquette, 2019; Nesterov, 2007).

In the stochastic setting of the form (1), Wang et al. (2017a;b) proposed stochastic compositional gradient descent methods to solve more general forms than (1), but they required a set of stronger assumptions than Assumptions 1.1 and 1.2, including the smoothness of ϕ. These methods eventually belong to a gradient-based class. Other works in this direction include Lian et al. (2017); Yu & Huang (2017); Yang et al. (2019); Liu et al. (2017); Xu & Xu (2019), which also rely on a similar approach. Together with algorithms, convergence guarantees and stochastic oracle complexity bounds have also been estimated. For instance, Wang et al. (2017a) estimates $\mathcal{O}\left(\varepsilon^{-8}\right)$ oracle complexity for solving (1), while it is improved to $\mathcal{O}\left(\varepsilon^{-4.5}\right)$ in Wang et al. (2017b). Recent works such as Zhang & Xiao (2019a) further improve the complexity to $\mathcal{O}\left(\varepsilon^{-3}\right)$. However, these methods are completely different from GN and require much stronger assumptions, including the smoothness of ϕ and F.

One main challenge to design algorithms for solving (1) is the bias of stochastic estimators. Some researchers have tried to remedy this issue by proposing more sophisticated sampling schemes, see, e.g., Blanchet et al. (2017). Other works relies on biased estimators but using variance reduction techniques, e.g., Zhang & Xiao (2019a).

Challenges. The stochastic formulation (1) creates several challenges for developing numerical methods. First, it is often nonconvex. Many papers consider special cases when Ψ is convex. This only holds if ϕ is convex and F is linear, or ϕ is convex and monotone and F is convex or concave. Clearly, such a setting is almost unrealistic or very limited. One can assume weak convexity of Ψ and add a regularizer to make the resulting problem convex but this completely changes the model. Second, ϕ is often non-smooth such as norm, penalty, or gauge functions. This prevents the use of gradient-based methods. Third, even when both ϕ and F are smooth, to guarantee Lipschitz continuity of $\nabla\Psi$, it requires simultaneously F, F', ϕ, and $\nabla\phi$ to be Lipschitz continuous. This condition is very restrictive and often requires additional bounded constraints or bounded domain assumption. Otherwise, it fails to hold even for bilinear functions. Finally, in stochastic settings, it is very challenging to form unbiased estimate for gradients or subgradients of Ψ, making classical stochastic-based method inapplicable.

Our approach and contribution. Our main motivation is to overcome the above challenges by following a different approach.[1] We extend the GN method from the deterministic setting (Lewis & Wright, 2016; Nesterov, 2007) to the stochastic setting (1). Our methods can be viewed as inexact variants of GN using stochastic estimators for both function values $F(x)$ and its Jacobian $F'(x)$. This approach allows us to cover a wide class of (1), while only requires standard assumptions as Assumptions 1.1 and 1.2. Our contribution can be summarized as follows:

(a) We develop an inexact GN framework to solve (1) and (2) using inexact estimations of F and its Jacobian F'. This framework is independent of approximation schemes for generating approximate estimators. We characterize approximate stationary points of (1) and (2) via prox-linear gradient mappings. Then, we prove global convergence guarantee of our method to a stationary point under appropriate inexact computation.

(b) We analyze stochastic oracle complexity of our GN algorithm when mini-batch stochastic estimators are used. We separate our analysis into two cases. The first variant is to solve (1), where we obtain convergence guarantee in expectation. The second variant is to solve (2), where we use adaptive mini-batches and obtain convergence guarantee with high probability.

(c) We also provide oracle complexity of this algorithm when mini-batch SARAH estimators in Nguyen et al. (2017; 2019) are used for both (1) and (2). Under an additional mild assumption, this estimator significantly improves the oracle complexity by an order of ε compared to the mini-batch stochastic one.

We believe that our methods are the first ones to achieve global convergence rates and stochastic oracle complexity

[1] When this paper was under review, Zhang & Xiao (2020) was brought to our attention, which presents similar methods.

for solving (1) and (2) under standard assumptions. It is completely different from existing works such as Wang et al. (2017a;b); Lian et al. (2017); Yu & Huang (2017); Yang et al. (2019); Zhang & Xiao (2019a), where we only use Assumptions 1.1 and 1.2, while not imposing any special structure on ϕ and F, including smoothness. When using SARAH estimators, we impose the Lipschitz continuity of F to achieve better oracle complexity. This additional assumption is still much weaker than the ones used in existing works. However, without this assumption, our GN scheme with SARAH estimators still converges (see Remark 4.1).

Content. Section 2 recalls some mathematical tools. Section 3 develops an inexact GN framework. Sections 4 analyzes convergence and complexity of the two stochastic GN variants using different stochastic estimators. Numerical examples are given in Section 5. All the proofs and discussion are deferred to Supplementary Document (Supp. Doc.).

2. Background and Mathematical Tools

We first characterize the optimality condition of (1). Next, we recall the prox-linear mapping of the compositional function $\Psi(x) := \phi(F(x))$ and its properties.

Basic notation. We work with Euclidean spaces $\mathbb{R}^p$ and $\mathbb{R}^q$. Given a convex set $\mathcal{X}$, $\mathrm{dist}\,(u, \mathcal{X}) := \inf_{x\in\mathcal{X}} \|u - x\|$ denotes the Euclidean distance from u to $\mathcal{X}$. For a convex function f, we denote ∂f its subdifferential, ∇f its gradient, and f^* its Fenchel conjugate. For a smooth function $F : \mathbb{R}^p \to \mathbb{R}^q$, $F'(\cdot)$ denotes its Jacobian. For vectors, we use Euclidean norms, while for matrices, we use spectral norms, i.e., $\|X\| := \sigma_{\max}(X)$. $\lfloor\cdot\rfloor$ stands for number rounding.

2.1. Exact and Approximate Stationary Points

The optimality condition of (1) can be written as

$$\begin{aligned} &0 \in \partial\Psi(x^\star) \equiv F'(x^\star)^\top \partial\phi(F(x^\star)),\\ &\text{or equivalently } \ \mathrm{dist}\,(0, \partial\Psi(x^\star)) = 0. \end{aligned} \tag{5}$$

Any $x^\star$ satisfying (5) is called a stationary point of (1) or (2).

Since ϕ is convex, let ϕ^* be its Fenchel conjugate and $y^\star \in \partial\phi(F(x^\star))$. Then, (5) can be rewritten as

$$0 = F'(x^\star)^\top y^\star \ \text{ and } \ 0 \in -F(x^\star) + \partial\phi^*(y^\star). \tag{6}$$

Now, if we define

$$\mathcal{E}(x, y) := \|F'(x)^\top y\| + \mathrm{dist}\,(0, -F(x) + \partial\phi^*(y))\,, \tag{7}$$

then the optimality condition (5) of (1) or (2) becomes

$$\mathcal{E}(x^\star, y^\star) = 0. \tag{8}$$

Note that once a stationary point $x^\star$ is available, we can compute $y^\star$ as any element $y^\star \in \partial\phi(F(x^\star))$ of $\phi \circ F$.

In practice, we can only find an approximate stationary point $\hat{x}$ and its dual $\hat{y}$ such that $(\hat{x}, \hat{y})$ approximates $(x^\star, y^\star)$ of (1) or (2) up to a given accuracy $\varepsilon \geq 0$ as follows:

Definition 2.1. Given $\varepsilon > 0$, we call $\hat{x} \in \mathbb{R}^p$ an ε-stationary point of (1) if there exists $\hat{y} \in \mathbb{R}^q$ such that

$$\mathcal{E}(\hat{x}, \hat{y}) \leq \varepsilon, \tag{9}$$

where $\mathcal{E}(\cdot)$ is defined by (7). This condition can be characterized in expectation, where $\mathbb{E}[\cdot]$ is taken over all the randomness generated by the problem and the corresponding stochastic algorithm, or with high probability $1 - \delta$. Such guarantees will be specified in the sequel.

2.2. Prox-Linear Operator and Its Properties

(a) *Prox-linear operator.* Since we assume that the Jacobian $F'(\cdot)$ of F is Lipschitz continuous with a Lipschitz constant $L_F \in (0, +\infty)$, and ϕ is M_ϕ-Lipschitz continuous as in Assumption 1.1, we have (see Supp. Doc. A):

$$\phi(F(z)) \leq \phi(F(x) + F'(x)(z - x)) + \frac{M_\phi L_F}{2}\|z - x\|^2, \tag{10}$$

for all $z, x \in \mathbb{R}^p$. Given $x \in \mathbb{R}^p$, let $\widetilde{F}(x) \approx F(x)$ and $\widetilde{J}(x) \approx F'(x)$ be a deterministic or stochastic approximation of $F(x)$ and its Jacobian $F'(x)$, respectively. We consider the following approximate prox-linear model:

$$\begin{aligned} \widetilde{T}_M(x) := \mathop{\arg\min}_{z\in\mathbb{R}^p} \big\{ &\widetilde{\mathcal{Q}}_M(z; x) := \phi(\widetilde{F}(x) + \widetilde{J}(x)(z - x))\\ &+ \tfrac{M}{2}\|z - x\|^2 \big\}, \end{aligned} \tag{11}$$

where $M > 0$ is a given constant. As usual, if $\widetilde{F}(x) = F(x)$ and $\widetilde{J}(x) = F'(x)$, then

$$\begin{aligned} T_M(x) := \mathop{\arg\min}_{z\in\mathbb{R}^p} \big\{ &\mathcal{Q}_M(z; x) := \phi(F(x) + F'(x)(z - x))\\ &+ \tfrac{M}{2}\|z - x\|^2 \big\} \end{aligned} \tag{12}$$

is the exact prox-linear operator of Ψ. In this context, we also call $\widetilde{T}_M(\cdot)$ an approximate prox-linear operator of Ψ.

(b) *Prox-gradient mapping.* We also define the prox-gradient mapping and its approximation, respectively as

$$\left\{ \begin{array}{ll} G_M(x) & := M(x - T_M(x)),\\ \widetilde{G}_M(x) & := M(x - \widetilde{T}_M(x)). \end{array} \right. \tag{13}$$

Clearly if $\|G_M(x)\| = 0$, then $x = T_M(x)$ and x is a stationary point of (1). In our context, we can only compute $\widetilde{G}_M(x)$ as an approximation of $G_M(x)$.

(c) *Characterizing approximate stationary points.* The following lemma bounds the optimality error $\mathcal{E}(\cdot)$ defined by (7) via the approximate prox-gradient mapping $\widetilde{G}_M(x)$.

Lemma 2.1. *Let $\widetilde{T}_M(x)$ be computed by (11) and $\widetilde{G}_M(x)$ be defined by (13). Then, $\mathcal{E}(\widetilde{T}_M(x), y)$ of (1) defined by (7) with $y \in \partial\phi(F(\widetilde{T}_M(x)))$ is bounded by*

$$\mathcal{E}(\widetilde{T}_M(x), y) \le \left(1+\tfrac{M_\phi L_F}{M}\right)\|\widetilde{G}_M(x)\| + \tfrac{(1+L_F)}{2M^2}\|\widetilde{G}_M(x)\|^2 + \|\widetilde{F}(x) - F(x)\| + \tfrac{1}{2}\|\widetilde{J}(x) - F'(x)\|^2. \quad (14)$$

Clearly, if we use exact oracles $\widetilde{F}(x) = F(x)$ and $\widetilde{J}(x) = F'(x)$, then $\mathcal{E}(\widetilde{T}_M(x), y)$ is reduced to

$$\mathcal{E}(T_M(x), y) \le \left(1+\tfrac{M_\phi L_F}{M}\right)\|G_M(x)\| + \tfrac{(1+L_F)}{2M^2}\|G_M(x)\|^2.$$

Alternatively, from (14), if we can guarantee $\|\widetilde{F}(x) - F(x)\| \le \mathcal{O}(\varepsilon)$, $\|\widetilde{J}(x) - F'(x)\| \le \mathcal{O}(\sqrt{\varepsilon})$, and $\|\widetilde{G}_M(x)\| \le \mathcal{O}(\varepsilon)$, then

$$\mathcal{E}(\widetilde{T}_M(x), y) \le \mathcal{O}(\varepsilon),$$

which shows that $\widetilde{T}_M(x)$ is a $\mathcal{O}(\varepsilon)$-stationary point of (1) in the sense of Definition 2.1. Our goal is to approximate F and F' and compute $\widetilde{G}_M(x)$ to guarantee these conditions.

3. Inexact Gauss-Newton Framework

In this section, we develop a conceptual inexact Gauss-Newton (iGN) framework for solving (1) and (2).

3.1. Descent Property and Approximate Conditions

Lemma 3.1 provides a key bound regarding (11), which will be used for convergence analysis of our algorithms.

Lemma 3.1. *Let Assumption 1.1 hold and $\widetilde{T}_M(x)$ be computed by (11). Then, for any $\beta_d > 0$, we also have*

$$\phi(F(\widetilde{T}_M(x))) \le \phi(F(x)) + 2L_\phi\|F(x) - \widetilde{F}(x)\| + \tfrac{M_\phi}{2\beta_d}\|F'(x) - \widetilde{J}(x)\|_F^2 - \tfrac{(2M - M_\phi L_F - \beta_d L_\phi)}{2}\|\widetilde{T}_M(x) - x\|^2. \quad (15)$$

Since we approximate both F and its Jacobian F' in our prox-linear model (11), we assume that this approximation satisfies one of the following two conditions:

- **Condition 1:** Given a tolerance $\varepsilon > 0$ and $M > \frac{1}{2}M_\phi(L_F + \beta_d)$, at each iterate $x_t \in \mathbb{R}^p$, it holds that

$$\begin{cases} \|\widetilde{F}(x_t) - F(x_t)\| \le \frac{C_g\varepsilon^2}{16M_\phi M^2}, \\ \|\widetilde{J}(x_t) - F'(x_t)\| \le \frac{\sqrt{\beta_d C_g}\varepsilon}{M\sqrt{2M_\phi}}, \end{cases} \quad (16)$$

  where $C_g := 2M - M_\phi(L_F + \beta_d) > 0$.

- **Condition 2:** Given $C_f > 0$, $C_d > 0$, and $\beta_d > 0$, let $C_g := 2M - M_\phi(L_F + \beta_d)$ and $C_a := 2M - M_\phi\big(L_F + \beta_d + 2\sqrt{C_f} + \frac{C_d}{2\beta_d}\big)$ such that $C_a > 0$. For $x_0 \in \mathbb{R}^p$, we assume that

$$\begin{cases} \|\widetilde{F}(x_0) - F(x_0)\| \le \frac{C_a\varepsilon^2}{16M_\phi M^2} \\ \|\widetilde{J}(x_0) - F'(x_0)\| \le \frac{\sqrt{\beta_d C_a}\varepsilon}{M\sqrt{2M_\phi}}, \end{cases} \quad (17)$$

  while, for any iterate $x_t \in \mathbb{R}^p$ ($t \ge 1$), we assume that

$$\begin{cases} \|\widetilde{F}(x_t) - F(x_t)\| \le \sqrt{C_f}\,\|x_t - x_{t-1}\|^2, \\ \|\widetilde{J}(x_t) - F'(x_t)\| \le \sqrt{C_d}\,\|x_t - x_{t-1}\|. \end{cases} \quad (18)$$

The condition (16) assumes that both $\widetilde{F}$ and $\widetilde{J}$ should respectively well approximate F and F' up to a given accuracy ε. Here, the function value F must have higher accuracy than its Jacobian F'. The condition (18) is adaptive, which depends on the norm $\|x_t - x_{t-1}\|$ of the iterates x_t and x_{t-1}. This condition is less conservative than (16).

3.2. The Inexact Gauss-Newton Algorithm

We first present a conceptual stochastic Gauss-Newton method as described in Algorithm 1.

Algorithm 1 (Inexact Gauss-Newton **(iGN)**)

1: **Initialization:** Choose $x_0 \in \mathbb{R}^p$ and $M > 0$.
2: **For** $t := 0, \cdots, T$ **do**
3: Form $\widetilde{F}(x_t)$ and $\widetilde{J}(x_t)$ satisfying either (16) or (18).
4: Update $x_{t+1} := \widetilde{T}_M(x_t)$ based on (11).
5: **End For**

Algorithm 1 remains conceptual since we have not specified how to form $\widetilde{F}(x_t)$ and $\widetilde{J}(x_t)$.

3.3. Convergence Analysis

Let us first state the convergence of Algorithm 1 under **Conditon 1** or **Condition 2** in the following theorem.

Theorem 3.1. *Assume that Assumptions 1.1 and 1.2 are satisfied. Let $\{x_t\}$ be generated by Algorithm 1 to solve either (1) or (2). Then, the following statements hold:*

(a) *If (16) holds for some $\varepsilon \ge 0$, then*

$$\frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \le \frac{2M^2\left[\Psi(x_0) - \Psi^\star\right]}{C_g(T+1)} + \frac{\varepsilon^2}{2}, \quad (19)$$

where $C_g := 2M - M_\phi(L_F + \beta_d)$ with $M > \frac{1}{2}M_\phi(L_F + \beta_d)$.

(b) *If (17) and (18) hold for given $C_a > 0$, then*

$$\frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \le \frac{2M^2\left[\Psi(x_0) - \Psi^\star\right]}{C_a(T+1)} + \frac{\varepsilon^2}{2}. \quad (20)$$

Consequently, the total number of iterations T to achieve $\frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \leq \varepsilon^2$ is at most

$$T := \left\lfloor \frac{4M^2\left[\Psi(x_0)-\Psi^\star\right]}{D\varepsilon^2} \right\rfloor = \mathcal{O}\left(\frac{1}{\varepsilon^2}\right),$$

where $D := C_g$ for (a) *and $D := C_a$ for* (b).

Remark 3.1. The guarantee $\frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \leq \varepsilon^2$ implies that $\liminf_{t\to\infty,\varepsilon\downarrow 0^+}\|\widetilde{G}_M(x_t)\| = 0$. That is there exists subsequence x_{t_k} of $\{x_t\}$ such that $\|\widetilde{G}_M(x_{t_k})\| \to 0$ as $k \to +\infty$ and $\varepsilon \to 0$.

4. Stochastic Gauss-Newton Methods

4.1. SGN with Mini-Batch Stochastic Estimators

As a natural instance of Algorithm 1, we propose to approximate $F(x_t)$ and $F'(x_t)$ in Algorithm 1 by mini-batch stochastic estimators as:

$$\left\{\begin{array}{lcl} \widetilde{F}(x_t) & := & \frac{1}{b_t}\sum_{\xi_i\in\mathcal{B}_t}\mathbf{F}(x_t,\xi_i), \\ \widetilde{J}(x_t) & := & \frac{1}{\hat{b}_t}\sum_{\zeta_j\in\hat{\mathcal{B}}_t}\mathbf{F}'(x_t,\zeta_j), \end{array}\right. \tag{21}$$

where the mini-batches $\mathcal{B}_t$ and $\hat{\mathcal{B}}_t$ are not necessarily independent, $b_t := |\mathcal{B}_t|$, and $\hat{b}_t := |\hat{\mathcal{B}}_t|$. Using (21) we prove our first result in expectation on stochastic oracle complexity of Algorithm 1 for solving (1).

In practice, we may not need to explicitly form $\widetilde{J}(x_t)$, but its matrix-vector product $\widetilde{J}(x_t)d$ for some vector d, when evaluating the prox-linear operator $\widetilde{\mathcal{T}}_M(x_t)$. This requires $\mathbf{F}'(x_t,\zeta_j)d$, which can be evaluated efficiently by using, e.g., automatic differentiation techniques.

Theorem 4.1. *Suppose that Assumptions 1.1 and 1.2 hold for (1). Let $\widetilde{F}_t$ and $\widetilde{J}_t$ defined by (21) be mini-batch stochastic estimators of $F(x_t)$ and $F'(x_t)$, respectively. Let $\{x_t\}$ be generated by Algorithm 1 (called **SGN**) to solve (1). Assume that b_t and $\hat{b}_t$ in (21) are chosen as*

$$\left\{\begin{array}{lcll} b_t & := & \left\lfloor \frac{256M_\phi^2M^4\sigma_F^2}{C_g^2\varepsilon^4} \right\rfloor & = \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^4}\right) \\ \hat{b}_t & := & \left\lfloor \frac{2M_\phi M^2\sigma_D^2}{\beta_d C_g\varepsilon^2} \right\rfloor & = \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^2}\right), \end{array}\right. \tag{22}$$

for some constant $C_f > 0$ and $C_d > 0$. Furthermore, let $\widehat{x}_T$ be chosen uniformly at random in $\{x_t\}_{t=0}^T$ as the output of Algorithm 1 after T iterations. Then

$$\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] \leq \frac{2M^2\left[\Psi(x_0)-\Psi^\star\right]}{C_g(T+1)} + \frac{\varepsilon^2}{2}, \tag{23}$$

where $C_g := 2M - M_\phi(L_F+\beta_d) > 0$.

Moreover, the number $\mathcal{T}_f$ of function evaluations $\mathbf{F}(x_t,\xi)$ and the number $\mathcal{T}_d$ of Jacobian evaluations $\mathbf{F}'(x_t,\zeta)$ to achieve $\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] \leq \varepsilon^2$ do not exceed

$$\left\{\begin{array}{lcl} \mathcal{T}_f & := & \left\lfloor \frac{1024M^6M_\phi^2\sigma_F^2\left[\Psi(x_0)-\Psi^\star\right]}{C_g^3\varepsilon^6} \right\rfloor = \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^6}\right), \\ \mathcal{T}_d & := & \left\lfloor \frac{8M^4M_\phi\sigma_D^2\left[\Psi(x_0)-\Psi^\star\right]}{\beta_d C_g^2\varepsilon^4} \right\rfloor = \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^4}\right). \end{array}\right. \tag{24}$$

Note that if we replace b_t and $\hat{b}_t$ in (22) by $\min\{b_t, n\}$ and $\min\{\hat{b}_t, n\}$, respectively, then the result of Theorem 4.1 still holds for (2) since it is a special case of (1).

Now, we derive the convergence result of Algorithm 1 for solving (2) using adaptive mini-batches. However, our convergence guarantee is obtained with high probability.

Theorem 4.2. *Suppose that Assumptions 1.1 and 1.2 hold for (2). Let $\widetilde{F}_t$ and $\widetilde{J}_t$ defined by (21) be mini-batch stochastic estimators to approximate $F(x_t)$ and $F'(x_t)$, respectively. Let $\{x_t\}$ be generated by Algorithm 1 for solving (2). Assume that b_t and $\hat{b}_t$ in (21) are chosen such that $b_t := \min\{n, \bar{b}_t\}$ and $\hat{b}_t := \min\{n, \hat{\bar{b}}_t\}$ for $t \geq 0$, with*

$$\left\{\begin{array}{l} \bar{b}_0 := \left\lfloor \frac{32M_\phi M^2\sigma_F\left(48\sigma_F M_\phi M^2 + C_a\varepsilon^2\right)}{3C_a^2\varepsilon^4} \cdot \log\left(\frac{p+1}{\delta}\right) \right\rfloor \\ \hat{\bar{b}}_0 := \left\lfloor \frac{4M\sqrt{2M_\phi}\sigma_D\left(3M\sqrt{2M_\phi}\sigma_D + \sqrt{\beta_d C_a}\varepsilon\right)}{\beta_d C_a\varepsilon^2} \cdot \log\left(\frac{p+q}{\delta}\right) \right\rfloor, \\ \bar{b}_t := \left\lfloor \frac{\left(6\sigma_F^2 + 2\sigma_F\sqrt{C_f}\|x_t - x_{t-1}\|^2\right)}{3C_f^2\|x_t - x_{t-1}\|^4} \cdot \log\left(\frac{p+1}{\delta}\right) \right\rfloor \quad (t \geq 1) \\ \hat{\bar{b}}_t := \left\lfloor \frac{\left(6\sigma_D^2 + 2\sigma_D\sqrt{C_d}\|x_t - x_{t-1}\|\right)}{3C_d\|x_t - x_{t-1}\|^2} \cdot \log\left(\frac{p+q}{\delta}\right) \right\rfloor \quad (t \geq 1), \end{array}\right. \tag{25}$$

*for $\delta \in (0,1)$, and C_f, C_d, and C_a given in **Condition 2**. Then, with probability at least $1-\delta$, the bound (20) in Theorem 3.1 still holds.*

Moreover, the total number $\mathcal{T}_f$ of stochastic function evaluations $\mathbf{F}(x_t,\xi)$ and the total number $\mathcal{T}_d$ of stochastic Jacobian evaluations $\mathbf{F}'(x_t,\zeta)$ to guarantee $\frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \leq \varepsilon^2$ do not exceed

$$\left\{\begin{array}{lcl} \mathcal{T}_f & := & \mathcal{O}\left(\frac{\sigma_F^2\left[\Psi(x_0)-\Psi^\star\right]}{\varepsilon^6}\cdot\log\left(\frac{p+1}{\delta}\right)\right), \\ \mathcal{T}_d & := & \mathcal{O}\left(\frac{\sigma_D^2\left[\Psi(x_0)-\Psi^\star\right]}{\varepsilon^4}\cdot\log\left(\frac{p+q}{\delta}\right)\right). \end{array}\right. \tag{26}$$

To the best of our knowledge, the oracle complexity bounds stated in Theorems 4.1 and 4.2 are the first results for the stochastic Gauss-Newton methods described in Algorithm 1 under Assumptions 1.1 and 1.2. Whereas there exist several methods for solving (1), these algorithms are either not in the form of GN schemes as ours or rely on a different set of assumptions. For instance, Duchi & Ruan (2018); Duchi et al. (2011) considers a different model and uses stochastic subgradient methods, while Zhang & Xiao (2019a;b) directly applies a variance reduction gradient descent method and requires a stronger set of assumptions.

4.2. SGN with SARAH Estimators

Algorithm 1 with mini-batch stochastic estimators (21) has high oracle complexity bounds when ε is sufficiently small, especially for function evaluations $\mathbf{F}(\cdot,\xi)$. We attempt to reduce this complexity by exploiting a biased estimator called SARAH in Nguyen et al. (2017) in this subsection.

More concretely, we approximate $F(x_t)$ and $F'(x_t)$ by using the following SARAH estimators, respectively:

$$\begin{cases} \widetilde{F}_t := \widetilde{F}_{t-1} + \frac{1}{b_t}\sum_{\xi_j\in\mathcal{B}_t}\left(\mathbf{F}(x_t,\xi_j) - \mathbf{F}(x_{t-1},\xi_j)\right), \\ \widetilde{J}_t := \widetilde{J}_{t-1} + \frac{1}{\hat{b}_t}\sum_{\xi_i\in\hat{\mathcal{B}}_t}\left(\mathbf{F}'(x_t,\zeta_i) - \mathbf{F}'(x_{t-1},\zeta_i)\right), \end{cases} \tag{27}$$

where the snapshots $\widetilde{F}_0$ and $\widetilde{J}_0$ are given, and $\mathcal{B}_t$ and $\hat{\mathcal{B}}_t$ are two mini-batches of size $b_t := |\mathcal{B}_t|$ and $\hat{b}_t := |\hat{\mathcal{B}}_t|$.

Using both the standard stochastic estimators (21) and these SARAH estimators (27), we modify Algorithm 1 to obtain the following double-loop variant as in Algorithm 2.

Algorithm 2 (SGN with SARAH estimators **(SGN2)**)

Initialization: Choose $\widetilde{x}^0 \in \mathbb{R}^p$ and $M > 0$.
For $s := 1, \cdots, S$ **do**
  Generate mini-batches $\mathcal{B}_s$ (size b_s) and $\hat{\mathcal{B}}_s$ (size $\hat{b}_s$).
  Evaluate $F_0^{(s)}$ and $J_0^{(s)}$ at $x_0^{(s)} := \widetilde{x}^{s-1}$ from (21).
  Update $x_1^{(s)} := \widetilde{T}_M(x_0^{(s)})$ based on (11).
  Inner Loop: For $t := 1, \cdots, m$ **do**
    Generate mini-batches $\mathcal{B}_t^{(s)}$ and $\hat{\mathcal{B}}_t^{(s)}$.
    Evaluate $F_t^{(s)}$ and $J_t^{(s)}$ from (27).
    Update $x_{t+1}^{(s)} := \widetilde{T}_M(x_t^{(s)})$ based on (11).
  End of Inner Loop
  Set $\widetilde{x}^s := x_{m+1}^{(s)}$.
End For

In Algorithm 2, every outer iteration s, we take a snapshot $\widetilde{x}^s$ using (21). Then, we run Algorithm 2 up to m iterations in the inner loop t but using SARAH estimators (27). Unlike (21), we are unable to exploit matrix-vector products for $\widetilde{J}_t$ in (27) due to its dependence on $\widetilde{J}_{t-1}$.

Let us first prove convergence and oracle complexity estimates in expectation of Algorithm 2 for solving (1). However, we require an additional assumption for this case:

Assumption 4.1. $\mathbf{F}$ is M_F-average Lipschitz continuous, i.e. $\mathbb{E}_\xi\left[\|\mathbf{F}(x,\xi) - \mathbf{F}(y,\xi)\|^2\right] \le M_F^2\|x-y\|^2$ for all x, y.

Though Assumption 4.1 is relatively strong, it has been used in several models, including neural network training under a bounded weight assumption.

Given a tolerance $\varepsilon > 0$ and $C > 0$, we first choose $M > 0$, $\delta_d > 0$, and two constants $\gamma_1 > 0$ and $\gamma_2 > 0$ such that

$$\begin{cases} \theta_F := 2M - M_\phi(L_F + \delta_d) - \gamma_1 M_F^2 - \gamma_2 L_F^2 > 0, \\ m := \left\lfloor \frac{8\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F C\varepsilon} \right\rfloor. \end{cases} \tag{28}$$

Next, we choose the mini-batch sizes of $\mathcal{B}_s$, $\hat{\mathcal{B}}_s$, $\mathcal{B}_t^{(s)}$, and $\hat{\mathcal{B}}_t^{(s)}$, respectively as follows:

$$\begin{cases} b_s := \frac{2CM_\phi^2\sigma_F^2}{\theta_F^2\varepsilon^3} & \hat{b}_s := \frac{4CM_\phi\sigma_D^2}{\theta_F\delta_D\varepsilon} \\ b_t^{(s)} := \frac{8M_\phi^2(m+1-t)}{\theta_F\gamma_1\varepsilon^2} & \hat{b}_t^{(s)} := \frac{M_\phi(m+1-t)}{\gamma_2\delta_d}. \end{cases} \tag{29}$$

Then, the following theorem states the convergence and oracle complexity bounds of Algorithm 2.

Theorem 4.3. *Suppose that Assumptions 1.1 and 1.2, and 4.1 are satisfied for (1). Let $\{x_t^{(s)}\}_{t=0\to m}^{s=1\to S}$ be generated by Algorithm 2 to solve (1). Let θ_F and m be chosen by (28), and the mini-batches b_s, $\hat{b}_s$, $b_t^{(s)}$, and $\hat{b}_t^{(s)}$ be set as in (29). Assume that the output $\widehat{x}_T$ of Algorithm 2 is chosen uniformly at random in $\{x_t^{(s)}\}_{t=0\to m}^{s=1\to S}$. Then:*

(a) *The following bound holds*

$$\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\mathbb{E}\left[\|\widetilde{G}_M(x_t)\|^2\right] \le \varepsilon^2. \tag{30}$$

(b) *The total number of iterations T to obtain $\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] \le \varepsilon^2$ is at most*

$$T := S(m+1) = \left\lfloor \frac{8M^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F\varepsilon^2} \right\rfloor = \mathcal{O}\left(\frac{1}{\varepsilon^2}\right).$$

Moreover, the total stochastic oracle calls $\mathcal{T}_f$ and $\mathcal{T}_d$ for evaluating stochastic estimators of $\mathbf{F}(x_t,\xi)$ and its Jacobian $\mathbf{F}'(x_t,\zeta)$, respectively do not exceed:

$$\begin{cases} \mathcal{T}_f & := \mathcal{O}\left(\frac{M_\phi^2\sigma_F^2}{\theta_F^2\varepsilon^4} + \frac{M^4M_\phi^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F^2\varepsilon^5}\right) \\ \mathcal{T}_d & := \mathcal{O}\left(\frac{M_\phi\sigma_D^2}{\theta_F\varepsilon^2} + \frac{M^2M_\phi\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F\varepsilon^3}\right). \end{cases} \tag{31}$$

Finally, we show that x_t computed by our methods is indeed an approximate stationary point of (1) or (2).

Corollary 4.1. *If x_t satisfies $\|\widetilde{G}_M(x_t)\| \le \varepsilon$ for given $\varepsilon > 0$, then under either **Condition 1** or **Condition 2**, and for any $y_t \in \partial\phi(F(x_{t+1}))$, we have $\mathcal{E}(x_{t+1}, y_t) \le \mathcal{O}(\varepsilon)$. Consequently, x_{t+1} is a $\mathcal{O}(\varepsilon)$-stationary point of (1) or (2).*

Proof. From Lemma 2.1, we have

$$\mathcal{E}(x_{t+1},y_t) \le \left(1 + \frac{M_\phi L_F}{M}\right)\|\widetilde{G}_M(x_t)\| + \tfrac{1}{2}\|\widetilde{J}_t - F'(x_t)\|^2 + \frac{(1+L_F)}{2M^2}\|\widetilde{G}_M(x_t)\|^2 + M_\phi\|\widetilde{F}_t - F(x_t)\|.$$

Under either **Condition 1** or **Condition 2**, we have $\|\widetilde{F}_t - F(x_t)\| \le \mathcal{O}\left(\varepsilon^2\right)$ and $\|\widetilde{J}_t - F'(x_t)\|^2 \le \mathcal{O}\left(\varepsilon^2\right)$. Hence, if $\|\widetilde{G}_M(x_t)\| \le \varepsilon$, then using these three bounds into the last estimate, one can show that $\mathcal{E}(x_{t+1}, y_t) \le \mathcal{O}(\varepsilon)$. Consequently, x_{t+1} is a $\mathcal{O}(\varepsilon)$-stationary point of (1) or (2). $\square$

Remark 4.1 (**Algorithm 2 without Assumption 4.1**). We claim that Algorithm 2 still converges without Assumption 4.1. However, its oracle complexity remains $\mathcal{O}\left(\sigma_F^2\varepsilon^{-6}\right)$ for F and $\mathcal{O}\left(\sigma_D^2\varepsilon^{-4}\right)$ for F' as in Algorithm 1. We therefore omit the proof of this statement.

Another main step of both Algorithms 1 and 2 is to compute $\widetilde{T}_M(x_t)$. We will provide different routines in Sup. Doc. E to efficiently compute $\widetilde{T}_M(x_t)$.

4.3. Extension to The Regularization Setting (3)

It is straight forward to extend our methods to handle a regularizer g as in (3). If g is nonsmooth and convex, then we can modify (11) as follows:

$$\widetilde{T}_M(x_t) := \arg\min_{z\in\mathbb{R}^p}\Big\{\widetilde{\mathcal{Q}}_M(z;x_t) := \phi(\widetilde{F}(x_t) + \widetilde{J}(x_t)(z-x_t)) + g(z) + \tfrac{M}{2}\|z-x_t\|^2\Big\}. \tag{32}$$

Then, we obtain variants of Algorithms 1 and 2 for solving (3), where our theoretical guarantees in this paper remain preserved. This subproblem can efficiently be solved by primal-dual methods as presented in Supp. Doc. E. If g is L_g-smooth, then we can replace g in (11) by its quadratic surrogate $g(x_t) + \langle\nabla g(x_t), z - x_t\rangle + \frac{L_g}{2}\|z-x_t\|^2$.

5. Numerical Experiments

We conduct two numerical experiments to evaluate the performance of Algorithm 1 (SGN) and Algorithm 2 (SGN2). Further details of our experiments are in Supp. Doc. F.

5.1. Stochastic Nonlinear Equations

We consider a nonlinear equation: $\mathbb{E}_\xi[\mathbf{F}(x,\xi)] = 0$ as the expectation of a stochastic function $\mathbf{F}: \mathbb{R}^p \times \Omega \to \mathbb{R}^q$. This equation can be viewed as a natural extension of nonlinear equations from a deterministic setting to a stochastic setting, including stochastic dynamical systems and PDEs. It can also present as the first-order optimality condition $\mathbb{E}_\xi[\nabla\mathbf{G}(x,\xi)] = 0$ of a stochastic optimization problem $\min_x \mathbb{E}_\xi[\mathbf{G}(x,\xi)]$. Moreover, it can be considered as a special case of stochastic variational inequality in the literature, see, e.g., Rockafellar & Wets (2017).

Instead of directly solving $\mathbb{E}_\xi[\mathbf{F}(x,\xi)] = 0$, we can formulate it into the following minimization problem:

$$\min_{x\in\mathbb{R}^p}\Big\{\Psi(x) := \|\mathbb{E}_\xi[\mathbf{F}(x,\xi)]\|\Big\}, \tag{33}$$

where $\mathbf{F}(x,\xi) := (\mathbf{F}_1(x,\xi), \mathbf{F}_2(x,\xi), \cdots, \mathbf{F}_q(x,\xi))^\top$ such that $F_j : \mathbb{R}^p \to \mathbb{R}$ is the expectation of $\mathbf{F}_j(\cdot,\xi)$, i.e., $F_j(x) := \mathbb{E}_\xi[\mathbf{F}_j(x,\xi)]$ for $j = 1, \cdots, q$, and $\|\cdot\|$ is a given norm (e.g., ℓ_2-norm or ℓ_1-norm).

Assume that we take average approximation of $\mathbb{E}_\xi[\mathbf{F}(x,\xi)]$ to obtain a finite sum $F(x) = \frac{1}{n}\sum_{i=1}^n F(x,\xi_i)$ for sufficiently large n. In the following experiments, we choose $q = 4$, and for $i = 1, \cdots, n$, we choose $\mathbf{F}_j(x,\xi_i)$ as

$$\begin{cases}\mathbf{F}_1(x,\xi_i) := (1 - \tanh(y_i(a_i^\top x + b_i)),\\ \mathbf{F}_2(x,\xi_i) := \left(1 - (1 + \exp(-y_i(a_i^\top x + b_i)))^{-1}\right)^2,\\ \mathbf{F}_3(x,\xi_i) := \log(1 + \exp(-y_i(a_i^\top x + b_i)))\\ \qquad\qquad - \log(1 + \exp(-y_i(a_i^\top x + b_i) - 1)),\\ \mathbf{F}_4(x,\xi_i) := \log(1 + (y_i(a_i^\top x + b_i) - 1)^2),\end{cases}$$

where a_i is the i-th row of an input matrix $A \in \mathbb{R}^{n\times p}$, and $y \in \{-1,1\}^n$, $b \in \mathbb{R}^n$ are two input vectors, and $\xi_i := (a_i, b_i, y_i)$. These functions were used in binary classification involving nonconvex losses, e.g., Zhao et al. (2010). Since they are nonnegative, if we use the ℓ_1-norm, then (33) can be viewed as a model average of 4 different losses in binary classification (see Supp. Doc. F).

We implement both Algorithms 1 (SGN) and 2 (SGN2) to solve (33). We also compare them with the baseline using the full samples instead of calculating $\widetilde{J}$ and $\widetilde{F}$ as in (21) and (27). We call it the deterministic GN scheme (GN).

Experiment setup. We test three algorithms on four standard datasets: `w8a`, `ijcnn1`, `covtype`, and `url_combined` from LIBSVM[2]. Further information about these dataset is described in Supp. Doc. F.

To find appropriate batch sizes for $\widetilde{J}$ and $\widetilde{F}$, we perform a grid search over different combinations of mini-batch sizes to select the best ones. More information about this process can be found in Supp. Doc. F.

We evaluate these algorithms on instances of (33) using $\phi(\cdot) = \|\cdot\|_2$. We use $M := 1$ and $\rho := 1$ for all datasets. The performance of three algorithms is shown in Figure 1 for the `w8a` and `ijcnn1` datasets. This figure depicts the relative objective residuals $\frac{\Psi(x_t) - \Psi^\star}{|\Psi^\star|}$ over the number of epochs, where $\Psi^\star$ is the lowest objective value obtained when running three algorithms until the relative residuals falls below 10^{-6}. In both cases, SGN2 works best while SGN is still much better than the baseline GN in terms of sample efficiency.

For `covtype` and `url_combined` datasets, we obverse similar behavior as shown in Figure 2, where SGN2 is more efficient than SGN, and both SGN schemes outperform GN. This experiment shows that both SGN algorithms are indeed much more sample efficient than the baseline GN algorithm.

[2] Available online at https://www.csie.ntu.edu.tw/∼cjlin/libsvm/

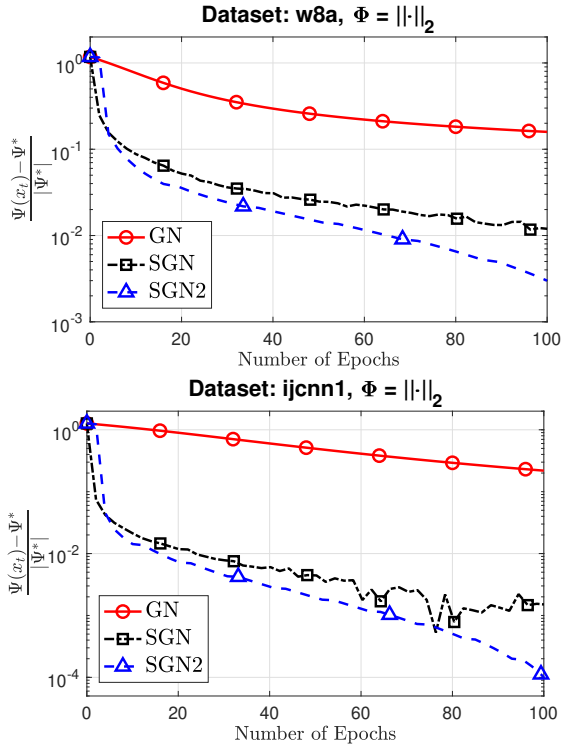

Figure 1. The performance of 3 algorithms on the `w8a` and `ijcnn1`.

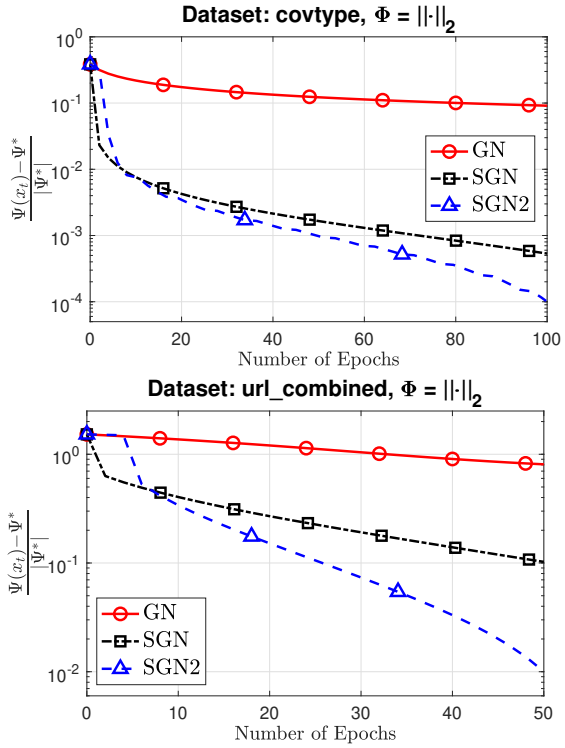

Figure 2. The performance of 3 algorithms on `covtype` and `url_combined`.

In order to compare with existing algorithms, we use a smooth objective function in (33) with a Huber loss, $\phi(u) = \frac{1}{2}u^2$ for $|u| \leq \delta$ and $\phi(u) = \delta(|u| - \frac{1}{2}\delta)$ otherwise, and $\delta := 1.0$. We implement the nested SPIDER method in Zhang & Xiao (2019a, Algorithm 3), denoted as N-SPIDER, and the stochastic compositional gradient descent in Wang et al. (2017a, Algorithm 1), denoted as SCGD.

We run 5 algorithms: GN, SGN, SGN2, N-SPIDER, and SCGD on 4 datasets as in the previous test. We choose $M := 1$ and $\rho := 1$ for all datasets. We tune the learning rate for both N-SPIDER and SCGD and finally obtain $\eta := 1.0$ for both algorithms. We also set $\varepsilon = 10^{-1}$ for N-SPIDER, see Zhang & Xiao (2019a, Algorithm 3). In addition, we conduct similar grid search as before to choose the suitable parameters for these algorithms. The chosen parameters are presented in Supp. Doc. F. The results on these datasets are depicted in Figure 3 and Figure 4.

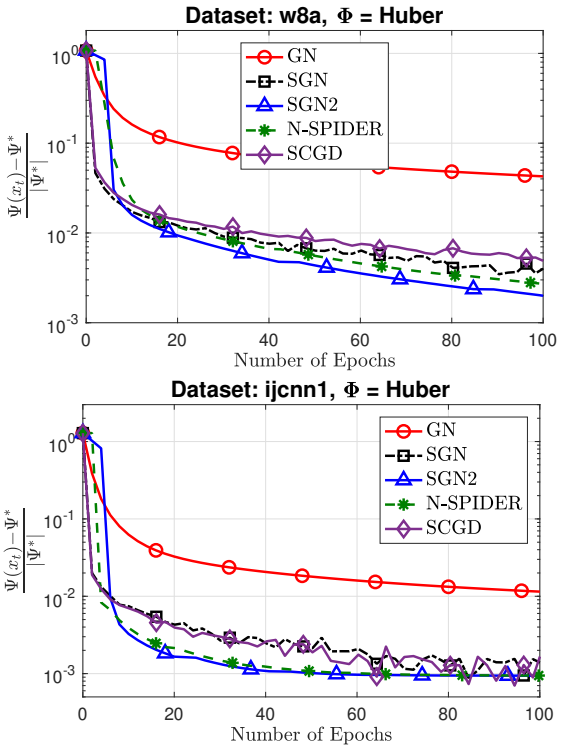

Figure 3. The performance of 5 algorithms on `w8a` and `ijcnn1` datasets.

From both figures, SGN2 seems to perform best in all datasets. N-SPIDER is better than SGN and comparable with SGN in `ijcnn1` and `covtye`. SGN is comparable with SCGD in `ijcnn1` dataset while having better performance in the remaining ones. GN still perform poorly in these cases since it use full samples to compute $\tilde{F}$ and $\tilde{J}$.

5.2. Optimization Involving Expectation Constraints

We consider the following optimization problem:

$$\min_{x\in\mathbb{R}^p} \Big\{ g(x) \;\text{ s.t. }\; \mathbb{E}_\xi\left[\mathbf{F}(x,\xi)\right] \leq 0 \Big\}, \tag{34}$$

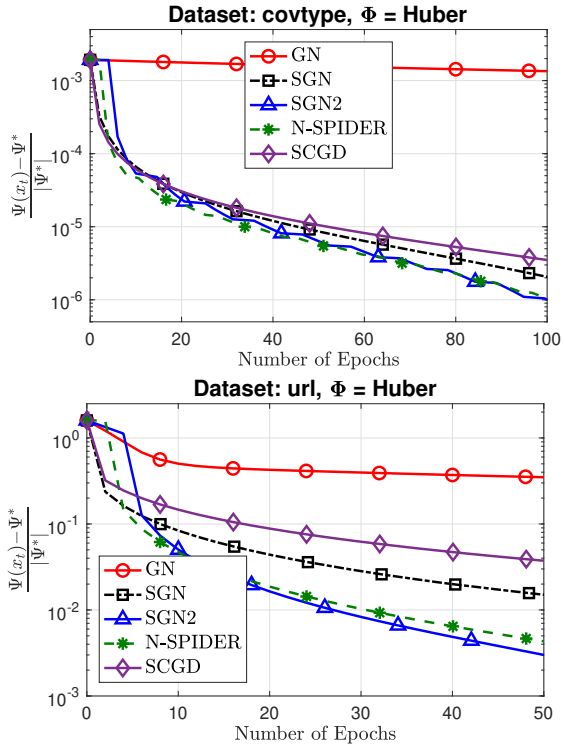

Figure 4. The performance of 5 algorithms on `covtype` and `url_combined` datasets.

where $g : \mathbb{R}^p \to \mathbb{R} \cup \{+\infty\}$ is a convex function, possibly nonsmooth, and $\mathbf{F} : \mathbb{R}^p \times \Omega \to \mathbb{R}^q$ is a smooth stochastic function. This problem has various applications such as optimization with conditional value at risk (CVaR) constraints and metric learning (Lan & Zhou, 2016) among others. Let us consider an exact penalty formulation of (34) as

$$\min_{x\in\mathbb{R}^p}\Big\{\Psi(x) := g(x) + \phi(\mathbb{E}_\xi\left[\mathbf{F}(x,\xi)\right])\Big\}, \qquad (35)$$

where $\phi(u) := \rho\sum_{i=1}^q [u_i]_+$ with $[u]_+ := \max\{0, u\}$ and $\rho > 0$ is a given penalty parameter. Clearly, (35) coincides with (3), an extension of (1).

We evaluate 3 algorithms on the asset allocation problem (Rockafellar & Uryasev, 2000) as an instance of (34):

$$\begin{cases} \min\limits_{\tau\in[\underline{\tau},\bar{\tau}],z\in\mathbb{R}^p} -c^\top z + \phi\left(\tau + \frac{1}{\beta n}\sum_{i=1}^n [-\xi_i^\top z - \tau]_+\right) \\ \text{s.t} \qquad\qquad z \in \Delta_p := \left\{\hat{z}\in\mathbb{R}^p_+ \mid \sum_{i=1}^p \hat{z}_i = 1\right\}. \end{cases} \qquad (36)$$

To apply our methods, we need to smooth $[u]_+$ by $\frac{1}{2}\left[u + (u^2+\gamma^2)^{1/2} - \gamma\right]$ for a sufficiently small value $\gamma > 0$. If we introduce $x := (z,\tau)$, $F(x,\xi_i) := \tau + \frac{1}{2\beta}\left([(\xi_i^\top z + \tau)^2 + \gamma^2]^{1/2} - \xi_i^\top z - \tau - \gamma\right)$ for $i = 1,\cdots,n$, and $g(x) = -c^\top z + \delta_{\Delta_p}(x)$, then we can reformulate the smoothed approximation of (36) into (3), where $\delta_{\Delta_p\times[\underline{\tau},\bar{\tau}]}$ is the indicator of $\Delta_p \times [\underline{\tau},\bar{\tau}]$. Note that $F'(\cdot,\xi_i)$ is Lipschitz continuous with the Lipschitz constant $L_i := \frac{\|\xi_i\|^2}{2\beta\gamma}$. In our experiments, we choose $[\underline{\tau},\bar{\tau}]$ to be $[0,1]$, $\beta := 0.1$, and $\gamma := 10^{-3}$. We were experimenting different ρ and M, and eventually set $\rho := 5$ and $M := 5$.

We test three algorithms: GN, SGN, and SGN2 on both synthetic and real datasets. We follow the procedures from Lan et al. (2012) to generate synthetic data with $n = 10^5$ and $p \in \{300, 500, 700\}$. We also obtain real datasets of US stock prices for 889, 865, and 500 types of stocks described, e.g., in Sun & Tran-Dinh (2019) then bootstrap them to obtain different datasets of sizes $n = 10^5$. The details and additional results are given in Supp. Doc. F.

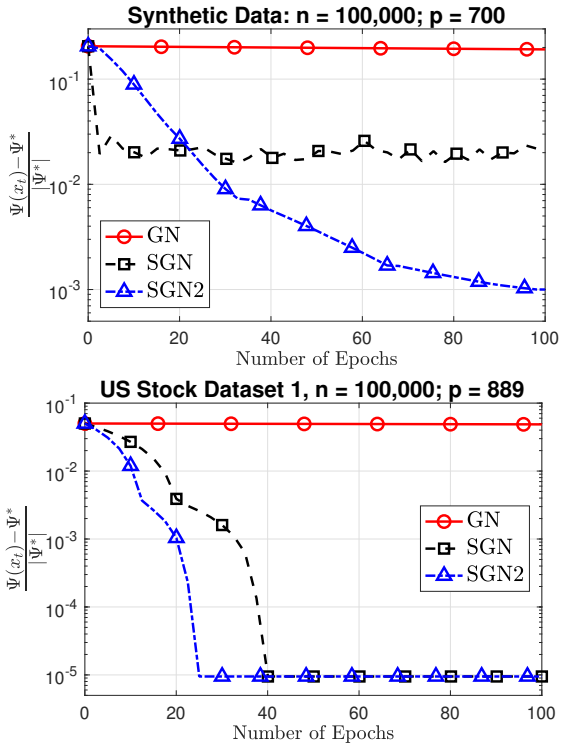

Figure 5. The performance of 3 algorithms on two datasets.

The performance of three algorithms on these datasets is depicted in Figure 5. SGN is still much better than GN in both experiments while SGN2 is the best among three. With the large amount of samples per iteration, GN performs poorly in these experiments.

Numerical results have confirmed the advantages of SGN and SGN2 which well align with our theoretical analysis.

Acknowledgements

The work of Q. Tran-Dinh has partially been supported by the National Science Foundation (NSF), award No. DMS-1619884 and the Office of Naval Research (ONR), grant No. N00014-20-1-2088 (2020–2023). The authors are thankful to Deyi Liu for providing some parts of Python codes used in the experiment section.

References

Bauschke, H. H. and Combettes, P. *Convex analysis and monotone operators theory in Hilbert spaces*. Springer-Verlag, 2nd edition, 2017.

Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2(1):183–202, 2009.

Blanchet, J., Goldfarb, D., Iyengar, G., Li, F., and Zhou, C. Unbiased simulation for optimizing stochastic function compositions. *arXiv preprint arXiv:1711.07564*, 2017.

Chambolle, A. and Pock, T. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.*, 40(1):120–145, 2011.

Davis, D. and Drusvyatskiy, D. Stochastic model-based minimization of weakly convex functions. *SIAM Journal Optim.*, 29(1):207–239, 2019.

Davis, D. and Grimmer, B. Proximally guided stochastic subgradient method for nonsmooth, nonconvex problems. *SIAM J. Optim.*, 29(3):1908–1930, 2019.

Drusvyatskiy, D. and Paquette, C. Efficiency of minimizing compositions of convex functions and smooth maps. *Math. Program.*, 178(1-2):503–558, 2019.

Duchi, J. and Ruan, F. Stochastic methods for composite and weakly convex optimization problems. *SIAM J. Optim.*, 28(4):3229–3259, 2018.

Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.

Esser, J. E. *Primal-dual algorithm for convex models and applications to image restoration, registration and nonlocal inpainting*. PhD Thesis, University of California, Los Angeles, Los Angeles, USA, 2010.

Ghadimi, S. and Lan, G. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Math. Program.*, 156(1-2):59–99, 2016.

Goldstein, T., Esser, E., and Baraniuk, R. Adaptive primal-dual hybrid gradient methods for saddle point problems. *Tech. Report.*, pp. 1–26, 2013. http://arxiv.org/pdf/1305.0546v1.pdf.

Lan, G. and Zhou, Z. Algorithms for stochastic optimization with functional or expectation constraints. *Comput. Optim. Appl.*, 76: 461–498, 2020.

Lan, G., Nemirovski, A., and Shapiro, A. Validation analysis of mirror descent stochastic approximation method. *Math. Program.*, 134(2):425–458, 2012.

Lewis, A. S. and Wright, S. J. A proximal method for composite minimization. *Math. Program.*, 158(1-2):501–546, 2016.

Lian, X., Wang, M., and Liu, J. Finite-sum composition optimization via variance reduced gradient descent. In *Artificial Intelligence and Statistics*, pp. 1159–1167, 2017.

Liu, L., Liu, J., and Tao, D. Variance reduced methods for non-convex composition optimization. *arXiv preprint arXiv:1711.04416*, 2017.

Lohr, S. L. *Sampling: Design and Analysis*. Nelson Education, 2009.

Nesterov, Y. *Introductory lectures on convex optimization: A basic course*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, 2004.

Nesterov, Y. Modified Gauss-Newton scheme with worst case guarantees for global performance. *Optim. Method Softw.*, 22(3):469–483, 2007.

Nguyen, L. M., Liu, J., Scheinberg, K., and Takáč, M. SARAH: A novel method for machine learning problems using stochastic recursive gradient. *ICML*, 2017.

Nguyen, L. M., van Dijk, M., Phan, D. T., Nguyen, P. H., Weng, T.-W., and Kalagnanam, J. R. Optimal finite-sum smooth non-convex optimization with SARAH. *arXiv preprint arXiv:1901.07648*, 2019.

Pham, H. N., Nguyen, M. L., Phan, T. D., and Tran-Dinh, Q. ProxSARAH: An efficient algorithmic framework for stochastic composite nonconvex optimization. *J. Mach. Learn. Res.*, 21:1–48, 2020.

Rockafellar, R. T. and Wets, R. J. B. Stochastic variational inequalities: single-stage to multistage. *Math. Program.*, 165(1):331–360, 2017.

Rockafellar, T. R. and Uryasev, S. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42, 2000.

Sun, T. and Tran-Dinh, Q. Generalized Self-Concordant Functions: A Recipe for Newton-Type Methods. *Math. Program.*, 178:145–213, 2019.

Tran-Dinh, Q. Proximal Alternating Penalty Algorithms for Constrained Convex Optimization. *Comput. Optim. Appl.*, 72(1):1–43, 2019.

Tran-Dinh, Q. and Diehl, M. Proximal methods for minimizing the sum of a convex function and a composite function. *Tech. Report*, KU Leuven, OPTEC and ESAT/SCD, Belgium, May 2011.

Tran-Dinh, Q., Fercoq, O., and Cevher, V. A smooth primal-dual optimization framework for nonsmooth composite convex minimization. *SIAM J. Optim.*, 28(1):96–134, 2018.

Tropp, J. A. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12 (4):389–434, 2012.

Wang, M., Fang, E., and Liu, L. Stochastic compositional gradient descent: algorithms for minimizing compositions of expected-value functions. *Math. Program.*, 161 (1-2):419–449, 2017a.

Wang, M., Liu, J., and Fang, E. X. Accelerating stochastic composition optimization. *The Journal of Machine Learning Research*, 18(1):3721–3743, 2017b.

Xu, Y. and Xu, Y. Katyusha acceleration for convex finite-sum compositional optimization. *arXiv preprint arXiv:1910.11217*, 2019.

Yang, S., Wang, M., and Fang, E. X. Multilevel stochastic gradient methods for nested composition optimization. *SIAM J. Optim.*, 29(1):616–659, 2019.

Yu, Y. and Huang, L. Fast stochastic variance reduced admm for stochastic composition optimization. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pp. 3364–3370. AAAI Press, 2017.

Zhang, J. and Xiao, L. Multi-level composite stochastic optimization via nested variance reduction. *arXiv preprint arXiv:1908.11468*, 2019a.

Zhang, J. and Xiao, L. A stochastic composite gradient method with incremental variance reduction. *Advances in Neural Information Processing Systems*, 28:9078–9088, 2019b.

Zhang, J. and Xiao, L. Stochastic variance-reduced prox-linear algorithms for nonconvex composite optimization. *arXiv preprint arXiv:2004.04357*, 2020.

Zhao, L., Mammadov, M., and Yearwood, J. From convex to nonconvex: a loss function analysis for binary classification. In *IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 1281–1288. IEEE, 2010.

SUPPLEMENTARY DOCUMENT

Stochastic Gauss-Newton Algorithms for Nonconvex Compositional Optimization

A. The Proof of Technical Results in Section 2: Mathematical Tools

This section provides the full proof of technical results in Section 2. Let us first recall the bound (10). The proof of this bound can be found, e.g., in Nesterov (2007). However, for completeness, we prove it here.

The proof of (10). Since F' is L_F-Lipschitz continuous with a Lipschitz constant L_F, we have $\|F(y)-F(x)-F'(x)(y-x)\| \le \frac{L_F}{2}\|y-x\|^2$ for any $x, y \in \mathbb{R}^p$. On the other hand, since ϕ is M_ϕ-Lipschitz continuous, we have $\phi(u) \le \phi(v) + M_\phi\|u-v\|$ for any $u, v \in \mathbb{R}^q$. Hence, we have

$$\begin{aligned}\phi(F(y)) &\le \phi(F(x)+F'(x)(y-x)) + M_\phi\|F(y)-F(x)-F'(x)(y-x)\| \\ &\le \phi(F(x)+F'(x)(y-x)) + \tfrac{M_\phi L_F}{2}\|y-x\|^2,\end{aligned}$$

which proves (10). □

A.1. The Proof of Lemma 2.1: Approximate Optimality Condition

Lemma. 2.1. *Suppose that Assumption 1.1 holds. Let $\widetilde{T}_M(x)$ be computed by (11) and $\widetilde{G}_M(x)$ be defined by (13). Then, $\mathcal{E}(\widetilde{T}_M(x), y)$ of (1) or (2) defined by (7) with $y \in \partial\phi(F(\widetilde{T}_M(x)))$ is bounded by*

$$\begin{aligned}\mathcal{E}(\widetilde{T}_M(x), y) &:= \operatorname{dist}\Big(0, -F(\widetilde{T}_M(x)) + \partial\phi^*(y)\Big) + \|F'(\widetilde{T}_M(x))^\top y\| \\ &\le \Big(1+\tfrac{M_\phi L_F}{M}\Big)\|\widetilde{G}_M(x)\| + \tfrac{(1+L_F)}{2M^2}\|\widetilde{G}_M(x)\|^2 + \|\widetilde{F}(x)-F(x)\| + \tfrac{1}{2}\|\widetilde{J}(x)-F'(x)\|^2.\end{aligned} \tag{14}$$

Proof. First, the optimality condition of (11) becomes

$$0 \in \widetilde{J}(x)^\top\partial\phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x)-x)) + M(\widetilde{T}_M(x)-x). \tag{37}$$

We can rewrite this optimality condition as

$$r_F(x) = F'(\widetilde{T}_M(x))^\top y \qquad \text{and} \qquad r_D(x) \in -F(\widetilde{T}_M(x)) + \partial\phi^*(y),$$

where

$$\begin{cases} r_F(x) &:= M(x-\widetilde{T}_M(x)) + (F'(\widetilde{T}_M(x)) - \widetilde{J}(x))^\top y, \\ r_D(x) &:= \widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x)-x) - F(\widetilde{T}_M(x)).\end{cases}$$

Next, since $y \in \partial\phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x)-x))$ and ϕ is M_ϕ-Lipschitz continuous, we can bound y as $\|y\| \le M_\phi$. Now, we need to bound r_F as follows:

$$\begin{aligned}\|r_F(x)\| &= \|M(x-\widetilde{T}_M(x)) + (F'(\widetilde{T}_M(x)) - \widetilde{J}(x))^\top y\| \\ &= \|M(x-\widetilde{T}_M(x)) + (F'(\widetilde{T}_M(x)) - F'(x))^\top y + (F'(x) - \widetilde{J}(x))^\top y\| \\ &\le M\|x-\widetilde{T}_M(x)\| + \|F'(\widetilde{T}_M(x)) - F'(x)\|_F\|y\| + \|F'(x)-\widetilde{J}(x)\|_F\|y\| \\ &\le \|\widetilde{G}_M(x)\| + M_\phi\|F'(\widetilde{T}_M(x)) - F'(x)\|_F + M_\phi\|F'(x)-\widetilde{J}(x)\| \\ &\le \Big(1+\tfrac{M_\phi L_F}{M}\Big)\|\widetilde{G}_M(x)\| + M_\phi\|F'(x)-\widetilde{J}(x)\|.\end{aligned}$$

Similarly, we can also bound r_D as

$$\begin{aligned}\|r_D(x)\| &= \|\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x) - F(\widetilde{T}_M(x))\| \\ &= \|\widetilde{F}(x) - F(x) + F(x) + F'(x)(\widetilde{T}_M(x) - x) - F(\widetilde{T}_M(x)) + [\widetilde{J}(x) - F'(x)](\widetilde{T}_M(x) - x)\| \\ &\le \|\widetilde{F}(x) - F(x)\| + \|F(x) + F'(x)(\widetilde{T}_M(x) - x) - F(\widetilde{T}_M(x))\| + \|[\widetilde{J}(x) - F'(x)](\widetilde{T}_M(x) - x)\| \\ &\le \|\widetilde{F}(x) - F(x)\| + \tfrac{L_F}{2}\|\widetilde{T}_M(x) - x\|^2 + \tfrac{1}{2}\|F'(x) - \widetilde{J}(x)\|^2 + \tfrac{1}{2}\|\widetilde{T}_M(x) - x\|^2 \\ &= \|\widetilde{F}(x) - F(x)\| + \tfrac{1}{2}\|F'(x) - \widetilde{J}(x)\|^2 + \tfrac{(1+L_F)}{2M^2}\|\widetilde{G}_M(x)\|^2.\end{aligned}$$

Combining these bounds, we can show that

$$\begin{aligned}\mathcal{E}(\widetilde{T}_M(x), y) &:= \|F'(\widetilde{T}_M(x))^\top y\| + \mathrm{dist}\left(0, -F(\widetilde{T}_M(x)) + \partial\phi^*(y)\right) \\ &\le \left(1 + \tfrac{M_\phi L_F}{M}\right)\|\widetilde{G}_M(x)\| + \tfrac{(1+L_F)}{2M^2}\|\widetilde{G}_M(x)\|^2 + \|\widetilde{F}(x) - F(x)\| + \tfrac{1}{2}\|F'(x) - \widetilde{J}(x)\|^2,\end{aligned}$$

which is exactly (14). □

B. The Proof of Technical Results in Section 3: Convergence of Inexact GN Framework

This appendix provides the full proof of technical results in Section 3 on convergence of the inexact Gauss-Newton framework, Algorithm 1.

B.1. The Proof of Lemma 3.1: Descent Property

Lemma. 3.1. *Let Assumption 1.1 hold, $\widetilde{T}_M(x)$ be computed by (11), and $\widetilde{G}_M(x) := M(x - \widetilde{T}_M(x))$ be the prox-gradient mapping of F. Then, for any $z \in \mathbb{R}^p$, we have*

$$\phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)) \le \phi(\widetilde{F}(x) + \widetilde{J}(x)(z - x)) - \langle \widetilde{G}_M(x), z - x\rangle - \tfrac{1}{M}\|\widetilde{G}_M(x)\|^2. \tag{38}$$

For any $\beta_d > 0$, we also have

$$\begin{aligned}\phi(F(\widetilde{T}_M(x))) &\le \phi(F(x)) + 2L_\phi\|F(x) - \widetilde{F}(x)\| + M_\phi\|F'(x) - \widetilde{J}(x)\|\|x - \widetilde{T}_M(x)\| - \tfrac{(2M - M_\phi L_F)}{2}\|\widetilde{T}_M(x) - x\|^2 \\ &\le \phi(F(x)) + 2L_\phi\|F(x) - \widetilde{F}(x)\| + \tfrac{M_\phi}{2\beta_d}\|F'(x) - \widetilde{J}(x)\|_F^2 - \tfrac{(2M - M_\phi L_F - \beta_d L_\phi)}{2M^2}\|\widetilde{G}_M(x)\|^2.\end{aligned} \tag{15}$$

Proof. The optimality condition (37) can be written as

$$\widetilde{J}(x)^\top y = M(x - \widetilde{T}_M(x)) \quad \text{and} \quad y \in \partial\phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)).$$

By convexity of ϕ, using the above relations, we have

$$\begin{aligned}\phi(\widetilde{F}(x) + \widetilde{J}(x)(z - x)) &\ge \phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)) + \langle y, \widetilde{F}(x) + \widetilde{J}(x)(z - x) - (\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x))\rangle \\ &\ge \phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)) + \langle \widetilde{J}(x)^\top y, z - \widetilde{T}_M(x)\rangle \\ &= \phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)) + M\langle z - \widetilde{T}_M(x), x - \widetilde{T}_M(x)\rangle \\ &= \phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)) + M\langle x - \widetilde{T}_M(x), z - x\rangle + M\|x - \widetilde{T}_M(x)\|^2 \\ &= \phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)) + \langle \widetilde{G}_M(x), z - x\rangle + \tfrac{1}{M}\|\widetilde{G}_M(x)\|^2,\end{aligned}$$

which implies (38).

Now, combining (10) and (38), we can show that

$$\begin{aligned}\phi(F(\widetilde{T}_M(x))) &\overset{(10)}{\leq} \phi(F(x) + F'(x)(\widetilde{T}_M(x) - x)) + \tfrac{M_\phi L_F}{2}\|\widetilde{T}_M(x) - x\|^2\\ &\leq \phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x)) + \tfrac{M_\phi L_F}{2}\|\widetilde{T}_M(x) - x\|^2\\ &\quad + |\phi(F(x) + F'(x)(\widetilde{T}_M(x) - x)) - \phi(\widetilde{F}(x) + \widetilde{J}(x)(\widetilde{T}_M(x) - x))|\\ &\overset{(38)}{\leq} \phi(\widetilde{F}(x) + \widetilde{J}(x)(z - x)) - M\langle x - \widetilde{T}_M(x), z - x\rangle - \tfrac{(2M - M_\phi L_F)}{2}\|\widetilde{T}_M(x) - x\|^2\\ &\quad + M_\phi\|F(x) - \widetilde{F}(x) + [F'(x) - \widetilde{J}(x)](\widetilde{T}_M(x) - x)\|\\ &\leq \phi(F(x)) - \tfrac{(2M - M_\phi L_F)}{2}\|\widetilde{T}_M(x) - x\|^2 + M_\phi\|F(x) - \widetilde{F}(x)\|\\ &\quad + M_\phi\|F(x) - \widetilde{F}(x) - \widetilde{J}(x)(z - x)\| - M\langle x - \widetilde{T}_M(x), z - x\rangle\\ &\quad + M_\phi\|(F'(x) - \widetilde{J}(x))(\widetilde{T}_M(x) - x)\|.\end{aligned}$$

Substituting $z = x$ into this estimate, we obtain

$$\begin{aligned}\phi(F(\widetilde{T}_M(x))) \leq \phi(F(x)) - \tfrac{(2M - M_\phi L_F)}{2}\|\widetilde{T}_M(x) - x\|^2 + 2M_\phi\|F(x) - \widetilde{F}(x)\|\\ + M_\phi\|(F'(x) - \widetilde{J}(x))(\widetilde{T}_M(x) - x)\|.\end{aligned} \tag{39}$$

Using the Cauchy-Schwarz inequality, we have

$$\|(F'(x) - \widetilde{J}(x))(\widetilde{T}_M(x) - x)\| \leq \|F'(x) - \widetilde{J}(x)\|\|\widetilde{T}_M(x) - x\|.$$

Next, applying Young's inequality to the right hand side of this inequality, for any $\beta_d > 0$, we obtain

$$\|(F'(x) - \widetilde{J}(x))(\widetilde{T}_M(x) - x)\| \leq \|F'(x) - \widetilde{J}(x)\|_F\|\widetilde{T}_M(x) - x\| \leq \frac{1}{2\beta_d}\|F'(x) - \widetilde{J}(x)\|^2 + \frac{\beta_d}{2}\|\widetilde{T}_M(x) - x\|^2. \tag{40}$$

Finally, plugging (40) into (39), we have

$$\begin{aligned}\phi(F(\widetilde{T}_M(x))) &\leq \phi(F(x)) - \tfrac{(2M - M_\phi L_F)}{2}\|\widetilde{T}_M(x) - x\|^2 + 2L_\phi\|F(x) - \widetilde{F}(x)\| + M_\phi\|F'(x) - \widetilde{J}(x)\|\|\widetilde{T}_M(x) - x\|\\ &\leq \phi(F(x)) - \tfrac{(2M - M_\phi L_F - \beta_d L_\phi)}{2}\|\widetilde{T}_M(x) - x\|^2 + 2L_\phi\|F(x) - \widetilde{F}(x)\| + \tfrac{M_\phi}{2\beta_d}\|F'(x) - \widetilde{J}(x)\|^2,\end{aligned}$$

for any $\beta_d > 0$, which exactly implies (15). □

B.2. The Proof of Theorem 3.1: Convergence Rate of Algorithm 1

Theorem. 3.1. *Assume that Assumptions 1.1 and 1.2 are satisfied. Let $\{x_t\}$ be generated by Algorithm 1 to solve either (1) or (2). Then, the following statements hold:*

(a) *If (16) holds for some $\varepsilon \geq 0$, then*

$$\min_{0\leq t\leq T}\|\widetilde{G}_M(x_t)\|^2 \leq \frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \leq \frac{2M^2\left[\Psi(x_0) - \Psi^\star\right]}{C_g(T+1)} + \frac{\varepsilon^2}{2}, \tag{19}$$

where $C_g := 2M - M_\phi(L_F + \beta_d)$ for $M > \frac{1}{2}M_\phi(L_F + \beta_d)$.

(b) *If (17) and (18) hold for given $C_a > 0$, then*

$$\min_{0\leq t\leq T}\|\widetilde{G}_M(x_t)\|^2 \leq \frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \leq \frac{2M^2\left[\Psi(x_0) - \Psi^\star\right]}{C_a(T+1)} + \frac{\varepsilon^2}{2}. \tag{20}$$

Consequently, with $\varepsilon > 0$, the total number of iterations T to achieve $\frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \leq \varepsilon^2$ is at most

$$T := \left\lfloor \frac{4M^2\left[\Psi(x_0) - \Psi^\star\right]}{D\varepsilon^2} \right\rfloor = \mathcal{O}\left(\frac{\left[\Psi(x_0) - \Psi^\star\right]}{\varepsilon^2}\right),$$

where $D := C_g$ for the case (a) and $D := C_a$ for the case (b).

Proof. Using the second inequality of (15) with $x := x_t$ and $T_M(x) = x_{t+1}$, we have

$$\phi(F(x_{t+1})) \le \phi(F(x_t)) - \frac{(2M - M_\phi(L_F + \beta_d))}{2}\|x_{t+1} - x_t\|^2 + 2M_\phi\|F(x_t) - \widetilde{F}_t\| + \frac{M_\phi\|F'(x_t) - \widetilde{J}_t\|^2}{2\beta_d}. \quad (41)$$

(a) If (16) holds for some $\varepsilon \ge 0$, then using (16) into (41), we have

$$\phi(F(x_{t+1})) \le \phi(F(x_t)) - \frac{C_g}{2}\|x_{t+1} - x_t\|^2 + 2M_\phi \cdot \frac{C_g\varepsilon^2}{16M_\phi M^2} + \frac{M_\phi}{2\beta_d} \cdot \frac{\beta_d C_g\varepsilon^2}{4M_\phi M^2},$$

where $C_g := 2M - M_\phi(L_F + \beta_d) > 0$. Since $\Psi(x) = \phi(F(x))$, the last estimate leads to

$$\Psi(x_{t+1}) \le \Psi(x_t) - \frac{C_g}{2}\|x_{t+1} - x_t\|^2 + \frac{C_g\varepsilon^2}{4M^2}.$$

By induction, $\widetilde{G}_M(x_t) := M(x_t - \widetilde{T}_M(x_t))$, and $\Psi(x_{T+1}) \ge \Psi^\star$, we can show that

$$\frac{1}{M^2(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 = \frac{1}{T+1}\sum_{t=0}^{T}\|x_{t+1} - x_t\|^2 \le \frac{2\left[\Psi(x_0) - \Psi^\star\right]}{C_g(T+1)} + \frac{\varepsilon^2}{2M^2}, \quad (42)$$

which leads to (19).

(b) If (17) and (18) are used, then from (41) and (18), we have

$$\phi(F(x_{t+1})) \le \phi(F(x_t)) - \frac{C_1}{2}\|x_{t+1} - x_t\|^2 + \frac{C_2}{2}\|x_t - x_{t-1}\|^2, \quad \forall t \ge 1.$$

where $C_1 := 2M - M_\phi L_F - \beta_d M_\phi$ and $C_2 := 2M_\phi\sqrt{C_f} + \frac{M_\phi C_d}{2\beta_d}$. For $t = 0$, it follows from (41) and (17) that

$$\phi(F(x_1)) \le \phi(F(x_0)) - \frac{C_1}{2}\|x_1 - x_0\|^2 + \frac{(C_1 - C_2)\varepsilon^2}{4M^2}.$$

Now, note that $\Psi(x) = \phi(F(x))$, the last two estimates respectively become

$$\Psi(x_{t+1}) \le \Psi(x_t) - \frac{C_1}{2}\|x_{t+1} - x_t\|^2 + \frac{C_2}{2}\|x_t - x_{t-1}\|^2, \quad \forall t \ge 1,$$

and for $t = 0$, it holds that

$$\Psi(x_1) \le \Psi(x_0) - \frac{C_1}{2}\|x_1 - x_0\|^2 + \frac{(C_1 - C_2)\varepsilon^2}{4M^2}.$$

By induction and $\Psi^\star \le \Psi(x_{T+1})$, this estimate leads to

$$\begin{aligned}\Psi^\star \le \Psi(x_{T+1}) \quad &\le \Psi(x_0) - \tfrac{(C_1-C_2)}{2}\textstyle\sum_{t=0}^{T}\|x_{t+1} - x_t\|^2 + \tfrac{(C_1-C_2)\varepsilon^2}{4M^2} \\ &- \tfrac{C_2}{2}\|x_{T+1} - x_T\|^2.\end{aligned}$$

Since $C_1 > C_2$, if we define $C_a := C_1 - C_2 > 0$, then the last inequality implies

$$\frac{1}{M^2(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 = \frac{1}{(T+1)}\sum_{t=0}^{T}\|x_{t+1} - x_t\|^2 \le \frac{2\left[\Psi(x_0) - \Psi^\star\right]}{C_a(T+1)} + \frac{\varepsilon^2}{4M^2},$$

which leads to (20). The last statement of this theorem is a direct consequence of either (19) or (20), and we omit the detailed derivation here. □

C. High Probability Inequalities and Variance Bounds

Since our methods are stochastic, we recall some mathematical tools from high probability and concentration theory, as well as variance bounds that will be used for our analysis. First, we need the following lemmas to estimate sample complexity of our algorithms.

Lemma C.1 (Matrix Bernstein inequality (Tropp, 2012)(Theorem 1.6)). *Let $X_1, X_2, \cdots, X_n$ be independent random matrices in $\mathbb{R}^{p_1\times p_2}$. Assume that $\mathbb{E}[X_i] = 0$ and $\|X_i\| \le R$ a.s. for $i = 1, \cdots, n$ and given $R > 0$, where $\|\cdot\|$ is the spectral norm. Define $\sigma_X^2 := \max\left\{\left\|\sum_{i=1}^n \mathbb{E}\left[X_iX_i^\top\right]\right\|, \left\|\sum_{i=1}^n \mathbb{E}\left[X_i^\top X_i\right]\right\|\right\}$. Then, for any $\epsilon > 0$, we have*

$$\mathbf{Prob}\left(\left\|\sum_{i=1}^n X_i\right\| \ge \epsilon\right) \le (p_1+p_2)\exp\left(-\frac{3\epsilon^2}{6\sigma_X^2 + 2R\epsilon}\right).$$

As a consequence, if $\sigma_X^2 \le \bar{\sigma}_X^2$ for a given $\bar{\sigma}_X^2 > 0$, then

$$\mathbf{Prob}\left(\left\|\sum_{i=1}^n X_i\right\| \le \epsilon\right) \ge 1 - (p_1+p_2)\exp\left(-\frac{3\epsilon^2}{6\bar{\sigma}_X^2 + 2R\epsilon}\right).$$

Lemma C.2 (Lohr (2009)). *Let $\widetilde{F}(x_t)$ and $\widetilde{J}(x_t)$ be the mini-batch stochastic estimators of $F(x_t)$ and $F'(x_t)$ defined by (21), respectively, and $\mathcal{F}_t := \sigma(x_0, x_1, \cdots, x_{t-1})$ be the σ-field generated by $\{x_0, x_1, \cdots, x_{t-1}\}$. Then, these are unbiased estimators, i.e., $\mathbb{E}\left[\widetilde{F}(x_t) \mid \mathcal{F}_t\right] = F(x_t)$ and $\mathbb{E}\left[\widetilde{J}(x_t) \mid \mathcal{F}_t\right] = F'(x_t)$. Moreover, under Assumption 1.2, we have*

$$\mathbb{E}\left[\|\widetilde{F}(x_t) - F(x_t)\|^2 \mid \mathcal{F}_t\right] \le \frac{\sigma_F^2}{b_t} \quad \text{and} \quad \mathbb{E}\left[\|\widetilde{J}(x_t) - F'(x_t)\|^2 \mid \mathcal{F}_t\right] \le \frac{\sigma_D^2}{\hat{b}_t}. \tag{43}$$

Lemma C.3 (Nguyen et al. (2017); Pham et al. (2020)). *Let $\widetilde{F}_t$ and $\widetilde{J}_t$ be the mini-batch SARAH estimators of $F(x_t)$ and $F'(x_t)$, respectively defined by (27), and $\mathcal{F}_t := \sigma(x_0, x_1, \cdots, x_{t-1})$ be the σ-field generated by $\{x_0, x_1, \cdots, x_{t-1}\}$. Then, we have the following estimate*

$$\mathbb{E}\left[\|\widetilde{F}_t - F(x_t)\|^2 \mid \mathcal{F}_t\right] = \|\widetilde{F}_{t-1} - F(x_{t-1})\|^2 + \rho_t\mathbb{E}_\xi\left[\|\mathbf{F}(x_t,\xi) - \mathbf{F}(x_{t-1},\xi)\|^2\right] - \rho_t\|F(x_t) - F(x_{t-1})\|^2, \tag{44}$$

where $\rho_t := \frac{n-b_t}{(n-1)b_t}$ if $F(x) := \frac{1}{n}\sum_{i=1}^n F_i(x)$, and $\rho_t := \frac{1}{b_t}$, otherwise, i.e., $F(x) = \mathbb{E}_\xi[\mathbf{F}(x,\xi)]$.

Similarly, we also have

$$\mathbb{E}\left[\|\widetilde{J}_t - F'(x_t)\|^2 \mid \mathcal{F}_t\right] = \|\widetilde{J}_{t-1} - F'(x_{t-1})\|^2 + \hat{\rho}_t\mathbb{E}_\xi\left[\|\mathbf{F}'(x_t,\xi) - \mathbf{F}'(x_{t-1},\xi)\|^2\right] - \hat{\rho}_t\|F'(x_t) - F'(x_{t-1})\|^2, \tag{45}$$

where $\hat{\rho}_t := \frac{n-\hat{b}_t}{(n-1)\hat{b}_t}$ if $F(x) := \frac{1}{n}\sum_{i=1}^n F_i(x)$, and $\hat{\rho}_t := \frac{1}{\hat{b}_t}$, otherwise, i.e., $F(x) = \mathbb{E}_\xi[\mathbf{F}(x,\xi)]$.

D. The Proof of Technical Results in Section 4

This appendix provides the full proof of technical results in Section 4 on our stochastic Gauss-Newton methods.

D.1. The Proof of Theorem 4.1: Convergence of The Stochastic Gauss-Newton Method for Solving (1)

Theorem. 4.1. *Suppose that Assumptions 1.1 and 1.2 hold for (1). Let $\widetilde{F}_t$ and $\widetilde{J}_t$ defined by (21) be mini-batch stochastic estimators of $F(x_t)$ and $F'(x_t)$, respectively. Let $\{x_t\}$ be generated by Algorithm 1 (called **SGN**) to solve (1). For a given tolerance $\varepsilon > 0$, assume that b_t and $\hat{b}_t$ in (21) are chosen as*

$$\left\{\begin{array}{lcll} b_t &:=& \left\lfloor \frac{256M_\phi^2M^4\sigma_F^2}{C_g^2\varepsilon^4}\right\rfloor &= \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^4}\right),\\ \hat{b}_t &:=& \left\lfloor \frac{2M_\phi M^2\sigma_D^2}{\beta_d C_g\varepsilon^2}\right\rfloor &= \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^2}\right). \end{array}\right. \tag{22}$$

Furthermore, let $\widehat{x}_T$ be chosen uniformly at random in $\{x_t\}_{t=0}^T$ as the output of Algorithm 1 after T iterations. Then

$$\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] = \frac{1}{(T+1)}\sum_{t=0}^{T}\mathbb{E}\left[\|\widetilde{G}_M(x_t)\|^2\right] \leq \frac{2M^2\left[\Psi(x_0)-\Psi^\star\right]}{C_g(T+1)} + \frac{\varepsilon^2}{2}, \tag{23}$$

where $C_g := 2M - M_\phi(L_F+\beta_d)$ with $M > \frac{1}{2}M_\phi(L_F+\beta_d)$. Moreover, the total number $\mathcal{T}_f$ of function evaluations $\mathbf{F}(x_t,\xi)$ and the total number $\mathcal{T}_d$ of Jacobian evaluations $\mathbf{F}'(x_t,\zeta)$ to achieve $\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] \leq \varepsilon^2$ do not exceed

$$\left\{\begin{array}{lcll} \mathcal{T}_f &:=& \left\lfloor\frac{1024M^6M_\phi^2\sigma_F^2\left[\Psi(x_0)-\Psi^\star\right]}{C_g^3\varepsilon^6}\right\rfloor &= \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^6}\right),\\ \mathcal{T}_d &:=& \left\lfloor\frac{8M^4M_\phi\sigma_D^2\left[\Psi(x_0)-\Psi^\star\right]}{\beta_dC_g^2\varepsilon^4}\right\rfloor &= \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^4}\right).\end{array}\right. \tag{24}$$

Proof. Let $\mathcal{F}_t := \sigma(x_0,x_1,\cdots,x_{t-1})$ be the σ-field generated by $\{x_0,x_1,\cdots,x_{t-1}\}$. By repeating a similar proof as of (19), but taking the full expectation overall the randomness with $\mathbb{E}[\cdot] = \mathbb{E}[\mathbb{E}[\cdot]\mid\mathcal{F}_{t+1}]$, we have

$$\frac{1}{(T+1)}\sum_{t=0}^{T}\mathbb{E}\left[\|\widetilde{G}_M(x_t)\|^2\right] \leq \frac{2M^2\left[\Psi(x_0)-\Psi^\star\right]}{C_g(T+1)} + \frac{\varepsilon^2}{2}, \tag{46}$$

where $C_g := 2M - M_\phi(L_F+\beta_d)$ with $M > \frac{1}{2}M_\phi(L_F+\beta_d)$. Moreover, by the choice of $\widehat{x}_T$, we have $\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] = \frac{1}{(T+1)}\sum_{t=0}^{T}\mathbb{E}\left[\|\widetilde{G}_M(x_t)\|^2\right]$. Combining this relation and (46), we proves (23).

Next, by Lemma C.2, to guarantee the condition (16) in expectation, i.e.:

$$\left\{\begin{array}{lcl}\mathbb{E}\left[\|\widetilde{F}(x_t)-F(x_t)\|^2\mid\mathcal{F}_t\right] &\leq& \frac{C_g^2\varepsilon^4}{256M_\phi^2M^4},\\ \mathbb{E}\left[\|\widetilde{J}(x_t)-F'(x_t)\|^2\mid\mathcal{F}_t\right] &\leq& \frac{\beta_dC_g\varepsilon^2}{2M^2M_\phi},\end{array}\right.$$

we have to choose $\frac{\sigma_F^2}{b_t} \leq \frac{C_g^2\varepsilon^4}{256M_\phi^2M^4}$ and $\frac{\sigma_D^2}{\hat{b}_t} \leq \frac{\beta_dC_g\varepsilon^2}{2M_\phi M^2}$, which respectively lead to

$$b_t \geq \frac{256M_\phi^2M^4\sigma_F^2}{C_g^2\varepsilon^4} \quad\text{and}\quad \hat{b}_t \geq \frac{2M_\phi M^2\sigma_D^2}{\beta_dC_g\varepsilon^2}.$$

By rounding to the nearest integer, we obtain (22). Using (19), we can see that since $\mathbb{E}\left[\|\widetilde{G}_M(\hat{x}_T)\|^2\right] = \frac{1}{(T+1)}\sum_{t=0}^{T}\mathbb{E}\left[\|\widetilde{G}_M(x_t)\|^2\right]$, to guarantee $\mathbb{E}\left[\|\widetilde{G}_M(\hat{x}_T)\|^2\right] \leq \varepsilon^2$, we impose $\frac{2M^2[\Psi(x_0)-\Psi^\star]}{C_g(T+1)} \leq \frac{\varepsilon^2}{2}$, which leads to $T := \left\lfloor\frac{4M^2[\Psi(x_0)-\Psi^\star]}{C_g\varepsilon^2}\right\rfloor$. Hence, the total number $\mathcal{T}_f$ of stochastic function evaluations $\mathbf{F}(x_t,\xi)$ can be bounded by

$$\mathcal{T}_f := Tb_t = \left\lfloor\frac{1024M^6M_\phi^2\sigma_F^2\left[\Psi(x_0)-\Psi^\star\right]}{C_g^3\varepsilon^6}\right\rfloor = \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^6}\right).$$

Similarly, the total number $\mathcal{T}_d$ of stochastic Jacobian evaluations $\mathbf{F}'(x_t,\zeta)$ can be bounded by

$$\mathcal{T}_d := T\hat{b}_t = \left\lfloor\frac{8M^4M_\phi\sigma_D^2\left[\Psi(x_0)-\Psi^\star\right]}{\beta_dC_g^2\varepsilon^4}\right\rfloor = \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^4}\right).$$

These two last estimates prove (24). □

D.2. The Proof of Theorem 4.2: Convergence of The Stochastic Gauss-Newton Method for Solving (2)

Theorem. 4.2. *Suppose that Assumptions 1.1 and 1.2 hold for (2). Let $\widetilde{F}_t$ and $\widetilde{J}_t$ defined by (21) be mini-batch stochastic estimators to approximate $F(x_t)$ and $F'(x_t)$, respectively. Let $\{x_t\}$ be generated by Algorithm 1 for solving (2). Assume that b_t and $\hat{b}_t$ in (21) are chosen such that $b_t := \min\{n, \bar{b}_t\}$ and $\hat{b}_t := \min\left\{n, \hat{\bar{b}}_t\right\}$ for $t \geq 0$, where*

$$\left\{\begin{array}{lcl} \bar{b}_0 &:=& \left\lfloor \frac{32M_\phi M^2\sigma_F\left[48\sigma_F M_\phi M^2 + C_a\varepsilon^2\right]}{3C_a^2\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right)\right\rfloor,\\ \hat{\bar{b}}_0 &:=& \left\lfloor \frac{4M\sqrt{2M_\phi}\sigma_D\left(3M\sqrt{2M_\phi}\sigma_D + \sqrt{\beta_d C_a}\varepsilon\right)}{\beta_d C_a\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right\rfloor,\\ \bar{b}_t &:=& \left\lfloor \frac{\left(6\sigma_F^2 + 2\sigma_F\sqrt{C_f}\|x_t - x_{t-1}\|^2\right)}{3C_f^2\|x_t - x_{t-1}\|^4}\cdot\log\left(\frac{p+1}{\delta}\right)\right\rfloor \quad (t\geq 1),\\ \hat{\bar{b}}_t &:=& \left\lfloor \frac{\left(6\sigma_D^2 + 2\sigma_D\sqrt{C_d}\|x_t - x_{t-1}\|\right)}{3C_d\|x_t - x_{t-1}\|^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right\rfloor \quad (t\geq 1), \end{array}\right. \tag{25}$$

*for $\delta \in (0,1)$, and C_f and C_d given in **Condition 2**, where $\varepsilon > 0$ is a given tolerance.*
Then, we have the following conclusions:

- *With probability at least $1-\delta$, the bound (20) in Theorem 3.1 still holds.*
- *Moreover, the total number $\mathcal{T}_f$ of stochastic function evaluations $\mathbf{F}(x_t,\xi)$ and the total number $\mathcal{T}_d$ of stochastic Jacobian evaluations $\mathbf{F}'(x_t,\zeta)$ to guarantee $\frac{1}{(T+1)}\sum_{t=0}^{T}\|\widetilde{G}_M(x_t)\|^2 \leq \varepsilon^2$ do not exceed*

$$\left\{\begin{array}{lcl} \mathcal{T}_f &:=& \mathcal{O}\left(\frac{\sigma_F^2\left[\Psi(x_0) - \Psi^\star\right]}{\varepsilon^6}\cdot\log\left(\frac{p+1}{\delta}\right)\right),\\ \mathcal{T}_d &:=& \mathcal{O}\left(\frac{\sigma_D^2\left[\Psi(x_0) - \Psi^\star\right]}{\varepsilon^4}\cdot\log\left(\frac{p+q}{\delta}\right)\right). \end{array}\right. \tag{26}$$

Proof. We first use Lemma C.1 to estimate the total number of samples for $F(x_t)$ and $F'(x_t)$. Let $\mathcal{F}_t := \sigma(x_0, x_1, \cdots, x_{t-1})$ be the σ-field generated by $\{x_0, x_1, \cdots, x_{t-1}\}$. We define $X_i := F_i(x_t) - F(x_t) \in \mathbb{R}^p$ for $i \in \mathcal{B}_t$. Conditioned on $\mathcal{F}_t$, due to the choice of $\mathcal{B}_t$, $\{X_i\}_{i\in\mathcal{B}_t}$ are independent vector-valued random variables and $\mathbb{E}[X_i] = 0$. Moreover, by Assumption 1.2, we have $\|F_i(x) - F(x)\| \leq \sigma_F$ for all $i \in [n]$. This implies that $\|X_i\| \leq \sigma_F$ a.s. and $\mathbb{E}\left[\|X_i\|^2\right] \leq \sigma_F^2$. Hence, the conditions of Lemma C.1 hold. In addition, we have

$$\sigma_X^2 := \max\left\{\left\|\sum_{i\in\mathcal{B}_t}\mathbb{E}\left[X_iX_i^\top\right]\right\|, \left\|\sum_{i\in\mathcal{B}_t}\mathbb{E}\left[X_i^\top X_i\right]\right\|\right\} \leq \sum_{i\in\mathcal{B}_t}\mathbb{E}\left[\|X_i\|^2\right] \leq b_t\sigma_F^2.$$

Since $\widetilde{F}_t := \frac{1}{b_t}\sum_{i\in\mathcal{B}_t}F_i(x_t)$, by Lemma C.1, we have

$$\begin{array}{lcl} \mathbf{Prob}\left(\|\widetilde{F}_t - F(x_t)\| \leq \epsilon\right) &=& \mathbf{Prob}\left(\left\|\sum_{i\in\mathcal{B}_t}X_i\right\| \leq b_t\epsilon\right)\\ &\geq& 1 - (p+1)\exp\left(-\frac{3b_t^2\epsilon^2}{6b_t\sigma_F^2 + 2\sigma_F b_t\epsilon}\right)\\ &=& 1 - (p+1)\exp\left(-\frac{3b_t\epsilon^2}{6\sigma_F^2 + 2\sigma_F\epsilon}\right). \end{array}$$

Let us choose $\delta \in (0,1]$ such that $\delta \geq (p+1)\exp\left(-\frac{3b_t\epsilon^2}{6\sigma_F^2+2\sigma_F\epsilon}\right)$ and $\delta \leq 1$, then $\mathbf{Prob}\left(\|\widetilde{F}_t - F(x_t)\| \leq \epsilon\right) \geq 1-\delta$. Hence, we have $b_t \geq \left(\frac{6\sigma_F^2+2\sigma_F\epsilon}{3\epsilon^2}\right)\cdot\log\left(\frac{p+1}{\delta}\right)$.

To guarantee the first condition of (17), we choose $\epsilon := \frac{C_a\varepsilon^2}{16M_\phi M^2}$. Then, the condition on b_0 leads to $b_0 \geq \frac{32M_\phi M^2\sigma_F\left(48\sigma_F M_\phi M^2 + C_a\varepsilon^2\right)}{3C_g^2\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right)$. To guarantee the first condition of (18), we choose $\epsilon := \sqrt{C_f}\,\|x_t - x_{t-1}\|^2$.

Then, the condition on b_t leads to $b_t \geq \frac{(6\sigma_F^2+2\sigma_F\sqrt{C_f}\|x_t-x_{t-1}\|^2}{3C^2\|x_t-x_{t-1}\|^4}\cdot\log\left(\frac{p+1}{\delta}\right)$. Rounding both b_0 and b_t, we obtain

$$\left\{\begin{array}{lcll}\bar{b}_0 &:=& \left\lfloor\frac{32M_\phi M^2\sigma_F\left[48\sigma_F M_\phi M^2+C_a\varepsilon^2\right]}{3C_a^2\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right)\right\rfloor &= \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^4}\cdot\log\left(\frac{p}{\delta}\right)\right),\\ \bar{b}_t &:=& \left\lfloor\frac{(6\sigma_F^2+2\sigma_F\sqrt{C_f}\|x_t-x_{t-1}\|^2}{3C^2\|x_t-x_{t-1}\|^4}\cdot\log\left(\frac{p+1}{\delta}\right)\right\rfloor &= \mathcal{O}\left(\frac{\sigma_F^2}{\|x_t-x_{t-1}\|^4}\cdot\log\left(\frac{p}{\delta}\right)\right),\quad \forall t\geq 1.\end{array}\right.$$

Since $b_t \leq n$ for all $t\geq 0$, we have $b_t := \min\left\{n, \bar{b}_t\right\}$ for $t\geq 0$, which proves the first part of (25).

Next, we estimate a sample size for $\widetilde{J}_t$. Let us define $Y_i := F_i'(x_t) - F'(x_t)$. Then, similar to the above proof of X_i for F, we have $\widetilde{J}_t - F'(x_t) = \frac{1}{\hat{b}_t}\sum_{i\in\hat{\mathcal{B}}_t}(F_i'(x_t)-F'(x_t)) = \frac{1}{\hat{b}_t}\sum_{i\in\hat{\mathcal{B}}_t}Y_i$. Under Assumption 1.2, the sequence $\{Y_i\}$ satisfies all conditions of Lemma C.1. Hence, we obtain

$$\mathbf{Prob}\left(\|\widetilde{J}_t - F'(x_t)\|\leq\epsilon\right)\geq 1-(p+q)\exp\left(\frac{-3\hat{b}_t\epsilon^2}{6\sigma_D^2+2\sigma_D\epsilon}\right).$$

Hence, we can choose $\hat{b}_t \geq \left\lceil\frac{6\sigma_D^2+2\sigma_D\epsilon}{3\epsilon^2}\right\rceil\cdot\log\left(\frac{p+q}{\delta}\right)$. From the second condition of (17), if we choose $\epsilon := \frac{\sqrt{\beta_d C_a}\varepsilon}{M\sqrt{2M_\phi}}$, then we have $\hat{b}_0 \geq \frac{4M\sqrt{2M_\phi}\sigma_D\left(3M\sqrt{2M_\phi}\sigma_D+\sqrt{\beta_d C_a}\varepsilon\right)}{\beta_d C_a\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right)$. From the second condition of (18), if we choose $\epsilon := \sqrt{C_d}\|x_t-x_{t-1}\|$, then we have $\hat{b}_t \geq \frac{\left(6\sigma_D^2+2\sigma_D\sqrt{C_d}\|x_t-x_{t-1}\|\right)}{3C_d\|x_t-x_{t-1}\|^2}\cdot\log\left(\frac{p+q}{\delta}\right)$. Rounding $\hat{b}_t$, we obtain

$$\left\{\begin{array}{lcll}\hat{\bar{b}}_0 &:=& \left\lfloor\frac{4M\sqrt{2M_\phi}\sigma_D\left(3M\sqrt{2M_\phi}\sigma_D+\sqrt{\beta_d C_a}\varepsilon\right)}{\beta_d C_a\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right\rfloor &= \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right),\\ \hat{\bar{b}}_t &:=& \left\lfloor\frac{\left(6\sigma_D^2+2\sigma_D\sqrt{C_d}\|x_t-x_{t-1}\|\right)}{3C_d\|x_t-x_{t-1}\|^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right\rfloor &= \mathcal{O}\left(\frac{\sigma_D^2}{\|x_t-x_{t-1}\|^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right),\quad t\geq 1.\end{array}\right.$$

Since $\hat{b}_t \leq n$ for all $t\geq 0$, combining these conditions, we obtain $\hat{b}_t := \min\{n, \hat{\bar{b}}_t\}$ for $t\geq 0$, which proves the second part of (25).

For $t\geq 1$, we have $\|\widetilde{G}_M(x_{t-1})\| = M\|x_t-x_{t-1}\| > \varepsilon$. Otherwise, the algorithm has been terminated. Therefore, we can even bound b_t and $\hat{b}_t$ as

$$b_t \leq \frac{2M^2\sigma_F(3M^2\sigma_F+\sqrt{C_f}\varepsilon^2)}{3C^2\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right)\quad\text{and}\quad \hat{b}_t \leq \frac{M\left(6M\sigma_D^2+2\sigma_D\sqrt{C_d}\varepsilon\right)}{3C_d\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right).$$

From (20), to guarantee $\frac{1}{(T+1)}\sum_{t=0}^T\|\widetilde{G}_M(x_t)\|^2 \leq \varepsilon^2$, we impose $\frac{2M^2[\Psi(x_0)-\Psi^\star]}{C_a(T+1)} \leq \frac{\varepsilon^2}{2}$, which leads to $T := \left\lfloor\frac{4M^2[\Psi(x_0)-\Psi^\star]}{C_a\varepsilon^2}\right\rfloor$. Hence, the total number $\mathcal{T}_f$ of stochastic function evaluations $\mathbf{F}(\cdot,\xi)$ can be bounded by

$$\begin{array}{lcl}\mathcal{T}_f &:=& b_0+(T-1)b_t\\ &\leq& \left[\frac{32M_\phi M^2\sigma_F\left(48\sigma_F M_\phi M^2+C_a\varepsilon^2\right)}{3C_a^2\varepsilon^4}+\frac{8M^4\sigma_F(3M^2\sigma_F+\sqrt{C_f}\varepsilon^2)[\Psi(x_0)-\Psi^\star]}{3C^2C_a\varepsilon^6}\right]\cdot\log\left(\frac{p+1}{\delta}\right)\\ &=& \mathcal{O}\left(\frac{\sigma_F^2[\Psi(x_0)-\Psi^\star]}{\varepsilon^6}\cdot\log\left(\frac{p+1}{\delta}\right)\right).\end{array}$$

Similarly, the total number $\mathcal{T}_d$ of stochastic Jacobian evaluations $\mathbf{F}'(\cdot,\zeta)$ can be bounded by

$$\begin{array}{lcl}\mathcal{T}_d &:=& \hat{b}_0+(T-1)\hat{b}_t\\ &\leq& \left[\frac{4M\sqrt{2M_\phi}\sigma_D\left(3M\sqrt{2M_\phi}\sigma_D+\sqrt{\beta_d C_a}\varepsilon\right)}{\beta_d C_a\varepsilon^2}+\frac{4M^3[\Psi(x_0)-\Psi^\star]\left(6M\sigma_D^2+2\sigma_D\sqrt{C_d}\varepsilon\right)}{3C_dC_a\varepsilon^4}\right]\cdot\log\left(\frac{p+q}{\delta}\right)\\ &=& \mathcal{O}\left(\frac{\sigma_D^2[\Psi(x_0)-\Psi^\star]}{\varepsilon^4}\cdot\log\left(\frac{p+q}{\delta}\right)\right).\end{array}$$

Taking the upper bounds, these two last estimates prove (26). □

D.3. The Proof of Theorem 4.3: Convergence and Complexity Analysis of Algorithm 2 for (1)

Theorem. 4.3. *Suppose that Assumptions 1.1 and 1.2, and 4.1 are satisfied for (1). Let $\{x_t^{(s)}\}_{t=0\to m}^{s=1\to S}$ be generated by Algorithm 2 to solve (1). Let θ_F and m be chosen by (28), and the mini-batches b_s, $\hat{b}_s$, $b_t^{(s)}$, and $\hat{b}_t^{(s)}$ be set as in (29). Assume that the output $\widehat{x}_T$ of Algorithm 2 is chosen uniformly at random in $\{x_t^{(s)}\}_{t=0\to m}^{s=1\to S}$. Then:*

(a) *For a given tolerance $\varepsilon > 0$, the following bound holds*

$$\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] = \frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\mathbb{E}\left[\|\widetilde{G}_M(x_t)\|^2\right] \le \varepsilon^2. \tag{30}$$

(b) *The total number of iterations T to obtain $\mathbb{E}\left[\|\widetilde{G}_M(\widehat{x}_T)\|^2\right] \le \varepsilon^2$ is at most*

$$T := S(m+1) = \left\lfloor \frac{8M^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F\varepsilon^2} \right\rfloor = \mathcal{O}\left(\frac{1}{\varepsilon^2}\right).$$

Moreover, the total numbers $\mathcal{T}_f$ and $\mathcal{T}_d$ of stochastic function evaluations $\mathbf{F}(x_t,\xi)$ and stochastic Jacobian evaluations $\mathbf{F}'(x_t,\zeta)$, respectively do not exceed:

$$\left\{\begin{array}{lcl} \mathcal{T}_f &:=& \mathcal{O}\left(\frac{M_\phi^2\sigma_F^2}{\theta_F^2\varepsilon^4} + \frac{M^4M_\phi^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F^2\varepsilon^5}\right), \\ \mathcal{T}_d &:=& \mathcal{O}\left(\frac{M_\phi\sigma_D^2}{\theta_F\varepsilon^2} + \frac{M^2M_\phi\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F\varepsilon^3}\right). \end{array}\right. \tag{31}$$

Proof. We first analyze the inner loop. Using (15) with $x := x_t^{(s)}$ and $T_M(x) = x_{t+1}^{(s)}$, and then taking the expectation conditioned on $\mathcal{F}_{t+1}^{(s)} := \sigma(x_0^{(s)}, x_1^{(s)}, \cdots, x_t^{(s)})$, we have

$$\begin{array}{lcl} \mathbb{E}\left[\phi(F(x_{t+1}^{(s)})) \mid \mathcal{F}_{t+1}^{(s)}\right] &\le& \phi(F(x_t^{(s)})) - \frac{(2M - M_\phi(L_F+\beta_d))}{2}\mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 \mid \mathcal{F}_{t+1}^{(s)}\right] \\ && + \frac{L_\phi}{\xi_t^s}\mathbb{E}\left[\|F(x_t^{(s)}) - \widetilde{F}(x_t^{(s)})\|^2 \mid \mathcal{F}_{t+1}^{(s)}\right] \\ && + \frac{M_\phi}{2\beta_d}\mathbb{E}\left[\|F'(x_t^{(s)}) - \widetilde{J}(x_t^{(s)})\|^2 \mid \mathcal{F}_{t+1}^{(s)}\right] + M_\phi\xi_t^s, \end{array}$$

for any $\xi_t^s > 0$, where we use $2ab \le a^2 + b^2$ and the Jensen inequality $\left(\mathbb{E}\left[\|F(x_t^{(s)}) - \widetilde{F}(x_t^{(s)})\| \mid \mathcal{F}_{t+1}^{(s)}\right]\right)^2 \le \mathbb{E}\left[\|F(x_t^{(s)}) - \widetilde{F}(x_t^{(s)})\|^2 \mid \mathcal{F}_{t+1}^{(s)}\right]$ in the second line. Taking the full expectation both sides of the last inequality, and noting that $\Psi(x) = \phi(F(x))$, we obtain

$$\begin{array}{lcl} \mathbb{E}\left[\Psi(x_{t+1}^{(s)})\right] &\le& \mathbb{E}\left[\Psi(x_t^{(s)})\right] - \frac{C_g}{2}\mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] + \frac{L_\phi}{\xi_t^s}\mathbb{E}\left[\|F(x_t^{(s)}) - \widetilde{F}(x_t^{(s)})\|^2\right] + M_\phi\xi_t^s \\ && + \frac{M_\phi}{2\beta_d}\mathbb{E}\left[\|F'(x_t^{(s)}) - \widetilde{J}(x_t^{(s)})\|^2\right], \end{array} \tag{47}$$

where $C_g := 2M - M_\phi(L_F + \beta_d) > 0$, and $\beta_d > 0$ and $\xi_t^s > 0$ are given.

Next, from Lemma C.3, using the Lipschitz continuity of F' in Assumption 1.2, we have

$$\mathbb{E}\left[\|\widetilde{J}_t^{(s)} - F'(x_t^{(s)})\|^2\right] \le \mathbb{E}\left[\|\widetilde{J}_{t-1}^{(s)} - F'(x_{t-1}^{(s)})\|^2\right] + \frac{L_F^2}{\hat{b}_t^{(s)}}\mathbb{E}\left[\|x_t^{(s)} - x_{t-1}^{(s)}\|^2\right]. \tag{48}$$

Similarly, using Lemma C.3, we also have

$$\mathbb{E}\left[\|\widetilde{F}_t^{(s)} - F(x_t^{(s)})\|^2 \mid \mathcal{F}_{t+1}^{(s)}\right] \le \|\widetilde{F}_{t-1}^{(s)} - F(x_{t-1}^{(s)})\|^2 + \frac{1}{b_t}\mathbb{E}_\xi\left[\|\mathbf{F}(x_t,\xi) - \mathbf{F}(x_{t-1},\xi)\|^2\right].$$

Taking the full expectation both sides of this inequality, and using Assumption 4.1, we obtain

$$\mathbb{E}\left[\|\widetilde{F}_t^{(s)} - F(x_t^{(s)})\|^2\right] \le \mathbb{E}\left[\|\widetilde{F}_{t-1}^{(s)} - F(x_{t-1}^{(s)})\|^2\right] + \frac{M_F^2}{b_t^{(s)}}\mathbb{E}\left[\|x_t^{(s)} - x_{t-1}^{(s)}\|^2\right]. \tag{49}$$

Let us define a Lyapunov function as

$$\mathcal{L}(x_t^{(s)}) := \mathbb{E}\left[\Psi(x_t^{(s)})\right] + \frac{a_t^s}{2}\mathbb{E}\left[\|\widetilde{F}_t^{(s)} - F(x_t^{(s)}\|^2\right] + \frac{c_t^s}{2}\mathbb{E}\left[\|\widetilde{J}_t^{(s)} - F'(x_t^{(s)}\|^2\right], \tag{50}$$

for some $a_t^s > 0$ and $c_t^s > 0$.

Combining (47), (48), and (49), and then using the definition of $\mathcal{L}$ in (50), we have

$$\begin{aligned}
\mathcal{L}(x_{t+1}^{(s)}) &= \mathbb{E}\left[\Psi(x_{t+1}^{(s)})\right] + \tfrac{a_{t+1}^s}{2}\mathbb{E}\left[\|\widetilde{F}_{t+1}^{(s)} - F(x_{t+1}^{(s)}\|^2\right] + \tfrac{c_{t+1}^s}{2}\mathbb{E}\left[\|\widetilde{J}_{t+1}^{(s)} - F'(x_{t+1}^{(s)}\|^2\right] \\
&\le \mathbb{E}\left[\Psi(x_t^{(s)})\right] - \left[\frac{C_g}{2} - \frac{M_F^2 a_{t+1}^s}{2b_{t+1}^{(s)}} - \frac{L_F^2 c_{t+1}^s}{2\hat{b}_{t+1}^{(s)}}\right]\mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] + M_\phi \xi_t^s \\
&\quad + \left(\tfrac{a_{t+1}^s}{2} + \tfrac{L_\phi}{\xi_t^s}\right)\mathbb{E}\left[\|F(x_t^{(s)}) - \widetilde{F}(x_t^{(s)})\|^2\right] + \left(\tfrac{c_{t+1}^s}{2} + \tfrac{M_\phi}{2\beta_d}\right)\mathbb{E}\left[\|F'(x_t^{(s)}) - \widetilde{J}(x_t^{(s)})\|^2\right].
\end{aligned} \tag{51}$$

If we assume that

$$a_t^s \ge a_{t+1}^s + \frac{M_\phi}{\xi_t^s} \quad \text{and} \quad c_t^s \ge c_{t+1}^s + \frac{M_\phi}{\beta_d}, \tag{52}$$

then, from (51), we have

$$\mathcal{L}(x_{t+1}^{(s)}) \le \mathcal{L}(x_t^{(s)}) - \frac{\rho_{t+1}^s}{2}\mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] + M_\phi \xi_t^s, \tag{53}$$

where $\rho_{t+1}^s := C_g - \frac{M_F^2 a_{t+1}^s}{b_{t+1}^{(s)}} - \frac{L_F^2 c_{t+1}^s}{\hat{b}_{t+1}^{(s)}}$.

Let us first fix $\xi_t^s := \xi > 0$. Next, we choose $a_t^s := (m+1-t)\frac{M_\phi}{\xi}$ and $c_t^s := (m+1-t)\frac{M_\phi}{\beta_d}$. Clearly, $a_{m+1}^s = c_{m+1}^s = 0$ and they both satisfy the condition (52). Then, we choose $b_t^{(s)} := \frac{1}{\gamma_1}a_t^s = (m+1-t)\frac{M_\phi}{\gamma_1\xi}$ and $\hat{b}_t^{(s)} = \frac{1}{\gamma_2}c_t^s = \frac{M_\phi}{\beta_d\gamma_2}(m+1-t)$ for some $\gamma_1 > 0$ and $\gamma_2 > 0$. In this case, we have $\rho_t^s = C_g - M_F^2\gamma_1 - L_F^2\gamma_2 \equiv \theta_F > 0$ due to (28) by appropriately choosing γ_1 and γ_2. Consequently, (53) reduces to

$$\mathcal{L}(x_{t+1}^{(s)}) \le \mathcal{L}(x_t^{(s)}) - \frac{\theta_F}{2}\mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] + M_\phi\xi.$$

Summing up this inequality from $t = 0$ to $t = m$, we obtain

$$\frac{\theta_F}{2}\sum_{t=0}^m \mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] \le \mathcal{L}(x_0^{(s)}) - \mathcal{L}(x_{m+1}^{(s)}) + (m+1)M_\phi\xi.$$

Using the fact that $\widetilde{x}^{s-1} = x_0^{(s)}$ and $\widetilde{x}^s = x_{m+1}^{(s)}$, we have

$$\frac{\theta_F}{2}\sum_{t=0}^m \mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] \le \mathcal{L}(\widetilde{x}^{s-1}) - \mathcal{L}(\widetilde{x}^s) + (m+1)M_\phi\xi.$$

Summing up this inequality from $s = 1$ to S and multiplying the result by $\frac{2}{\theta_F S(m+1)}$, we obtain

$$\frac{1}{S(m+1)}\sum_{s=1}^S\sum_{t=0}^m \mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] \le \frac{2\left[\mathcal{L}(\widetilde{x}^0) - \mathcal{L}(\widetilde{x}^S)\right]}{\theta_F S(m+1)} + \frac{2M_\phi\xi}{\theta_F}. \tag{54}$$

Since $\mathcal{L}(\widetilde{x}^0) = \Psi(\widetilde{x}^0) + \frac{(m+1)M_\phi}{2\xi}\mathbb{E}\left[\|\widetilde{F}_0 - F(\widetilde{x}^0)\|^2\right] + \frac{(m+1)M_\phi}{2\beta_d}\mathbb{E}\left[\|\widetilde{J}_0 - F'(\widetilde{x}^0)\|^2\right]$ and $\mathcal{L}(\widetilde{x}^S) = \mathbb{E}\left[\Psi(\widetilde{x}^S\right] \ge \Phi^\star$, we obtain from (54) that

$$\begin{aligned}
\frac{1}{S(m+1)}\sum_{s=1}^S\sum_{t=0}^m \mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] &\le \frac{2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F S(m+1)} + \frac{M_\phi}{\xi\theta_F S}\mathbb{E}\left[\|\widetilde{F}_0 - F(\widetilde{x}^0)\|^2\right] \\
&\quad + \frac{M_\phi}{\theta_F\beta_d S}\mathbb{E}\left[\|\widetilde{J}_0 - F'(\widetilde{x}^0)\|^2\right] + \frac{2M_\phi\xi}{\theta_F}.
\end{aligned} \tag{55}$$

Note that $\mathbb{E}\left[\|\widetilde{F}_0 - F(\widetilde{x}^0)\|^2\right] \le \frac{\sigma_F^2}{b}$ and $\mathbb{E}\left[\|\widetilde{J}_0 - F'(\widetilde{x}^0)\|^2\right] \le \frac{\sigma_D^2}{\hat{b}}$ due to the choice of $b_s = b > 0$ and $\hat{b}_s = \hat{b} > 0$ at Step 4 of Algorithm 2. Hence, we can further bound (55) as

$$\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\mathbb{E}\left[\|x_{t+1}^{(s)} - x_t^{(s)}\|^2\right] \le \frac{2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F S(m+1)} + \frac{M_\phi\sigma_F^2}{\xi\theta_F S b} + \frac{M_\phi\sigma_D^2}{\theta_F\beta_d S\hat{b}} + \frac{2M_\phi\xi}{\theta_F}.$$

Since $\|\widetilde{G}_M(x_t^{(s)})\| = M\|x_{t+1}^{(s)} - x_t^{(s)}\|$, to guarantee $\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\mathbb{E}\left[\|\widetilde{G}_M(x_t^{(s)})\|^2\right] \le \varepsilon^2$ for a given tolerance $\varepsilon > 0$, we need to set

$$\frac{2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F S(m+1)} + \frac{M_\phi\sigma_F^2}{\xi\theta_F S b} + \frac{M_\phi\sigma_D^2}{\theta_F\beta_d S\hat{b}} + \frac{2M_\phi\xi}{\theta_F} = \frac{\varepsilon^2}{M^2}.$$

Let us break this condition into

$$\frac{2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F S(m+1)} = \frac{\varepsilon^2}{4M^2} \quad \text{and} \quad \frac{M_\phi\sigma_F^2}{\xi\theta_F S b} = \frac{M_\phi\sigma_D^2}{\theta_F\beta_d S\hat{b}} = \frac{2M_\phi\xi}{\theta_F} = \frac{\varepsilon^2}{4M^2}.$$

Hence, we can choose $\xi := \frac{\theta_F\varepsilon^2}{8M^2M_\phi}$, $\hat{b} := \frac{4M_\phi\sigma_D^2}{\theta_F\beta_d M^2 S\varepsilon^2}$, $b := \frac{2M_\phi^2\sigma_F^2}{\theta_F^2 M^2 S\varepsilon^4}$, and $S(m+1) = \frac{8M^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F\varepsilon^2}$.

Now, let us choose $m+1 := \frac{\hat{C}}{\varepsilon}$ for some constant $\hat{C} > 0$. Then, we can estimate the total number $\mathcal{T}_f$ of stochastic function evaluations $\mathbf{F}(x_t^{(s)}, \xi)$ as follows:

$$\begin{aligned}
\mathcal{T}_f &:= \textstyle\sum_{s=1}^{S} b_s + \sum_{s=1}^{S}\sum_{t=0}^{m} b_t^{(s)} = Sb + \frac{M_\phi}{\gamma_1\xi}\sum_{s=1}^{S}\sum_{t=0}^{m}(m+1-t)\\
&= \frac{2M_\phi^2\sigma_F^2}{\theta_F^2 M^2\varepsilon^4} + \frac{8M^2M_\phi^2}{\gamma_1\theta_F\varepsilon^2}\cdot\frac{S(m+1)(m+2)}{2}\\
&= \frac{2M_\phi^2\sigma_F^2}{\theta_F^2 M^2\varepsilon^4} + \frac{8M^2M_\phi^2}{\gamma_1\theta_F\varepsilon^2}\cdot\frac{8M^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F\varepsilon^2}\cdot\frac{\hat{C}+\varepsilon}{2\varepsilon}\\
&= \mathcal{O}\left(\frac{M_\phi^2\sigma_F^2}{\theta_F^2\varepsilon^4} + \frac{M^4M_\phi^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F^2\varepsilon^5}\right).
\end{aligned}$$

Similarly, the total number $\mathcal{T}_d$ of stochastic Jacobian evaluations $\mathbf{F}'(x_t^{(s)}, \zeta)$ can be bounded as

$$\begin{aligned}
\mathcal{T}_d &:= \textstyle\sum_{s=1}^{S}\hat{b}_s + \sum_{s=1}^{S}\sum_{t=0}^{m}\hat{b}_t^{(s)} = S\hat{b} + \frac{M_\phi S}{\beta_d\gamma_2}\sum_{t=0}^{m}(m+1-t)\\
&\le \frac{4M_\phi\sigma_D^2}{\theta_F\beta_d M^2\varepsilon^2} + \frac{8M^2M_\phi\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\beta_d\gamma_2\theta_F\varepsilon^2}\cdot\frac{\hat{C}+\varepsilon}{2\varepsilon}\\
&= \mathcal{O}\left(\frac{M_\phi\sigma_D^2}{\theta_F\varepsilon^2} + \frac{M^2M_\phi\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\theta_F\varepsilon^3}\right).
\end{aligned}$$

Hence, taking the upper bounds, we have proven (31). □

E. Solution Routines for Computing Gauss-Newton Search Directions

One main step of SGN methods is to compute the Gauss-Newton direction by solving the subproblem (11). This subproblem is also called a **prox-linear operator**, which can be rewritten as

$$\min_{d\in\mathbb{R}^p}\left\{\phi(\widetilde{F}_t + \widetilde{J}_t d) + \hat{g}(d) + \tfrac{M}{2}\|d\|_2^2\right\}, \tag{56}$$

where $\widetilde{F}_t \approx F(x_t)$, $\widetilde{J}_t \approx F'(x_t)$, $d := x - x_t$, ϕ is convex, $\hat{g}(d) := g(x_t + d)$, and $M > 0$ is given. This is a basic convex problem, and we can apply different methods to solve it. Here, we describe two methods for solving (56).

E.1. Accelerated Dual Proximal-Gradient Method

For accelerated dual proximal-gradient method, we consider the case $\hat{g}(d) = 0$ for simplicity. Using Fenchel's conjugate of ϕ, we can write $\phi(\widetilde{F}_t + \widetilde{J}_t d) = \max\left\{\langle \widetilde{F}_t + \widetilde{J}_t d, u\rangle - \phi^*(u)\right\}$. Assume that strong duality holds for (56), then using this expression, we can write it as

$$\min_d \max_u \left\{\langle \widetilde{F}_t + \widetilde{J}_t d, u\rangle - \phi^*(u) + \frac{M}{2}\|d\|_2^2\right\} \quad \Leftrightarrow \quad \max_u \left\{ \min_d \left\{\langle \widetilde{F}_t + \widetilde{J}_t d, u\rangle + \frac{M}{2}\|d\|_2^2\right\} - \phi^*(u)\right\}.$$

Solving the inner problem $\min_d \left\{\langle \widetilde{F}_t + \widetilde{J}_t d, u\rangle + \frac{M}{2}\|d\|_2^2\right\}$, we obtain $d^*(u) := -\frac{1}{M}\widetilde{J}^\top u$. Substituting it into the objective, we eventually obtain the dual problem as follows:

$$\min_u \left\{\frac{1}{2M}\|\widetilde{J}_t^\top u\|_2^2 - \langle \widetilde{F}_t, u\rangle + \phi^*(u)\right\}. \tag{57}$$

We can solve this problem by an accelerated proximal-gradient method (Beck & Teboulle, 2009; Nesterov, 2004), which is described as follows.

Algorithm 3 (Accelerated Dual Proximal-Gradient **(ADPG)**)

1: **Initialization:** Choose $u_0 \in \mathbb{R}^m$. Set $\tau_0 := 1$ and $\hat{u}_0 := u_0$. Evaluate $L := \frac{1}{M}\|\widetilde{J}_t^\top \widetilde{J}_t\|$.
2: **For** $k := 0, \cdots, k_{\max}$ **do**
3: $\quad u_{k+1} := \mathrm{prox}_{(1/L)\phi^*}\left(\hat{u}_k - \frac{1}{L}(\frac{1}{M}\widetilde{J}_t\widetilde{J}_t^\top \hat{u}_k - \widetilde{F}_t)\right)$.
4: $\quad \tau_{k+1} := \frac{1+\sqrt{1+4\tau_k^2}}{2}$.
5: $\quad \hat{u}_{k+1} := u_{k+1} + \left(\frac{\tau_k - 1}{\tau_{k+1}}\right)(u_{k+1} - u_k)$.
6: **End For**
7: **Output:** Reconstruct $d^* := -\frac{1}{M}\widetilde{J}_t^\top u_k$ as an approximate solution of (56).

Note that in Algorithm 3, we use the proximal operator $\mathrm{prox}_{\lambda\phi^*}$ of ϕ^*. However, by Moreau's identity, $\mathrm{prox}_{\lambda\phi^*}(v) + \lambda\mathrm{prox}_{\phi/\lambda}(v/\lambda) = v$, we can again use the proximal operator $\mathrm{prox}_{\phi/\lambda}$ of ϕ.

E.2. Primal-Dual First-Order Methods

We can apply any primal-dual algorithm from the literature (Bauschke & Combettes, 2017; Chambolle & Pock, 2011; Esser, 2010; Goldstein et al., 2013; Tran-Dinh et al., 2018; Tran-Dinh, 2019) to solve (56). Here, we describe the well-known Chambolle-Pock's primal-dual method (Chambolle & Pock, 2011) to solve (56).

Let us define $\hat{\phi}(z) := \phi(z + F_k)$ and $\hat{\psi}(d) := \hat{g}(d) + \frac{M}{2}\|d\|^2$. Since (56) is strongly convex with the strong convexity parameter $\mu_{\hat{\psi}} := M$, we can apply the strongly convex primal-dual variant as follows.

Choose $\sigma_0 > 0$ and $\tau_0 > 0$ such that $\tau_0\sigma_0 \le \frac{1}{\|\widetilde{J}_t^\top \widetilde{J}_t\|}$. For example, we can choose $\sigma_0 = \tau_0 = \frac{1}{\|\widetilde{J}_t\|}$, or we choose $\sigma_0 > 0$ first, and choose $\tau_0 := \frac{1}{\sigma_0\|\widetilde{J}_t^\top \widetilde{J}_t\|}$. Choose $d_0 \in \mathbb{R}^p$ and $u_0 \in \mathbb{R}^m$ and set $\bar{d}_0 := d_0$. Then, at each iteration $k \ge 0$, we update

$$\begin{cases} u_{k+1} & := & \mathrm{prox}_{\sigma_k\hat{\phi}^*}\left(u_k + \sigma_k\widetilde{J}_t\bar{d}_k\right), \\ d_{k+1} & := & \mathrm{prox}_{\tau_k\hat{\psi}}\left(d_k - \tau_k\widetilde{J}_t^\top u_{k+1}\right), \\ \theta_k & := & 1/\sqrt{1 + 2M\tau_k}, \\ \tau_{k+1} & := & \theta_k\tau_k, \\ \sigma_{k+1} & := & \sigma_k/\theta_k, \\ \bar{d}_{k+1} & := & d_{k+1} + \theta_k(d_{k+1} - d_k). \end{cases} \tag{58}$$

Alternatively to the Accelerated Dual Proximal-Gradient and the primal-dual methods, we can also apply the alternating direction method of multipliers (ADMM) to solve (56). However, this method requires to solve a linear system, that may not scale well when the dimension p is large.

F. Details of The Experiments in Section 5

In this supplementary document, we provide the details of our experiments in Section 5, including modeling, data generating routines, and experiment configurations. We also provide more experiments for both examples. All algorithms are implemented in Python 3.6 running on a Macbook Pro with 2.3 GHz Quad-Core, 8 GB RAM and on a Linux-based computing node, called Longleaf, where each node has 24 physical cores, 2.50 GHz processors, and 256 GB RAM.

F.1. Stochastic Nonlinear Equations

Our goal is to solve the following nonlinear equation in expectation as described in Subsection 5.1:

$$F(x) = 0, \quad \text{where} \quad F(x) := \mathbb{E}_\xi \left[\mathbf{F}(x, \xi)\right]. \tag{59}$$

Here, $\mathbf{F}$ is a stochastic vector function from $\mathbb{R}^p \times \Omega \to \mathbb{R}^q$. As discussed in the main text, (59) covers the first-order optimality condition $\mathbb{E}_\xi \left[\nabla_x \mathbf{G}(x, \xi)\right] = 0$ of a stochastic optimization problem $\min_x \mathbb{E}_\xi \left[\mathbf{G}(x, \xi)\right]$ as a special case. More generally, it also covers the KKT condition of a stochastic optimization problem with equality constraints. However, these problems may not have stationary point, which leads to an inconsistency of (59). As a remedy, we can instead consider

$$\min_x \left\{\Psi(x) := \|\mathbb{E}_\xi \left[\mathbf{F}(x, \xi)\right]\|\right\}, \tag{60}$$

for a given norm $\|\cdot\|$ (e.g., ℓ_1-norm or ℓ_2-norm). Problem (59) also covers the expectation formulation of stochastic nonlinear equations such as stochastic ODEs or PDEs.

In our experiment from Subsection 5.1, we only consider one instance of (60) by choosing $q = 4$ and $\mathbf{F}_j$ $(j = 1, \cdots, q)$ as

$$\begin{cases} \mathbf{F}_1(x, \xi_i) & := & (1 - \tanh(y_i(a_i^\top x + b_i)), \\ \mathbf{F}_2(x, \xi_i) & := & \left(1 - \frac{1}{1+\exp(-y_i(a_i^\top x + b_i))}\right)^2, \\ \mathbf{F}_3(x, \xi_i) & := & \log(1 + \exp(-y_i(a_i^\top x + b_i))) - \log(1 + \exp(-y_i(a_i^\top x + b_i) - 1)), \\ \mathbf{F}_4(x, \xi_i) & := & \log(1 + (y_i(a_i^\top x + b_i) - 1)^2), \end{cases} \tag{61}$$

where a_i is the i-row of an input matrix $A \in \mathbb{R}^{n\times p}$, $y \in \{-1, 1\}^n$ is a vector of labels, $b \in \mathbb{R}^n$ is a bias vector in binary classification, and $\xi_i := (a_i, b_i, y_i)$. Note that the binary classification problem with nonconvex loss has been widely studied in the literature, including Zhao et al. (2010), where one aims at solving:

$$\min_{x\in\mathbb{R}^p} \left\{H(x) := \frac{1}{n}\sum_{i=1}^n \ell(y_i(a_i^T x + b_i))\right\}, \tag{62}$$

for a given loss function ℓ. If ℓ is nonnegative, then instead of solving (62), we can solve $\min_x |H(x)|$. If we have q different losses ℓ_j for $j = 1, \cdots, q$ and we want to solve q problems of the form (62) for different losses simultaneously, then we can formulate such a problem into (60) to have $\min_x \|\mathbf{H}(x)\|$, where $\mathbf{H}(x) := (H_1(x), H_2(x), \cdots, H_q(x))^\top$. Since we use different losses, under the formulation (60), we can view it as a binary classification task with an averaging loss.

Table 1. Hyper-parameter configurations for the two algorithms on all datasets when using the $\|\cdot\|_2$ loss.

Algorithm	w8a			ijcnn1			covtype			url_combined		
	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations
SGN	256	512		512	1,024		1,024	4,096		20,000	50,000	
SGN2	64	128	2,000	128	256	1,000	256	512	2000	5,000	10,000	5,000
	a9a			rcv1_train.binary			real-sim			skin_nonskin		
	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations
SGN	512	1,024		512	1,024		1,024	4,096		512	1024	
SGN2	64	128	2000	128	256	1,000	256	512	2,000	128	256	5,000

Datasets. We test three algorithms: GN, SGN, and SGN2 on four real datasets: `w8a` ($\boldsymbol{n = 49,749; p = 300}$), `ijcnn1` ($\boldsymbol{n = 91,701; p = 22}$), `covtype` ($\boldsymbol{n = 581,012; p = 54}$), and `url_combined` ($\boldsymbol{n = 2,396,130; p = 3,231,961}$) from LIBSVM.

Table 2. Hyper-parameter configurations for the four algorithms on 4 datasets when using the Huber loss.

Algorithm	w8a			ijcnn1			covtype			url_combined		
	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations
SGN	256	512		512	1,024		512	1,024		20,000	50,000	
SCGD	256	512		512	1,024		512	1,024		20,000	50,000	
SGN2	64	128	5,000	128	256	2,000	128	256	5,000	5,000	10,000	5,000
N-SPIDER	64	128	5,000	128	256	2,000	128	256	5,000	5,000	10,000	5,000
	a9a			rcv1_train.binary			real-sim			news20.binary		
	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations
SGN	128	256		128	512		256	512		128	512	
SCGD	1,024	2,048		128	512		256	512		128	512	
SGN2	64	128	2,000	64	128	5,000	64	128	5,000	64	128	5,000
N-SPIDER	64	128	2,000	64	128	5,000	64	128	5,000	64	128	5,000

Parameter configuration. We can easily check that F defined by (61) satisfies Assumption 1.1 and Assumption 2. However, we do not accurately estimate the Lipschitz constant of F' since it depends on the dataset. We were instead experimenting with different choices of the parameter M and ρ, and eventually fix $\rho := 1$ and $M := 1$ for our tests. We also choose the mini-batch sizes for both $\widetilde{F}$ and $\widetilde{J}$ in SGN and SGN2 by sweeping over the set of $\{64, 128, 256, 512, 1024, 2048, 4096, 8192\}$ to estimate the best ones. Table 1 presents the chosen parameters for the instance when $\phi = ||\cdot||_2$.

In the case of smooth ϕ, i.e., using Huber loss, we add two competitors: N-SPIDER (Yang et al., 2019, Algorithm 3) and SCGD Wang et al. (2017a, Algorithm 1). The learning rates of N-SPIDER and SCGD are tuned from a set of different values: $\{0.01, 0.05, 0.1, 0.5, 1, 2\}$. Eventually we obtain $\eta := 1.0$ and set $\varepsilon := 10^{-1}$ for N-SPIDER, see (Yang et al., 2019, Algorithm 3). For SCGD, we use $\beta_k := 1$ and $\alpha_k := 1$, see Wang et al. (2017a, Algorithm 1). The mini-batch sizes of these algorithm are chosen using similar search as in the previous case. Table 2 reveals the parameter configuration of the algorithms when using the Huber loss.

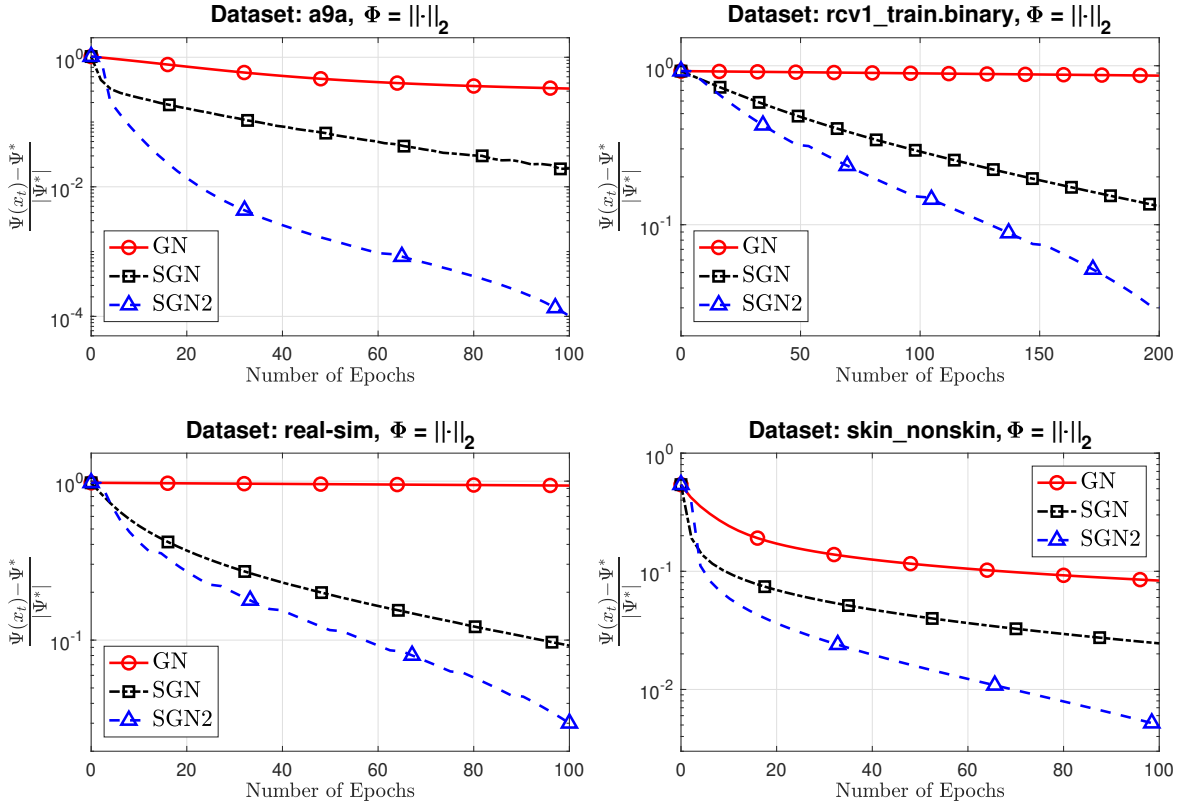

Figure 6. The performance of three algorithms on additional real datasets when $\phi(\cdot) = \|\cdot\|_2$.

Additional Experiments. When $\phi(\cdot) = \|\cdot\|_2$, we also run these algorithms on other classification datasets

from LIBSVM: `a9a` ($\boldsymbol{n = 32,561; p = 123}$), `rcv1_train.binary` ($\boldsymbol{n = 20,242; p = 47,236}$), `real-sim` ($\boldsymbol{n = 72,309; p = 20,958}$), and `skin_nonskin` ($\boldsymbol{n = 245,057; p = 3}$). We set $M := 1$ and $\rho := 1$ for three datasets. Other parameters are obtained via grid search and the results are shown in Table 1. The performance of three algorithms on these datasets are presented in Figure 6.

SGN2 appears to be the best among the 3 algorithms while SGN is much better than the baseline GN. SGN appears to have advantage in the early stage but SGN2 makes better progress later on.

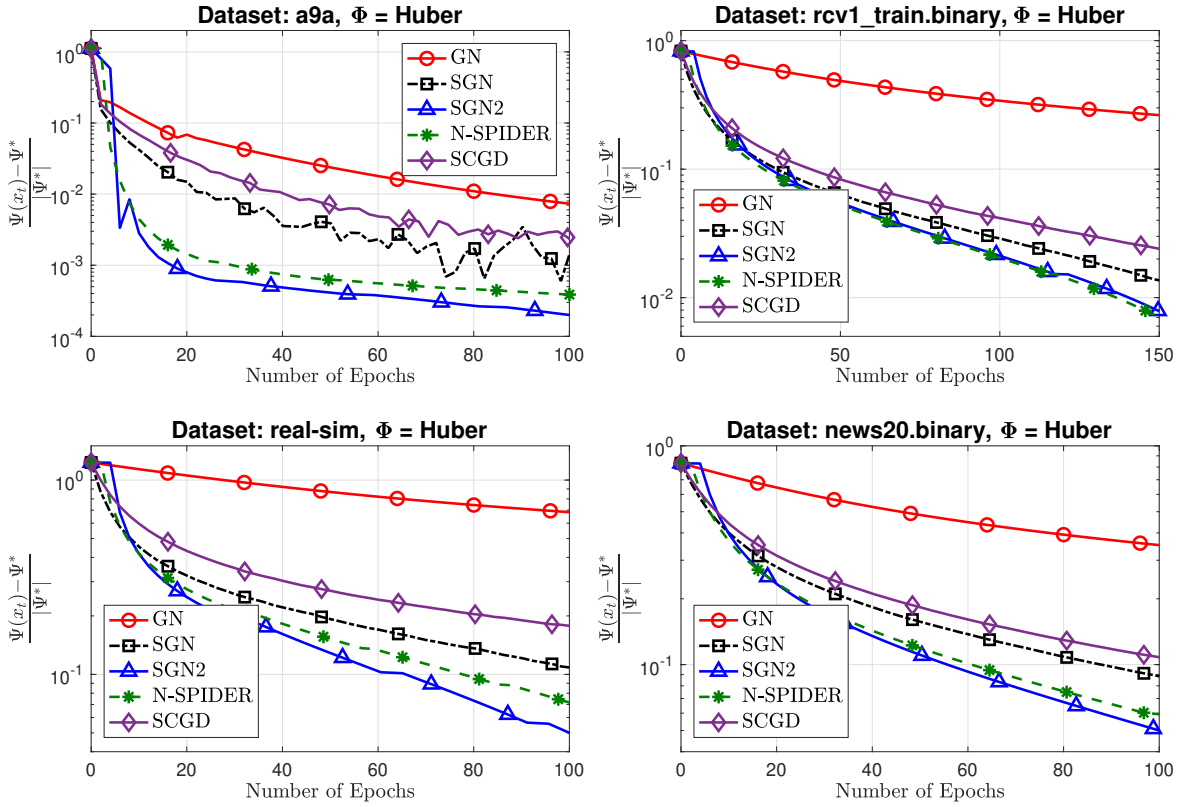

Figure 7. The performance of three algorithms on additional real datasets when using Huber loss.

In addition, we also run 5 algorithms on these datasets in the smooth case when using the Huber loss. We still tune the parameters for these algorithms and obtain the learning rate of 1.0 for both N-SPIDER and SCGD. We again use $\varepsilon = 10^{-1}$ for N-SPIDER. More details about other parameters selection are presented in Table 2 and the performance of these algorithms are shown in Figure 7.

From Figure 7, SGN2 performs better than other algorithms in most cases while N-SPIDER is better than SGN and somewhat comparable with SGN2 in the `rcv1_train.binary` and `news20.binary` datasets. SGN and SCGD appear to have similar behavior, but SGN is slightly better than SCGD in these datasets.

F.2. Optimization Involving Expectation Constraints

We consider an optimization problem involving expectation constraints as described in (34). As mentioned, this problem has various applications in different fields, including optimization with conditional value at risk (CVaR) constraints and metric learning, see, e.g., Lan & Zhou (2016) for detailed discussion.

Instead of solving the constrained setting (34), we consider its exact penalty formulation (35):

$$\min_{x\in\mathbb{R}^p} \Big\{\Psi(x) := g(x) + \phi(\mathbb{E}_\xi\left[\mathbf{F}(x,\xi)\right])\Big\}, \tag{35}$$

where $\phi(u) := \rho\sum_{i=1}^q [u_i]_+$ with $[u]_+ := \max\{0, u\}$ is a penalty function, and $\rho > 0$ is a given penalty parameter. It is

well-known that under mild conditions and ρ sufficiently large (e.g., $\rho > \|y^\star\|^*$, the dual norm of the optimal Lagrange multiplier $y^\star$), if $x^\star$ is a stationary point of (35) and it is feasible to (34), then it is also a stationary point of (34).

As a concrete instance of (34), we solve the following asset allocation problem studied in Rockafellar & Uryasev (2000); Lan & Zhou (2016):

$$\left\{\begin{array}{ll} \min\limits_{z\in\mathbb{R}^p,\tau\in[\underline{\tau},\bar{\tau}]} & -c^\top z \\ \text{s.t} & \tau + \frac{1}{\beta n}\sum_{i=1}^n [-\xi_i^\top z - \tau]_+ \le 0, \\ & z \in \Delta_p := \left\{\hat{z}\in\mathbb{R}^p_+ \mid \sum_{i=1}^p \hat{z}_i = 1\right\}. \end{array}\right. \tag{63}$$

Here, Δ_p denotes the standard simplex in $\mathbb{R}^p$, and $[\underline{\tau},\bar{\tau}]$ is a given range of τ. The exact penalty formulation of (63) is given by (36):

$$\min_{z\in\Delta^p,\tau\in[\underline{\tau},\bar{\tau}]} \left\{-c^\top z + \phi\left(\tau + \frac{1}{\beta n}\sum_{i=1}^n [-\xi_i^\top z - \tau]_+\right)\right\}, \tag{36}$$

where $\phi(u) := \rho[u]_+$ with given $\rho > 0$. However, since $[-\xi_i^\top z - \tau]_+$ is nonsmooth, we smooth it by $\sqrt{(\xi_i^\top z + \tau)^2 + \gamma^2} - \gamma - \xi_i^\top z - \tau$ for sufficiently small value of $\gamma > 0$. Hence, (36) can be approximated by

$$\min_{z\in\Delta_p,\tau\in[\underline{\tau},\bar{\tau}]} \left\{-c^\top z + \phi\left(\tau + \frac{1}{\beta n}\sum_{i=1}^n \left[\sqrt{(\xi_i^\top z + \tau)^2 + \gamma^2} - \gamma - \xi_i^\top z - \tau\right]\right)\right\}. \tag{64}$$

If we introduce $x := (z,\tau)$, $\mathbf{F}(x,\xi) := \tau + \frac{1}{2\beta}\left(\sqrt{(\xi_i^\top z + \tau)^2 + \gamma^2} - \gamma - \xi_i^\top z - \tau\right)$ for $i = 1,\cdots,n$, and $g(x) = -c^\top z + \delta_{\Delta_p\times[\underline{\tau},\bar{\tau}]}(x)$, where $\delta_{\mathcal{X}}$ is the indicator of $\mathcal{X}$, then we can reformulate (64) into (3). It is obvious to check that $\mathbf{F}(\cdot,\xi)$ is Lipschitz continuous with $M_i := 1 + \frac{\|\xi_i\|+1}{\beta\gamma}$ and its gradient $\mathbf{F}'(\cdot,\zeta)$ is also Lipschitz continuous with $L_i := \frac{\|\xi_i\|^2}{2\beta\gamma}$. Hence, Assumptions 1.1 and 4.1 hold.

Datasets. We consider both synthetic and US stock datasets. For the synthetic datasets, we follow the procedures from Lan et al. (2012) to generate the data with $n = 10^5$ and $p \in \{300, 500, 700\}$. We obtain real datasets of US stock prices for 889, 865, and 500 types of stocks as described, e.g., Sun & Tran-Dinh (2019). Then, we apply a bootstrap strategy to resample in order to obtain three corresponding new datasets of sizes $n = 10^5$.

Table 3. Hyper-parameter configuration of the two algorithms on 6 datasets in the asset allocation example.

Algorithm	Synthetic: p = 300			Synthetic: p = 500			Synthetic: p = 700		
	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations
SGN	1,024	2,048		1,024	2,048		1,024	2,048	
SGN2	128	256	5,000	128	256	2,000	256	512	2,000
Algorithm	US Stock 1: p = 889			US Stock 1: p = 865			US Stock 1: p = 500		
	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations	$\hat{b}_t$	b_t	Inner Iterations
SGN	512	1,024		512	1,024		512	1,024	
SGN2	128	256	5,000	128	256	5,000	128	256	5,000

Parameter selection. We fix the smoothness parameter $\gamma := 10^{-3}$ and choose the range $[\underline{\tau},\bar{\tau}]$ to be $[0,1]$. The parameter $\beta := 0.1$ as discussed in Lan & Zhou (2016). Note that we do not use the theoretical values for M as in our theory since that value is obtained in the worst-case. We were instead experimenting different values for the penalty parameter ρ and M, and eventually get $\rho := 5$ and $M := 5$ as default values for this example.

Experiment setup. We implement our algorithms: SGN and SGN2, and also a baseline variant, the deterministic GN scheme (i.e., we exactly evaluate F and its Jacobian using the full batches) as in the first example. Similar to the first example, we sweep over the same set of possible mini-batch sizes, and the chosen parameters are reported in Table 3.

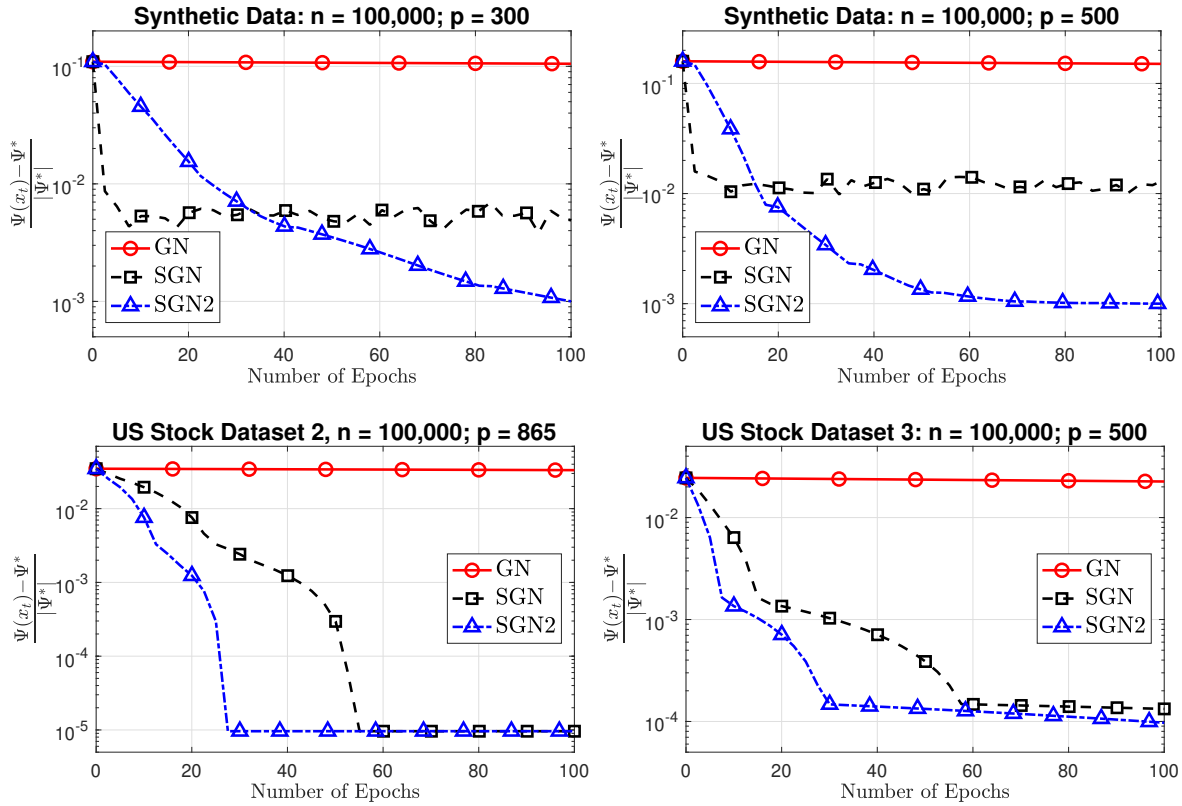

Figure 8. The performance of the three algorithms on two synthetic and two real datasets.

Additional experiments. We run three algorithms: GN, SGN, and SGN2 with 3 synthetic datasets, where the first one was reported in Figure 2 of the main text. We also use two other US Stock datasets and the performance of three algorithms on these synthetic and real datasets are revealed in Figure 8.

Clearly, SGN2 is the best, while SGN still outperforms GN in these two datasets. We believe that this experiment confirms our theoretical results presented in the main text.

G. Convergence of Algorithm 2 for the finite-sum case (2) without Assumption 4.1

Although Theorem 1 significantly improves stochastic oracle complexity of Algorithm 2 compared to Theorem 4.1, it requires additional assumption, Assumption 4.1. Assumption 4.1 is usually used in compositional models such as neural network and parameter estimations. However, we still attempt to establish a convergence and complexity result for Algorithm 2 to solve (2) without Assumption 4.1 in the following theorem.

Theorem G.1. *Suppose that Assumptions 1.1 and 1.2 are satisfied for (2). Let $\{x_t^{(s)}\}_{t=0\to m}^{s=1\to S}$ be generated by Algorithm 2 to solve (2). Let the mini-batches b_s, $\hat{b}_s$, $b_t^{(s)}$, and $\hat{b}_t^{(s)}$ be set as follows:*

$$\left\{\begin{array}{lcl} b_s & := & \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right)\right),\\ b_t^{(s)} & := & \mathcal{O}\left(\frac{m^2}{\varepsilon^2}\cdot\log\left(\frac{p+1}{\delta}\right)\right),\\ \hat{b}_s & := & \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right),\\ \hat{b}_t^{(s)} & := & \mathcal{O}\left(m^2\cdot\log\left(\frac{p+1}{\delta}\right)\right). \end{array}\right. \tag{65}$$

Then, with probability at least $1-\delta$, the following statements hold:

(a) *The following bound holds*

$$\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\|\widetilde{G}_M(x_t)\|^2 \le \mathcal{O}\left(\varepsilon^2\right).$$

(b) *The total number of iterations T to achieve*

$$\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\|\widetilde{G}_M(x_t)\|^2 \le \varepsilon^2$$

is at most $T := S(m+1) = \mathcal{O}\left(\frac{[\Psi(\widetilde{x}^0)-\Psi^\star]}{\varepsilon^2}\right)$. Moreover, the total stochastic oracle calls $\mathcal{T}_f$ and $\mathcal{T}_d$ to approximate F and its Jacobian F', respectively do not exceed

$$\left\{\begin{array}{lcl} \mathcal{T}_f &:=& \mathcal{O}\left(\left(\frac{\sigma_F^2}{\varepsilon^5}+\frac{1}{\varepsilon^6}\right)\left[\Psi(\widetilde{x}^0)-\Psi^\star\right]\cdot\log\left(\frac{p+1}{\delta}\right)\right),\\ \mathcal{T}_d &:=& \mathcal{O}\left(\left(\frac{\sigma_D^2}{\varepsilon^3}+\frac{1}{\varepsilon^4}\right)\left[\Psi(\widetilde{x}^0)-\Psi^\star\right]\cdot\log\left(\frac{p+q}{\delta}\right)\right).\end{array}\right.$$

Remark G.1. Although we do not gain an improvement on the worst-case oracle complexity through Theorem G.1, we observe in our experiment that Algorithm 1.2 highly outperforms **SGN**. There could be an artifact in our proof of Theorem G.1.

Proof. We first analyze the inner loop of Algorithm 2. For simplicity of notation, we drop the superscript $^{(s)}$ in the following derivations until it is recalled. We first verify the conditions (18) if we use the SARAH estimators (27) for $F'(x_t)$ and $F(x_t)$. Let $\mathcal{F}_t := \sigma(x_0, x_1, \cdots, x_{t-1})$ be the σ-field generated by $\{x_0, x_1, \cdots, x_{t-1}\}$. We define $X_i := F_i'(x_t) - F_i'(x_{t-1}) - [F'(x_t) - F'(x_{t-1})]$. Then, clearly, conditioned on $\mathcal{F}_t$, we have $\{X_i\}_{i\in\hat{\mathcal{B}}_t}$ is mutually independent and $\mathbb{E}[X_i \mid \mathcal{F}_{t-1}] = 0$. Moreover, by Assumption 1.2, we have

$$\begin{array}{lcl}\|X_i\| &=& \|F_i'(x_t) - F_i'(x_{t-1}) - [F'(x_t) - F'(x_{t-1})]\|\\ &\le& \|F_i'(x_t) - F_i'(x_{t-1})\| + \|F'(x_t) - F'(x_{t-1})\|\\ &\le& 2L_F\|x_t - x_{t-1}\| =: \hat{\sigma}_t.\end{array}$$

We consider $Z_t := \frac{1}{\hat{b}_t}\sum_{i\in\hat{\mathcal{B}}_t}[F_i'(x_t) - F_i'(x_{t-1}) - F'(x_t) + F'(x_{t-1})] = \frac{1}{\hat{b}_t}\sum_{i\in\hat{\mathcal{B}}_t} X_i$. We have

$$\sigma_X^2 := \max\left\{\left\|\sum_{i\in\hat{\mathcal{B}}_t}\mathbb{E}\left[X_iX_i^\top \mid \mathcal{F}_{t-1}\right]\right\|, \left\|\sum_{i\in\hat{\mathcal{B}}_t}\mathbb{E}\left[X_i^\top X_i \mid \mathcal{F}_{t-1}\right]\right\|\right\} \le \sum_{i\in\hat{\mathcal{B}}_t}\mathbb{E}\left[\|X_i\|^2 \mid \mathcal{F}_{t-1}\right] \le \hat{b}_t\hat{\sigma}_t^2.$$

For any $\hat{\epsilon} > 0$, we can apply Lemma C.1 to obtain the following bound

$$\begin{array}{lcl}\mathbf{Prob}\left(\|Z_t\| \le \hat{\epsilon}\hat{\sigma}_t\right) &=& \mathbf{Prob}\left(\|\sum_{i\in\hat{\mathcal{B}}_t} X_i\| \le \hat{b}_t\hat{\epsilon}\hat{\sigma}_t\right)\\ &\ge& 1 - (p+q)\exp\left(-\frac{3\hat{b}_t^2\hat{\epsilon}^2\hat{\sigma}_t^2}{6\hat{b}_t\hat{\sigma}_t^2 + 2\hat{\sigma}_t\hat{b}_t\hat{\epsilon}\hat{\sigma}_t}\right)\\ &=& 1 - (p+q)\exp\left(-\frac{3\hat{b}_t\hat{\epsilon}^2}{6+2\hat{\epsilon}}\right).\end{array}$$

Hence, if we choose $\delta \ge (p+q)\exp\left(-\frac{3\hat{b}_t\hat{\epsilon}^2}{6+2\hat{\epsilon}}\right)$ and $\delta \le 1$, we obtain $\mathbf{Prob}\left(\|Z_t\| \le \hat{\epsilon}\hat{\sigma}_t\right) \ge 1 - \delta$ for all $t \ge 0$. The condition in $\hat{b}_t$ leads to

$$\hat{b}_t \ge \frac{6+2\hat{\epsilon}}{3\hat{\epsilon}^2}\cdot\log\left(\frac{p+q}{\delta}\right).$$

By the update (27), we have $\widetilde{J}_t - F(x_t) = [\widetilde{J}_{t-1} - F'(x_{t-1})] + \frac{1}{\hat{b}_t}\sum_{i\in\hat{\mathcal{B}}_t}[F_i'(x_t) - F_i'(x_{t-1}) - F'(x_t) + F'(x_{t-1})] = [\widetilde{J}_{t-1} - F'(x_{t-1})] + Z_t$. Hence, by the triangle inequality, we get

$$\|\widetilde{J}_t - F'(x_t)\| = \|\widetilde{J}_0 - F'(x_0) + \sum_{j=1}^{t} Z_j\| \le \|\widetilde{J}_0 - F'(x_0)\| + \sum_{j=1}^{t}\|Z_j\|.$$

On the other hand, by the update (21) of $\widetilde{J}_0$ as $\widetilde{J}_0 := \frac{1}{\hat{b}}\sum_{i\in\hat{\mathcal{B}}} F_i'(x_0)$, where $\hat{b} := \hat{b}_s$ and $\hat{\mathcal{B}} := \hat{\mathcal{B}}_s$, with a similar proof as of Theorem 4.2, we can show that if we choose $\hat{b} \geq \frac{6\sigma_D^2+2\sigma_D\hat{\epsilon}_0}{3\hat{\epsilon}_0^2}\log\left(\frac{p+q}{\delta}\right)$ then $\mathbf{Prob}\left(\|\widetilde{J}_0 - F'(x_0)\| \leq \hat{\epsilon}_0\right) \geq 1-\delta$. Then with probability at least $1-\delta$, we have

$$\|\widetilde{J}_t - F'(x_t)\| \leq \|\widetilde{J}_0 - F'(x_0)\| + \sum_{j=1}^t \|Z_j\| \leq \hat{\epsilon}_0 + \hat{\epsilon}\sum_{s=1}^t \hat{\sigma}_s = \hat{\epsilon}_0 + 2L_F\hat{\epsilon}\sum_{s=1}^t \|x_s - x_{s-1}\|.$$

This inequality implies

$$\|\widetilde{J}_t - F'(x_t)\|^2 \leq 2\hat{\epsilon}_0^2 + 8L_F^2\hat{\epsilon}^2 t\sum_{s=1}^t \|x_s - x_{s-1}\|^2. \tag{66}$$

Our next step is to estimate the $\|\widetilde{F}_t - F(x_t)\|$. We define $Y_i := F_i(x_t) - F_i(x_{t-1}) - [F(x_t) - F(x_{t-1})]$ and $U_j := F_j(x_t) - F_j(x_{t-1}) - F_j'(x_{t-1})(x_t - x_{t-1})$ for $j \in [n]$. In this case, $\{Y_i\}_{i\in\mathcal{B}_t}$ is mutually independent and $\mathbb{E}[Y_i] = 0$. We also have

$$\begin{aligned}
\|Y_i\| &= \left\|F_i(x_t) - F_i(x_{t-1}) - \tfrac{1}{n}\textstyle\sum_{j=1}^n [F_j(x_t) - F_j(x_{t-1})]\right\| \\
&= \Big\|F_i(x_t) - F_i(x_{t-1}) - F_i'(x_{t-1})(x_t - x_{t-1}) + \tfrac{1}{n}\textstyle\sum_{j=1}^n [F_i'(x_{t-1}) - F_j'(x_{t-1})](x_t - x_{t-1}) \\
&\quad - \tfrac{1}{n}\textstyle\sum_{j=1}^n [F_j(x_t) - F_j(x_{t-1}) - F_j'(x_{t-1})(x_t - x_{t-1})]\Big\| \\
&\leq \tfrac{1}{n}\Big\|\textstyle\sum_{j=1,j\neq i}^n [U_i - U_j]\Big\| + \|[F_i'(x_{t-1}) - F'(x_{t-1})](x_t - x_{t-1})\| \\
&\leq \tfrac{1}{n}\textstyle\sum_{j=1,j\neq i} \|U_j\| + \frac{n-1}{n}\|U_i\| + \|F_i'(x_{t-1}) - F'(x_{t-1})\|\,\|x_t - x_{t-1}\| \\
&\leq \tfrac{(n-1)L_F}{n}\|x_t - x_{t-1}\|^2 + \sigma_D\|x_t - x_{t-1}\|.
\end{aligned}$$

Here, we use the facts that $\|U_j\| = \|F_j(x_t) - F_j(x_{t-1}) - F_j'(x_{t-1})(x_t - x_{t-1})\| \leq \frac{L^2}{2}\|x_t - x_{t-1}\|^2$ for $j \in [n]$ and $\|F_i'(x_{t-1}) - F'(x_{t-1})\| \leq \sigma_D$ from Assumption 1.2 into the last inequality. Moreover, we have

$$\sigma_Y^2 := \max\left\{\Big\|\sum_{i\in\mathcal{B}_t}\mathbb{E}\left[Y_iY_i^\top\right]\Big\|, \Big\|\sum_{i\in\mathcal{B}_t}\mathbb{E}\left[Y_i^\top Y_i\right]\Big\|\right\} \leq \sum_{i\in\mathcal{B}_t}\mathbb{E}\left[\|Y_i\|^2\right] \leq b_t\sigma_t^2,$$

where $\sigma_t := L_F\|x_t - x_{t-1}\|^2 + \sigma_D\|x_t - x_{t-1}\|$.

Now, we consider $W_t := \frac{1}{b_t}\sum_{i\in\mathcal{B}_t} Y_i = \frac{1}{b_t}\sum_{i\in\mathcal{B}_t}[F_i(x_t) - F_i(x_{t-1}) - F(x_t) + F(x_{t-1})]$. For any $\epsilon > 0$, we can apply Lemma C.1 to obtain the following bound

$$\begin{aligned}
\mathbf{Prob}\left(\|W_t\| \leq \epsilon\sigma_t\right) &= \mathbf{Prob}\left(\|\textstyle\sum_{i\in\mathcal{B}_t} Y_i\| \leq \epsilon b_t\sigma_t\right) \geq 1 - (p+1)\exp\left(-\frac{3b_t^2\epsilon^2\sigma_t^2}{6b_t\sigma_t^2 + 2\sigma_t b_t\epsilon\sigma_t}\right) \\
&= 1 - (p+1)\exp\left(-\tfrac{3b_t\epsilon^2}{6+2\epsilon}\right).
\end{aligned}$$

Hence, if we choose $\delta \geq (p+1)\exp\left(-\frac{3b_t\epsilon^2}{6+2\epsilon}\right)$ and $\delta \leq 1$, then we obtain $\mathbf{Prob}\left(\|W_t\| \leq \epsilon\sigma_t\right) \geq 1-\delta$ for all $t \geq 0$. The condition in b_t leads to $b_t \geq \frac{6+2\epsilon}{3\epsilon^2}\cdot\log\left(\frac{p+1}{\delta}\right)$.

Note that since $\widetilde{F}_0 := \frac{1}{b}\sum_{i\in\mathcal{B}} F_i(x_0)$ is updated by (21), to guarantee

$$\mathbf{Prob}\left(\|\widetilde{F}_0 - F(x_0)\| \leq \epsilon_0\right) \geq 1-\delta,$$

we choose the mini-batch size $b \geq \frac{6\sigma_F^2+2\sigma_F\epsilon_0}{3\epsilon_0^2}\log\left(\frac{p+1}{\delta}\right)$.

By the update of $\widetilde{F}_t$ from (27), we have $\widetilde{F}_t - F(x_t) = [\widetilde{F}_{t-1} - F(x_{t-1})] + \frac{1}{b_t}[F_i(x_t) - F_i(x_{t-1}) - F(x_t) + F(x_{t-1})] = [\widetilde{F}_{t-1} - F(x_{t-1})] + W_t$. Hence, by induction, it implies that $\widetilde{F}_t - F(x_t) = [\widetilde{F}_0 - F(x_0)] + \sum_{s=1}^t W_s$, which leads to

$$\|\widetilde{F}_t - F(x_t)\| \leq \|\widetilde{F}_0 - F(x_0)\| + \sum_{s=1}^t \|W_s\| \leq \epsilon_0 + \epsilon\sum_{s=1}^t \left[L_F\|x_s - x_{s-1}\|^2 + \sigma_D\|x_s - x_{s-1}\|\right]. \tag{67}$$

Now, we analyze the inner loop of $t = 0$ to m. Using (15) with $x := x_t^{(s)}$ and $\widetilde{T}_M(x) = x_{t+1}^{(s)}$, we have

$$\phi(F(x_{t+1}^{(s)})) \le \phi(F(x_t^{(s)})) - \frac{C_g}{2}\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 + 2M_\phi\|F(x_t^{(s)}) - \widetilde{F}(x_t^{(s)})\| + \frac{M_\phi}{2\beta_d}\|F'(x_t^{(s)}) - \widetilde{J}(x_t^{(s)})\|^2, \quad (68)$$

where $C_g := 2M - M_\phi(L_F + \beta_d) > 0$ and $\beta_d > 0$ is given. Combining (68), (66), and (67), we have

$$\begin{aligned}\phi(F(x_{t+1}^{(s)})) &\le \phi(F(x_t^{(s)})) - \tfrac{C_g}{2}\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 + 2M_\phi\left[\epsilon_0 + L_F\epsilon\textstyle\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|^2\right]\\ &\quad + \tfrac{L_F}{2\beta_d}\left[2\hat{\epsilon}_0^2 + 8L_F^2\hat{\epsilon}^2 t\textstyle\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|^2\right] + 2M_\phi\sigma_D\epsilon\textstyle\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|.\end{aligned}$$

Summing up this inequality from $t = 0$ to $t = m$, we obtain

$$\begin{aligned}\phi(F(x_{m+1}^{(s)})) &\le \phi(F(x_0^{(s)})) - \tfrac{C_g}{2}\textstyle\sum_{t=0}^m\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 + 2M_\phi(m+1)\epsilon_0 + \frac{M_\phi(m+1)\hat{\epsilon}_0^2}{\beta_d}\\ &\quad + 2M_\phi L_F\epsilon\textstyle\sum_{t=0}^m\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|^2 + \frac{4L_F^3\hat{\epsilon}^2}{\beta_d}\sum_{t=0}^m t\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|^2\\ &\quad + 2M_\phi\sigma_D\epsilon\textstyle\sum_{t=0}^m\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|.\end{aligned}$$

Since $\widetilde{x}^{s-1} = x_0^{(s)}$ and $\widetilde{x}^s = x_{m+1}^{(s)}$, the last inequality becomes

$$\phi(F(\widetilde{x}^s)) \le \phi(F(\widetilde{x}^{s-1})) - \frac{C_g}{4}\sum_{t=0}^m\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 + 2M_\phi(m+1)\epsilon_0 + \frac{M_\phi(m+1)m\epsilon^2}{2\gamma} + \frac{M_\phi(m+1)\hat{\epsilon}_0^2}{\beta_d} + \mathcal{T}_m^s, \quad (69)$$

where $\mathcal{T}_m^s$ is defined as

$$\begin{aligned}\mathcal{T}_m^s &:= 2M_\phi L_F\epsilon\textstyle\sum_{t=0}^m\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|^2 + 2M_\phi\sigma_D\epsilon\sum_{t=0}^m\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|\\ &\quad + \tfrac{4L_F^3\hat{\epsilon}^2}{\beta_d}\textstyle\sum_{t=0}^m t\sum_{j=1}^t\|x_j^{(s)} - x_{j-1}^{(s)}\|^2 - \frac{C_g}{4}\sum_{t=0}^m\|x_{t+1}^{(s)} - x_t^{(s)}\|^2.\end{aligned}$$

Let $u_{t-1} := \|x_t^{(s)} - x_{t-1}^{(s)}\|$. Then, we can rewrite $\mathcal{T}_m^s$ as

$$\begin{aligned}\mathcal{T}_m^s &= 2M_\phi L_F\epsilon\left[u_0^2 + (u_0^2 + u_1^2) + \cdots + (u_0^2 + u_1^2 + u_{m-1}^2)\right]\\ &\quad + 2M_\phi\sigma_D\epsilon\left[u_0 + (u_0 + u_1) + \cdots + (u_0 + u_1 + u_{m-1})\right]\\ &\quad + \tfrac{4L_F^3\hat{\epsilon}^2}{\beta_d}\left[u_0^2 + 2(u_0^2 + u_1^2) + \cdots + m(u_0^2 + u_1^2 + \cdots + u_{m-1}^2)\right]\\ &\quad - \tfrac{C_g}{4}\left[u_0^2 + u_1^2 + \cdots + u_m^2\right]\\ &= \left[2M_\phi L_F\epsilon m + \tfrac{4L_F^3\hat{\epsilon}^2}{\beta_d}m(m+1) - \tfrac{C_g}{4}\right]u_0^2 + \left[2M_\phi L_F\epsilon(m-1) + \tfrac{4L_F^3\hat{\epsilon}^2}{\beta_d}m(m-1) - \tfrac{C_g}{4}\right]u_1^2 + \cdots\\ &\quad + \left[2M_\phi L_F\epsilon + \tfrac{4L_F^3\hat{\epsilon}^2}{\beta_d} - \tfrac{C_g}{4}\right]u_{m-1}^2 - \tfrac{C_g}{4}u_m^2 + 2M_\phi\sigma_D\epsilon\left[mu_0 + (m-1)u_1 + \cdots + u_{m-1}\right]\\ &\le \left[2M_\phi L_F\epsilon m + \tfrac{4L_F^3\hat{\epsilon}^2 m(m+1)}{\beta_d} - \tfrac{C_g}{4}\right]u_0^2 + \left[2M_\phi L_F\epsilon(m-1) + \tfrac{4L_F^3\hat{\epsilon}^2 m(m-1)}{\beta_d} - \tfrac{C_g}{4}\right]u_1^2 + \cdots\\ &\quad + \left[2M_\phi L_F\epsilon + \tfrac{4L_F^3\hat{\epsilon}^2}{\beta_d} - \tfrac{C_g}{4}\right]u_{m-1}^2 - \tfrac{C_g}{4}u_m^2\\ &\quad + \tfrac{M_\phi\sigma_D m\epsilon}{\gamma}\left[u_0^2 + u_1^2 + \cdots + u_{m-1}^2\right] + M_\phi\sigma_D\epsilon\gamma m^2\\ &\le \left[2M_\phi L_F\epsilon m + M_\phi\sigma_D\sqrt{\epsilon}m + \tfrac{4L_F^3\hat{\epsilon}^2 m(m+1)}{\beta_d} - \tfrac{C_g}{4}\right]u_0^2\\ &\quad + \left[2M_\phi L_F\epsilon(m-1) + M_\phi\sigma_D\sqrt{\epsilon}m + \tfrac{4L_F^3\hat{\epsilon}^2 m(m-1)}{\beta_d} - \tfrac{C_g}{4}\right]u_1^2 + \cdots\\ &\quad + \left[2M_\phi L_F\epsilon + \tfrac{M_\phi\sigma_D\epsilon m}{\gamma} + \tfrac{4L_F^3\hat{\epsilon}^2}{\beta_d} - \tfrac{C_g}{4}\right]u_{m-1}^2 - \tfrac{C_g}{4}u_m^2 + M_\phi\sigma_D\epsilon\gamma m^2.\end{aligned}$$

If we impose the following condition

$$M_\phi\left(2L_F + \frac{\sigma_D}{\gamma}\right)\epsilon m + \frac{2L_F^3\hat{\epsilon}^2 m(m+1)}{\beta_d} \le \frac{C_g}{4}, \quad (70)$$

then $\mathcal{T}_m^s \le M_\phi\sigma_D\epsilon\gamma m^2$.

Under this condition, (69) reduces to

$$\phi(F(\widetilde{x}^s)) \le \phi(F(\widetilde{x}^{s-1})) - \frac{C_g}{4}\sum_{t=0}^{m}\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 + 2M_\phi(m+1)\epsilon_0 + \frac{M_\phi(m+1)\hat{\epsilon}_0^2}{\beta_d} + M_\phi\sigma_D\epsilon\gamma m^2.$$

Summing up this inequality from $s = 1$ to $S = s$ and rearranging the result, we obtain

$$\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 \le \frac{4}{C_g(m+1)S}\left[\phi(F(\widetilde{x}^0)) - \phi(F(\widetilde{x}^S))\right] + \frac{4M_\phi}{C_g}\left(2\epsilon_0 + \sigma_D m\gamma\epsilon + \frac{\hat{\epsilon}_0^2}{\beta_d}\right).$$

Using $\phi(F(\widetilde{x}^S) \ge \Psi^\star$ and $\Psi(x) = \phi(F(x))$, we obtain from the last inequality that

$$\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\|x_{t+1}^{(s)} - x_t^{(s)}\|^2 \le \frac{4}{C_g(m+1)S}\left[\Psi(\widetilde{x}^0) - \Psi^\star\right] + \frac{4M_\phi}{C_g}\left(2\epsilon_0 + \sigma_D m\gamma\epsilon + \frac{\hat{\epsilon}_0^2}{\beta_d}\right).$$

Clearly, if we choose $\epsilon_0 := C_0\varepsilon^2$, $\epsilon := \frac{C_1\varepsilon^2}{\gamma m}$, $\hat{\epsilon}_0 := \sqrt{\hat{C}_0}\varepsilon$, and $\hat{\epsilon}^2 := \frac{\hat{C}_1}{m(m+1)}$ for some positive constant C_0, C_1, $\hat{C}_0$, and $\hat{C}_1$, then we obtain from the last estimate that

$$\frac{1}{S(m+1)}\sum_{s=1}^{S}\sum_{t=0}^{m}\|\widetilde{G}_M(x_t^{(s)})\|^2 \le \frac{4M^2}{C_g(m+1)S}\left[\Psi(\widetilde{x}^0) - \Psi^\star\right] + \frac{M^2M_\phi}{C_g}\left(2C_0 + \sigma_D C_1 + \frac{\hat{C}_0}{\beta_d}\right)\varepsilon^2, \tag{71}$$

where we use the fact that $\widetilde{G}_M(x_t^{(s)}) = M(x_{t+1}^{(s)} - x_t^{(s)})$. Now, assume that the condition (70) is tight. Using the choice of accuracies, we obtain

$$M_\phi\left(2L_F + \frac{\sigma_D}{\gamma}\right)\frac{C_1\varepsilon^2}{\gamma} + \frac{2L_F^3\hat{C}_1}{\beta_d} = \frac{C_g}{4}.$$

If we choose $\gamma := \varepsilon$, then this condition becomes $2M_\phi\left(L_F\varepsilon + \sigma_D\right)C_1 + \frac{2L_F^3\hat{C}_1}{\beta_d} = \frac{C_g}{4}$ and $\epsilon := \frac{C_1\varepsilon}{m}$.

Now, with the choice of ϵ_0, ϵ, $\hat{\epsilon}_0$, and $\hat{\epsilon}$ as above, we can set the mini-batch sizes as follows:

$$\left\{\begin{array}{lcll} b_s & := & \left\lfloor\frac{6\sigma_F^2 + 2\sigma_F C_0\varepsilon^2}{3C_0^2\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right)\right\rfloor & = \mathcal{O}\left(\frac{\sigma_F^2}{\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right)\right), \\ b_t^{(s)} & := & \left\lfloor\frac{m[6m+2C_1\varepsilon]}{3C_1^2\varepsilon^2}\cdot\log\left(\frac{p+1}{\delta}\right)\right\rfloor & = \mathcal{O}\left(\frac{m^2}{\varepsilon^2}\cdot\log\left(\frac{p+1}{\delta}\right)\right), \\ \hat{b}_s & := & \left\lfloor\frac{\left[6\sigma_D^2 + 2\sigma_D\sqrt{\hat{C}_1}\varepsilon\right]}{3\hat{C}_1\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right\rfloor & = \mathcal{O}\left(\frac{\sigma_D^2}{\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right)\right), \\ \hat{b}_t^{(s)} & := & \left\lfloor\frac{\sqrt{m(m+1)}\left[6\sqrt{m(m+1)}+2\sqrt{\hat{C}_1}\right]}{3\hat{C}_1}\cdot\log\left(\frac{p+q}{\delta}\right)\right\rfloor & = \mathcal{O}\left(m^2\cdot\log\left(\frac{p+1}{\delta}\right)\right). \end{array}\right. \tag{72}$$

Since $S(m+1) = \frac{8M^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{C_g\varepsilon^2}$, if we choose $m := \frac{C}{\varepsilon}$, then $S = \frac{8M^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{CC_g\varepsilon}$. The total complexity is

$$\begin{aligned} \mathcal{T}_f &:= \textstyle\sum_{s=1}^{S} b_s + \sum_{s=1}^{S}\sum_{t=0}^{m} b_t^{(s)} = \frac{[6\sigma_F^2 + 2\sigma_F C_0\varepsilon^2]S}{3C_0^2\varepsilon^4}\cdot\log\left(\frac{p+1}{\delta}\right) + \frac{S(m+1)m[6m+2C_1\varepsilon]}{3C_1^2\varepsilon^2}\cdot\log\left(\frac{p+1}{\delta}\right) \\ &= \frac{8M^2[6\sigma_F^2 + 2\sigma_F C_0\varepsilon^2]\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{3CC_gC_0^2\varepsilon^5}\cdot\log\left(\frac{p+1}{\delta}\right) + \frac{8M^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]\left[6C^2 + 2CC_1\varepsilon^2\right]}{3C_gC_1^2\varepsilon^6}\cdot\log\left(\frac{p+1}{\delta}\right) \\ &= \mathcal{O}\left(\frac{\sigma_F^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\varepsilon^5}\cdot\log\left(\frac{p+1}{\delta}\right)\right) + \mathcal{O}\left(\frac{\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\varepsilon^6}\cdot\log\left(\frac{p+1}{\delta}\right)\right), \\ \mathcal{T}_d &:= \textstyle\sum_{s=1}^{S}\hat{b}_s + \sum_{s=1}^{S}\sum_{t=0}^{m}\hat{b}_t^{(s)} \\ &= \frac{\left[6\sigma_D^2 + 2\sigma_D\sqrt{\hat{C}_1}\varepsilon\right]S}{3\hat{C}_1\varepsilon^2}\cdot\log\left(\frac{p+q}{\delta}\right) + \frac{S(m+1)\sqrt{m(m+1)}\left[6\sqrt{m(m+1)}+2\sqrt{\hat{C}_1}\right]}{3\hat{C}_1}\cdot\log\left(\frac{p+q}{\delta}\right) \\ &= \mathcal{O}\left(\frac{\sigma_D^2\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\varepsilon^3}\cdot\log\left(\frac{p+q}{\delta}\right)\right) + \mathcal{O}\left(\frac{\left[\Psi(\widetilde{x}^0) - \Psi^\star\right]}{\varepsilon^4}\cdot\log\left(\frac{p+q}{\delta}\right)\right). \end{aligned}$$

This proves our theorem. □