Minibatch and Momentum Model-based Methods for Stochastic Non-smooth Non-convex Optimization

Anonymous Author(s)
Affiliation
Address
email

Abstract

Stochastic model-based methods have received increasing attention lately due to their appealing robustness to the stepsize selection and provable efficiency guarantee for non-smooth non-convex optimization. To further improve the performance of stochastic model-based methods, we make two important extensions. First, we propose a new minibatch algorithm which takes a set of samples to approximate the model function in each iteration. For the first time, we show that stochastic algorithms achieve linear speedup over the batch size even for non-smooth and non-convex problems. To this end, we develop a novel sensitivity analysis of the proximal mapping involved in each algorithm iteration. Our analysis can be of independent interests in more general settings. Second, motivated by the success of momentum techniques for convex optimization, we propose a new stochastic extrapolated model-based method to possibly improve the convergence in the non-smooth and non-convex setting. We obtain complexity guarantees for a fairly flexible range of extrapolation term. In addition, we conduct experiments to show the empirical advantage of our proposed methods.

1 Introduction

In this paper, we are interested in the following stochastic optimization problem

$$\min_{x \in \mathcal{X}} \quad f(x) = \mathbb{E}_{\xi \sim \Xi}\big[f(x, \xi)\big] \tag{1}$$

where $f(\cdot, \xi)$ stands for the loss function, sample ξ follows certain distribution Ξ, and $\mathcal{X}$ is a closed convex set. We assume that $f(\cdot, \xi)$ is weakly convex, namely, the sum of $f(x, \xi)$ and a quadratic function $\frac{\lambda}{2}\|x\|^2$ is convex ($\lambda > 0$). This type of non-smooth non-convex functions has a wide range of applications in signal processing and machine learning, such as phase retrieval, robust PCA and low rank decomposition [7]. To solve problem (1), we consider stochastic model-based methods (`SMOD`, [12, 8, 1]), which comprise a large class of stochastic algorithms (including stochastic (sub)gradient descent, proximal point, among others). In spite of the non-smoothness and non-convexity, `SMOD` exhibits promising convergence property [12, 8]: both asymptotic convergence and rate of convergence to certain stationarity measures have been established for the whole `SMOD` family. In addition, extensive empirical study [8, 13] has shown that `SMOD` often outperforms `SGD` due to its remarkable robustness to hyper-parameter tuning.

Despite much recent progress, it still remains to see whether `SMOD` is competitive against modern `SGD` in practice. We start by addressing the crucial limitations of the prior study and highlighting some remaining questions. First, despite the appealing robustness and stable convergence, the `SMOD` family is sequential in nature. It is unclear how minibatching, which is immensely used in training learning models, can improve the performance of `SMOD` when the problem is non-smooth. Particularly,

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

the current best complexity bound $\mathcal{O}(\frac{L^2}{\varepsilon^4})$ from [8], which is regardless of batchsize, is somewhat unsatisfactory. Were this bound tight, a sequential algorithm (using one sample per iteration) would be optimal: it offers the highest processing speed per iteration as well as the best iteration complexity. Therefore, it is crucial to know whether minibatching can improve the complexity bound of the SMOD family or the current one is tight. Second, in modern applications, momentum technique has been playing a vital role in large-scale non-convex optimization (see [28, 26]). In spite of its effectiveness, to the best of our knowledge, momentum technique has been provably efficient only in **1)** unconstrained smooth optimization [21, 9, 17]) and **2)** non-smooth optimization with a simple constraint [23], which constitute only a portion of the interesting applications. From the practical aspect, it is peculiarly desirable to know whether momentum technique is applicable beyond in SGD and whether it can benefit the SMOD algorithm family in a wider problem class.

Contributions. Motivated by the challenges for making model-based optimization more efficient, we make two extensions. First, we extend the stochastic model-based methods to minibatch setting and develop sharper rates of convergence to stationarity. Leveraging the tool of algorithm stability ([6, 25, 18]), we provide a nearly-complete recipe on when minibatching would be helpful even in presence of non-smoothness. Interestingly, our results show that the stochastic proximal point method can be linearly accelerated by minibatching and that stochastic proximal-linear method gets the same promising speedup for composition functions (illustrated in Section 3). To the best of our knowledge, this is the first analysis showing that these minibatch stochastic algorithms achieve linear speedup over the batchsize even for *non-smooth* and *non-convex* optimization problems.

Second, we present new extrapolated model-based methods by incorporating a Polyak-type momentum term. We develop a unified Lyapunov analysis to show that a worst-case complexity of $\mathcal{O}(1/\varepsilon^4)$ holds for all momentum SMOD algorithms. To the best of our knowledge, these are the first complexity results of momentum stochastic proximal-linear and proximal point algorithms for non-smooth non-convex optimization. Our theory also guarantees a similar complexity bound of momentum SGD and its proximal extension, thereby being more general than the recent work [23] which only proves the convergence of momentum projected SGD. A possible advantage of our work over [23] lies in composite optimization, where the non-smooth term is often involved via its proximal operator rather than subgradient. One such example is the Lasso problem where, to enhance solution sparsity, it is often favorable to invoke the proximal operator of ℓ_1 function (soft-thresholding). A summary of the complexity results is provided in Table 1.

Table 1: Complexity of SMOD methods for reaching $\mathbb{E}\,\|\nabla_{1/\rho} f\| \leq \varepsilon$ (M: minibatch; E: Extrapolation)

Algorithms	Problem	Current Best	Ours
M + SGD	f: nonsmooth	$\mathcal{O}(1/\varepsilon^4)$[8]	$\mathcal{O}(1/\varepsilon^4)$
M + Prox. SGD	$f = \ell + \omega$; ℓ:smooth	$\mathcal{O}(1/(m\varepsilon^4) + 1/\varepsilon^2)$[8]	$\mathcal{O}(1/(m\varepsilon^4) + 1/\varepsilon^2)$
M + SPL/SPP	f: nonsmooth	$\mathcal{O}(1/\varepsilon^4)$[8]	$\mathcal{O}(1/(m\varepsilon^4) + 1/\varepsilon^2)$
E + SGD	f: nonsmooth	$\mathcal{O}(1/\varepsilon^4)$[23]	$\mathcal{O}(1/\varepsilon^4)$
E + Prox. SGD	$f = \ell + \omega$; ℓ:smooth	—	$\mathcal{O}(1/\varepsilon^4)$
E + SPL/SPP	f: nonsmooth	—	$\mathcal{O}(1/\varepsilon^4)$
M + E + SGD	f: nonsmooth	$\mathcal{O}(1/\varepsilon^4)$[23]	$\mathcal{O}(1/\varepsilon^4)$
M + E + Prox. SGD	$f = \ell + \omega$; ℓ:smooth	—	$\mathcal{O}(1/(m\varepsilon^4) + 1/\varepsilon^2)$
M + E + SPL/SPP	f: nonsmooth	—	$\mathcal{O}(1/(m\varepsilon^4) + 1/\varepsilon^2)$

Other related work. For smooth and composite optimization, it is well known that batchsize can linearly reduce the iteration count of SGD. See [10, 15, 27]. Asi et al. [2] investigates minibatch stochastic model-based methods in the convex non-smooth setting but requires a strong assumption of restricted strong convexity. Hence their proof technique does not extend readily to the non-convex setting. Practically, the robustness and fast convergence of model-based optimization have been shown on various non-smooth non-convex statistical learning problems [7, 13, 1, 4, 14, 5]. Drusvyatskiy and Paquette [11] give a complete recipe of complexity analysis on accelerated proximal-linear methods for deterministic optimization. Momentum and accelerated methods for convex stochastic optimization can be referred from [22, 24]. The study [28, 21, 9] develop the convergence rate of stochastic momentum method for smooth nonconvex optimization.

2 Background

Throughout the whole paper, we use $\|\cdot\|$ to denote the Euclidean norm and $\langle\cdot,\cdot\rangle$ to denote the Euclidean inner product. We assume that $f(x)$ is bounded below. i.e., $\min_x f(x) > -\infty$. We say that the function $f(x)$ is λ-*weakly convex* if $f(x) + \frac{\lambda}{2}\|x\|^2$ is a convex function. . The subdifferential $\partial f(x)$ of function $f(x)$ is the set of vectors $v \in \mathbb{R}^d$ that satisfy: $f(y) \geq f(x) + \langle v, y-x\rangle + o(\|x-y\|)$, as $y \to x$. Any such vector in $\partial f(x)$ is called a subgradient and is denoted by $f'(x) \in \partial f(x)$ for simplicity. We say that a point x is stationary if $0 \in \partial f(x) + N_{\mathcal{X}}(x)$, where the normal cone $N_{\mathcal{X}}(x)$ is defined as $N_{\mathcal{X}}(x) \triangleq \{d : \langle d, y-x\rangle \leq 0, \forall y \in \mathcal{X}\}$. For a set S, define the set distance to 0 by: $\|\mathcal{S}\|_- \triangleq \inf\{\|x-0\|, x \in \mathcal{S}\}$. It is natural to use the quantity $\|\partial f(x) + N_{\mathcal{X}}(x)\|_-$ to measure the stationarity of point x.

Moreau-envelope. According to [3], the μ-Moreau-envelope of f is defined by $f_\mu(x) \triangleq \min_{y\in\mathcal{X}}\left\{f(y) + \frac{1}{2\mu}\|x-y\|^2\right\}$ and the proximal mapping associated with $f(\cdot)$ is defined by $\mathrm{prox}_{\mu f}(x) \triangleq \mathrm{argmin}_{y\in\mathcal{X}}\left\{f(y) + \frac{1}{2\mu}\|x-y\|^2\right\}$. Assume that $f(x)$ is λ-weakly convex, then for $\mu < \lambda^{-1}$, the Moreau envelope $f_\mu(\cdot)$ is differentiable and its gradient is $\nabla f_\mu(x) = \mu^{-1}(x - \mathrm{prox}_{\mu f}(x))$.

The SMOD family iteratively computes the proximal map associated with a model function $f_{x^k}(\cdot, \xi_k)$:

$$x^{k+1} = \operatorname*{argmin}_{x\in\mathcal{X}}\left\{f_{x^k}(x,\xi_k) + \frac{\gamma_k}{2}\|x - x^k\|^2\right\}, \tag{2}$$

where $\{\xi_k\}$ are i.i.d. samples. Typical algorithms and the accompanied models are described below.

Stochastic (Proximal) Gradient Descent: consider the composite function $f(x,\xi) = \ell(x,\xi) + \omega(x)$ where $\ell(x,\xi)$ is a data-driven loss term and $\omega(x)$ is a convex regularizer such as ℓ_1-penalty. SGD applies the model function:

$$f_y(x,\xi) = \ell(y,\xi) + \left\langle \ell'(y,\xi), x-y\right\rangle + \omega(x). \tag{3}$$

Stochastic Prox-linear (SPL): consider the composition function $f(x,\xi) = h(C(x,\xi))$ where $h(\cdot,\xi)$ is convex continuous and $C(x,\xi)$ is a continuously differentiable map. We perform partial linearization to obtain the model

$$f_y(x,\xi) = h\big(C(y,\xi) + \langle \nabla C(y,\xi), x-y\rangle\big). \tag{4}$$

Stochastic Proximal Point (SPP): compute (2) with full stochastic function:

$$f_y(x,\xi) = f(x,\xi). \tag{5}$$

Throughout the paper, we assume that $f(x,\xi)$ is continuous and μ-weakly convex, and that the model function $f_x(\cdot,\cdot)$ satisfies the following assumptions [8].

- **A1:** For any $\xi \sim \Xi$, the model function $f_x(y,\xi)$ is λ-weakly convex in y ($\lambda \geq 0$).
- **A2:** Tightness condition: $f_x(x,\xi) = f(x,\xi),\ \forall x \in \mathcal{X},\ \xi \sim \Xi$.
- **A3:** One-sided quadratic approximation: $f_x(y,\xi) - f(y,\xi) \leq \frac{\tau}{2}\|x-y\|^2,\ \forall x,y \in \mathcal{X}, \xi \sim \Xi$.
- **A4:** Lipschitz continuity: There exists $L > 0$ that $f_x(z,\xi) - f_x(y,\xi) \leq L\,\|z-y\|$, for any $x,y,z \in \mathcal{X},\ \xi \sim \Xi$.

Remark 1. Assumption A2 is quite standard and will be used only in the convergence proof. Combining A1 and A3, we immediately have that $f(x,\xi)$ is $(\lambda+\tau)$-weakly convex. Thus, it suffices to assume that $\mu < \tau + \lambda$. Assumptions A2-A4 can be slightly relaxed by replacing the uniform bound with a bound on expectation over ξ, leading to only a minor adjustment to the analysis.

Denote $\hat{x} \triangleq \mathrm{prox}_{f/\rho}(x) = \mathrm{argmin}_y\left\{f(y) + \frac{\rho}{2}\|y-x\|^2\right\}$ for some $\rho > \mu$. Davis and Drusvyatskiy [8] revealed a striking feature of Moreau envelope to characterize stationarity:

$$\|\hat{x} - x\| = \rho^{-1}\|\nabla f_{1/\rho}(x)\|, \text{ and } \|\partial f(\hat{x}) + N_{\mathcal{X}}(\hat{x})\|_- \leq \|\nabla f_{1/\rho}(x)\|.$$

Namely, a point x with small gradient norm $\|\nabla f_{1/\rho}(x)\|$ stays in the proximity of a nearly-stationary point $\hat{x}$. With this observation, they show the first complexity result of SMOD for non-smooth non-convex optimization: $\min_{1\leq k\leq K}\mathbb{E}\,\|\nabla f_{1/\rho}(x^k)\|^2 \leq \mathcal{O}(\frac{L}{\sqrt{K}})$. Note that this rate is regardless of the size of minibatches since it does not explicitly use any information of samples other than the Lispchitzness of the model function. Due to this limitation, it remains unclear whether minibatching can further improve the convergence rate of SMOD.

Algorithm 1 Stochastic Model-based Method with Minibatches (SMOD)

Input: x^1

for $k = 1$ **to** K **do**

Sample a minibatch $B_k = \{\xi_{k,1}, \ldots, \xi_{k,m_k}\}$ and update x^{k+1} by solving

$$\min_{x \in \mathcal{X}} \frac{1}{m_k} \sum_{i=1}^{m_k} f_{x^k}\left(x, \xi_{k,i}\right) + \frac{\gamma_k}{2} \left\|x - x^k\right\|^2 \tag{6}$$

end for

3 SMOD with Minibatches

In this section we present a minibatch SMOD method which takes a small batch of i.i.d. samples to estimate the model function. The overall procedure is detailed in Algorithm 1. Within each iteration, Algorithm 1 forms a stochastic model function $f_{x^k}(\cdot, B_k) = \frac{1}{m_k}\sum_{i=1}^{m_k} f_{x^k}(x, \xi_{k,i})$ parameterized at x^k by sampling over m_k i.i.d. samples $B_k = \xi_{k,1}, \ldots, \xi_{k,m_k}$. Then it performs proximal update to get the next iterate x^{k+1}. We will illustrate the main convergence results of Algorithm 1 but leave all the proof details in Appendix sections. But first, let us present a few additional assumptions.

A5: Two-sided quadratic bound: for any $x, y \in \mathcal{X}$, $\xi \sim \Xi$, $\left|f_x(y,\xi) - f(y,\xi)\right| \leq \frac{\tau}{2}\|x - y\|^2$.

A6: (Optional) Lipschitzness w.r.t. data: there exists $M > 0$ such that for any $x, y \in \mathcal{X}$, $\xi_1, \xi_2 \sim \Xi$, we have $|f_x(y,\xi_1) - f_x(y,\xi_2)| \leq M\,\|\xi_1 - \xi_2\|$.

A7: (Optional) Bounded second moment: there exists $D > 0$ such that $\mathbb{E}_{\xi\sim\Xi}\left[\|\xi\|^2\right] \leq \frac{1}{4}D^2$.

Remark 2. A5 is a vital piece for our improved convergence results. While it appears to be stronger than A3, A5 is indeed satisfied by the SMOD family in most contexts. **1)** For SPP, A5 is trivially satisfied by taking $f_x(y,\xi) = f(y,\xi)$. **2)** For SPL, we minimize a compound function $f(x,\xi) = h(C_\xi(x))$ where $h(\cdot)$ is a c_1-Lipschitz convex function and $C_\xi(\cdot)$ is a c_2-Lipschitz smooth map. In view of (4), A5 is verified with $|f_x(y,\xi) - f(y,\xi)| \leq c_1\left\|C_\xi(y) - C_\xi(x) - \nabla C_\xi(x)^{\mathrm{T}}(y-x)\right\| \leq \frac{c_1 c_2}{2}\|x-y\|^2$. **3)** For SGD, A5 is satisfied if $\ell(x,\xi)$ is L-Lipschitz smooth for some $L > 0$, as $|f_x(y,\xi) - f(y,\xi)| \leq |\ell(y,\xi) - \ell(x,\xi) - \nabla\ell(x,\xi)^{\mathrm{T}}(y-x)| \leq \frac{L}{2}\|x-y\|^2$. **4)** We note that A5 is not satisfied by SGD when the loss $\ell(\cdot,\xi)$ is also non-smooth. Unfortunately, there seems to be little hope to accelerate SGD in such case since the convergence rate of SGD already matches the rate of deterministic subgradient method.

Remark 3. A6 and A7 are benign assumptions on the data to establish the sharpest convergence rate possible. However, as will be shown, removing A6 and A7 (simply letting $M, D \to \infty$) does not affect our main convergence result.

We present an improved complexity analysis of minibatch algorithms by leveraging the framework of algorithm stability [6, 25]. In stark contrast to its standard use for estimating the algorithms' generalization ability, we perform stability analysis to determine how the variation of a minibatch affects the *estimation of the model function* in each algorithm iteration. Interestingly, stability analysis provides a highly general bound under mild conditions, even obviating the need of smoothness assumption in most analyses on minibatch SGD (e.g. [16]).

Notations. Let $B = \{\xi_1, \xi_2, \ldots, \xi_m\}$ be a batch of i.i.d. samples and $B_i = B \setminus \{\xi_i\} \cup \{\xi_i'\}$ by replacing ξ_i with an i.i.d. copy ξ_i'. We denote $B'\{\xi_1', \xi_2', \ldots, \xi_n'\}$. Let $h(\cdot,\xi)$ be a stochastic model function, and denote $h(y,B) = \frac{1}{m}\sum_{i=1}^m h(y,\xi_i)$. The stochastic proximal mapping associated with $h(\cdot,B)$ is defined by $\mathrm{prox}_{\rho h}(x,B) \triangleq \mathrm{argmin}_{y\in\mathcal{X}}\left\{h(y,B) + \frac{1}{2\rho}\|y-x\|^2\right\}$ for some $\rho > 0$. We denote $x_B^+ \triangleq \mathrm{prox}_{\rho h}(x,B)$ for brevity. We say that the stochastic proximal mapping $\mathrm{prox}_{\rho h}$ is ε-*stable* if, for any $x \in \mathcal{X}$, we have

$$\left|\mathbb{E}_{B,B',i}\left[h(x_{B_i}^+, \xi_i') - h(x_B^+, \xi_i')\right]\right| \leq \varepsilon, \tag{7}$$

where i is an index chosen from $\{1, 2, \ldots, m\}$ uniformly at random.

The next lemma exploits the stability of proximal mapping associated with the model function.

Lemma 3.1. *Let $f_z(\cdot, B)$ be a stochastic model function under the assumptions A1, A4, A6 and A7. Let $\gamma > \lambda$. For vectors z and y, the proximal mapping $\operatorname{prox}_{f_z/\gamma}(y, B) = \operatorname{argmin}_{x\in\mathcal{X}} \left\{ f_z(x, B) + \frac{\gamma}{2}\|x - y\|^2 \right\}$ is ε-stable with $\varepsilon = \min\left\{ \frac{2L^2}{m(\gamma-\lambda)}, L\sqrt{\frac{2MD}{m(\gamma-\lambda)}} \right\}$.*

Applying Lemma 3.1, we obtain the error bound for approximating the full model function. We summarize this result in the following theorem.

Theorem 3.2. *Under all the assumptions of Lemma 3.1, we have*

$$\left|\mathbb{E}_{B_k}\left[f_{x^k}(x^{k+1}, B_k) - \mathbb{E}_\xi f_{x^k}(x^{k+1}, \xi)|\sigma_k\right]\right| \le \varepsilon_k,\ \varepsilon_k = \min\left\{ \frac{2L^2}{m_k(\gamma_k-\lambda)}, L\sqrt{\frac{2MD}{m_k(\gamma_k-\lambda)}} \right\}. \quad (8)$$

where σ_k is the σ-algebra generating $\{B_i\}_{1\le i\le k-1}$.

It can be observed that the stochastic error is jointly bounded by the batchsize and stochastic noise. However, the advantage of stability analysis lies in the fact that when M and D get large ($M, D \to \infty$), the error is still bounded by $\mathcal{O}(1/(m_k\gamma_k))$, which is independent of the level of stochastic noise. This observation is the key for sharp analysis of minibatch stochastic algorithms. With all the tools at hands, we obtain the key descent property in the following theorem.

Theorem 3.3. *Under the assumptions of Lemma 3.1 and A5, assume that $\rho > \lambda + \tau$ and that $\gamma_k \ge \rho + \tau$, then we have*

$$\frac{(\rho - \lambda - \tau)}{\rho(\gamma_k + \rho - 2\lambda - \tau)}\|\nabla f_{1/\rho}(x^k)\|^2 \le f_{1/\rho}(x^k) - \mathbb{E}_k\left[f_{1/\rho}(x^{k+1})\right] + \frac{\rho\varepsilon_k}{\gamma_k + \rho - 2\lambda - \tau}, \quad (9)$$

where $\mathbb{E}_k[\cdot]$ abbreviates $\mathbb{E}_{B_k}\left[\,\cdot\,|\sigma_k\right]$ and ε_k is given by (8).

Next, we specify the rate of convergence to stationarity using a constant stepsize policy.

Theorem 3.4. *Under the assumptions of Theorem 3.3, let $\Delta = f_{1/\rho}(x^1) - \min_x f(x)$, $m_k = m$, and $\gamma_k = \gamma = \max\{\rho+\tau, \lambda+\eta\}$ where $\eta = \begin{cases} L\left(\frac{2\rho K}{m\Delta}\right)^{1/2} & \text{if } C_1 \le C_2 \\ \left(\frac{\rho L K}{\Delta}\right)^{2/3}\left(\frac{MD}{m}\right)^{1/3} & \text{o.w.} \end{cases}$, with $C_1 = 2L\sqrt{\frac{2\rho\Delta}{mK}}$ and $C_2 = 3\sqrt[3]{\frac{\rho^2L^2MD\Delta}{2mK}}$. Let k^* be an index chosen in $\{1, 2, \ldots, K\}$ uniformly at random, then*

$$\mathbb{E}\left[\|\nabla f_{1/\rho}(x^{k^*})\|^2\right] \le \frac{\rho}{\rho - \lambda - \tau}\left[\frac{(2\rho - \lambda)\Delta}{K} + \min\left\{C_1, C_2\right\}\right], \quad (10)$$

Remark 4. In view of Theorem 3.4, to obtain an iterate whose expected gradient norm is smaller than ε, the total iteration count is $\mathcal{T}_\varepsilon = \max\left\{\mathcal{O}(\frac{\Delta}{\varepsilon^2}), \min\{\mathcal{O}(\frac{L^2\Delta}{m\varepsilon^4}), \mathcal{O}(\frac{L^2MD\Delta}{m\varepsilon^6})\}\right\}$. For deterministic problems (i.e. $M = 0$), our complexity result can be further improved to $\mathcal{O}\left(\frac{\Delta}{\varepsilon^2}\right)$, matching the best complexity result for weakly convex optimization (see [11]). Taking $M, D \to \infty$, we have the total complexity: $\max\{\mathcal{O}(\frac{\Delta}{\varepsilon^2}), \mathcal{O}(\frac{L^2\Delta}{m\varepsilon^4})\}$ while [8] gives $\max\{\mathcal{O}(\frac{\Delta}{\varepsilon^2}), \mathcal{O}(\frac{L^2\Delta}{\varepsilon^4})\}$. Commonly, the second term in $\max(,)$ dominates, and our bound $\mathcal{O}(\frac{L^2\Delta}{m\varepsilon^4})$ is better than the result of [8] by a factor of m.

Remark 5. While the parameter setting appears to be complicated, it aims to provide the sharpest theoretical rate possible. In practice, we can simply take $M, D \to \infty$, and then we obtain nearly the same optimal bound, but only have one more tuning parameter (batchsize m) than [8]. Empirically, tuning stepsize γ and batchsize m warrants good performance. Please also see our experiments.

Remark 6. Theorem 3.4 implies the improved performance of minibatch `SGD` for composite problems (3) but leaves out the case where $\ell(x, \xi)$ is non-smooth. In the later, it seems difficult to improve the results further, since the bound $\mathcal{O}(\frac{L^2\Delta}{\varepsilon^4})$ of `SGD` for general non-smooth problems already matches the best result for deterministic subgradient method. Please refer to [8].

Remark 7. Our theoretical result focuses on iteration complexity of minibatch algorithms, skipping the discussion about how to process the batch samples in parallel. We refer to the appendix for efficient routines for the minibatch subproblems.

4 `SMOD` with Momentum

We present a new model-based method by incorporating an additional extrapolation step, and we record this stochastic extrapolated model-based method in Algorithm 2. Each iteration of Algorithm 2

Algorithm 2 Stochastic Extrapolated Model-Based Method

Input: x^0, x^1, β, γ
for $k = 1$ **to** K **do**
Sample data ξ^k and update:

$$y^k = x^k + \beta(x^k - x^{k-1}) \tag{11}$$

$$x^{k+1} = \operatorname*{argmin}_{x \in \mathcal{X}} \; f_{x^k}(x, \xi^k) + \frac{\gamma}{2}\|x - y^k\|^2 \tag{12}$$

end for

consists of two steps, first, an extrapolation is performed to get an auxiliary update y^k. Then a random sample ξ_k is collected and the proximal mapping, associated with the model function $f_{x^k}(\cdot, \xi_k)$, is computed at y^k to obtain the new point x^{k+1}. For ease of exposition, we take constant values of stepsize and extrapolation term.

Note that Algorithm 2 can be interpreted as an extension of the momentum SGD by replacing the gradient descent step with a broader class of proximal mappings. To see this intuition, we combine (11) and (12) to get

$$x^{k+1} = \operatorname*{argmin}_{x \in \mathcal{X}} \Big\{ f_{x^k}(x, \xi^k) + \gamma\beta\langle x^{k-1} - x^k, x - x^k\rangle + \frac{\gamma}{2}\|x - x^k\|^2 \Big\}, \tag{13}$$

If we choose the linear model (3), i.e., $f_{x^k}(x, \xi^k) = f(x, \xi^k) + \langle f'(x, \xi^k), x - x^k\rangle$, and assume $\mathcal{X} = \mathbb{R}^d$, then the update (13) has the following form:

$$x^{k+1} = x^k - \gamma^{-1} f'(x, \xi^k) - \beta(x^{k-1} - x^k). \tag{14}$$

Define $v^k \triangleq \gamma(x^{k-1} - x^k)$ and apply it to (14), then Algorithm 2 reduces to the heavy-ball method

$$v^{k+1} = f'(x, \xi^k) + \beta v^k, \tag{15}$$

$$x^{k+1} = x^k - \gamma^{-1} v^{k+1}. \tag{16}$$

Despite such relation, the gradient averaging view (15) only applies to SGD for unconstrained optimization, which limits the use of standard analysis of heavy-ball method ([28]) for our problem. To overcome this issue, we present a unified convergence analysis which can deal with all the model functions and is amenable to both constrained and composite problems.

Our theoretical analysis of Algorithm 2 relies on a different potential function from the one in previous section. More specifically, let us define

$$z^k \triangleq x^k + \frac{\beta}{1-\beta}(x^k - x^{k-1}). \tag{17}$$

The following lemma proves some approximate descent property by adopting the potential function

$$f_{1/\rho}(z^k) = \min_{x \in \mathcal{X}} f(x) + \frac{\rho}{2}\|x - z^k\|^2, \tag{18}$$

and measuring the quantity of $\|\nabla f_{1/\rho}(z^k)\|$.

Lemma 4.1. *Assume that $\rho \geq 2(\tau + \lambda)$ and $\beta \in [0, 1)$. Let $\theta = 1 - \beta$. Then we have*

$$\begin{aligned}\frac{(\rho - \lambda\theta)}{2\rho(\gamma\theta - \lambda\theta)}\|\nabla f_{1/\rho}(z^k)\|^2 \leq\; & f_{1/\rho}(z^k) - \mathbb{E}_k\big[f_{1/\rho}(z^{k+1})\big] + \frac{\rho L^2}{(\gamma\theta^2 - \rho\beta^2\theta^{-1})(\gamma\theta^2 - \lambda\theta^2)} \\ & + \frac{\rho(\gamma\beta + \rho\beta^2\theta^{-2})}{2(\gamma\theta - \lambda\theta)}\big(\|x^k - x^{k-1}\|^2 - \mathbb{E}_k[\|x^{k+1} - x^k\|^2]\big) \\ & - \frac{\rho(\gamma - \rho\beta^2\theta^{-3})}{4(\gamma - \lambda)}\mathbb{E}_k[\|x^{k+1} - x^k\|^2].\end{aligned} \tag{19}$$

Invoking Lemma 4.1 and specifying the stepsize policy, we obtain the main convergence result of Algorithm 2 in the following theorem.

Theorem 4.2. *Under assumptions of Lemma 4.1, if we choose $x^1 = x^0$, and set $\gamma = \gamma_0\theta^{-1}\sqrt{K} + \lambda + \rho\beta^2\theta^{-3}$ for some $\gamma_0 > 0$, then*

$$\mathbb{E}[\|\nabla f_{1/\rho}(z^{k^*})\|^2] \leq \frac{2\rho}{\rho-\lambda}\left[\frac{\rho\beta^2\theta^{-2}\Delta}{K} + \left(\gamma_0\Delta + \frac{\rho L^2}{\theta\gamma_0}\right)\frac{1}{\sqrt{K}}\right] \tag{20}$$

where k^* is an index chosen in $\{1, 2, \ldots, K\}$ uniformly at random.

Remark 8. Despite the fact that convergence is established for all $\gamma_0 > 0$, we can see that the optimal γ_0 would be $\gamma_0 = \sqrt{\frac{\rho}{\Delta\theta}}L$, which gives the bound $\mathbb{E}[\,\|\nabla f_{1/\rho}(z^{k^*})\|^2] \leq \frac{2\rho}{\rho-\lambda}\big(\frac{\rho\beta^2\theta^{-2}\Delta}{K} + 2L\sqrt{\frac{\rho\Delta}{\theta K}}\big)$. In practice we can set γ_0 to a suboptimal value, and obtain a possibly loose upper-bound.

Remark 9. Since z^k is an extrapolated solution, it may not be feasible. It is desirable to show optimality guarantee at iterates x^k. Note that using Lemma 4.1 and the parameters in Theorem 4.2, it is easy to show that $\mathbb{E}[\|x^{k^*} - x^{k^*-1}\|^2] = \mathcal{O}(\frac{1}{K})$. Based on (17) we have $\|z^{k^*} - x^{k^*}\|^2 = \beta^2\theta^{-2}\mathbb{E}[\|x^{k^*} - x^{k^*-1}\|^2] = \mathcal{O}(\frac{1}{K})$. Using Lipschitz smoothness of Moreau envelop, we can show $\mathbb{E}[\|\nabla f_{1/\rho}(x^{k^*})\|^2]$ converges at the same $\mathcal{O}(\frac{1}{\sqrt{K}})$ rate as is shown in Theorem 4.2.

Combining the momentum and minibatching techniques, we can develop a minibatch version of Algorithm 2 that takes a batch of samples B_k in each iteration. The convergence analysis of this new extension requires a more complicated potential function, and hence is more involving. We state the main result informally but leave the details in Appendix session.

Theorem 4.3 (Informal)**.** *In the momentum `SMOD` method with minibatching, if we additionally assume A5, set batchsize $|B_k| = m$ and take $\gamma = \mathcal{O}(\sqrt{K/m})$, then $\mathbb{E}[\,\|\nabla f_{1/\rho}(z^{k^*})\|^2] = \mathcal{O}(\frac{1}{K} + \sqrt{\frac{1}{mK}})$.*

5 Experiments

In this section, we examine the empirical performance of our proposed methods through experiments on the problem of robust phase retrieval. (Additional experiments on blind deconvolution are given in Appendix section). Given a set of vectors $a_i \in \mathbb{R}^d$ and nonnegative scalars $b_i \in \mathbb{R}_+$, the goal of phase retrieval is to recover the true signal x^* from the measurement $b_i = |\langle a_i, x^*\rangle|^2$. Due to the potential corruption in the dataset, we consider the following penalized formulation

$$\min_{x\,\in\,\mathbb{R}^d} \quad \frac{1}{n}\sum_{i=1}^{n}\left|\langle a_i, x\rangle^2 - b_i\right| \tag{21}$$

where we impose ℓ_1-loss to promote robustness and stability (cf. [13, 8, 23]).

Data Preparation. We conduct experiments on both synthetic and real datasets.

1) Synthetic Data. Synthetic data is generated following the setup in [23]. We set $n = 300, d = 100$ and select x^* from unit sphere uniformly at random. Moreover, we generate $A = QD$ where $Q \in \mathbb{R}^{n\times d}, q_{ij} \sim \mathcal{N}(0,1)$ and $D \in \mathbb{R}^d$ is a diagonal matrix whose diagonal entries are evenly distributed in $[1/\kappa, 1]$. Here $\kappa \geq 1$ plays the role of condition number (large κ makes problem hard). The measurements are generated by $b_i = \langle a_i, x^*\rangle^2 + \delta_i\zeta_i$ $(1 \leq i \leq n)$ with $\zeta_i \sim \mathcal{N}(0, 25)$, $\delta_i \sim \mathrm{Bernoulli}(p_{\text{fail}})$, where $p_{\text{fail}} \in [0,1]$ controls the fraction of corrupted observations on expectation.

2) Real Data. Furthermore, we consider `zipcode`, a dataset of 16×16 handwritten digits collected from [19]. Following the setup in [13], let $H \in \mathbb{R}^{256\times 256}$ be a normalized Hadamard matrix such that $h_{ij} \in \left\{\frac{1}{16}, -\frac{1}{16}\right\}, H = H^{\mathrm{T}}$ and $H = H^{-1}$. Then we generate $k = 3$ diagonal sign matrices S_1, S_2, S_3 such that each diagonal element of S_k is uniformly sampled from $\{-1, 1\}$. Last we set $A = [HS_1, HS_2, HS_3]^{\mathrm{T}} \in \mathbb{R}^{(3\times 256)\times 256}$. As for the true signal and measurements, each image is represented by a data matrix $X \in \mathbb{R}^{16\times 16}$ and gets vectorized to $x^* = \mathrm{vec}(X)$. To simulate the case of corruption, we set measurements $b = \phi_{p_{\text{fail}}}(Ax^*)$, where $\phi_{p_{\text{fail}}}(\cdot)$ denotes element-wise squaring and setting a fraction p_{fail} of entries to 0 on expectation.

In the first experiment, we illustrate that `SMOD` methods enjoy linear speedup with size of minibatches and exhibit strong robustness to the stepsize policy. We conduct comparison on `SPL` and `SGD` and describe the detailed experiment setup as follows.

1) Dataset Generation. We generate four testing cases: the synthetic datasets with $(\kappa, p_{\text{fail}}) = (10, 0.2)$, and $(10, 0.3)$; `zipcode` with digit images of id 2 and 24;

2) Initial Point. For all the algorithms, we set the initial point $x^1 (= x^0) \sim \mathcal{N}(0, I_d)$ for synthetic data and $x^1 = x^* + 2 \cdot \mathcal{N}(0, I_d)$ for `zipcode`;

3) Stepsize. We set the parameter $\gamma = \alpha_0^{-1}\sqrt{K/m}$ where m is the batchsize; For synthetic dataset, we test 10 evenly spaced α_0 values in range $[10^{-1}, 10^2]$ on logarithmic scale, and for `zipcode` dataset we set such range of α_0 to $[10^1, 10^3]$;

4) Maximum Iteration. For both datasets the maximum number of epochs is set 400;

5) Batchsize. The batchsize m is from the range $\{1, 4, 8, 16, 32, 64\}$;

6) Sub-problems The solution to proximal sub-problems is described in the appendix.

For each algorithm, speedup from minibatching is quantified as T_1^*/T_m^* where T_m^* is the total number of iterations for reaching the desired accuracy, with minibatchsize m and the best initial stepsize α_0 among values specified above. Specially, if an algorithm fails to reach desired accuracy after running out of 400 epochs, we set its iteration number to the maximum.

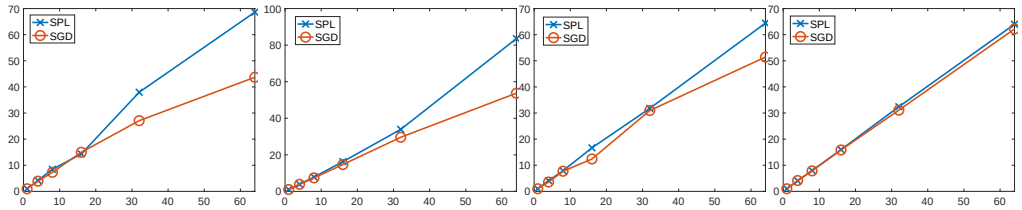

Figure 1: Speedup over minibatch sizes. The left two are for synthetic datasets $\kappa = 10, p_{\text{fail}} \in \{0.2, 0.3\}$; Digit datasets: digit image (id:24) with $p_{\text{fail}} \in \{0.2, 0.3\}$.

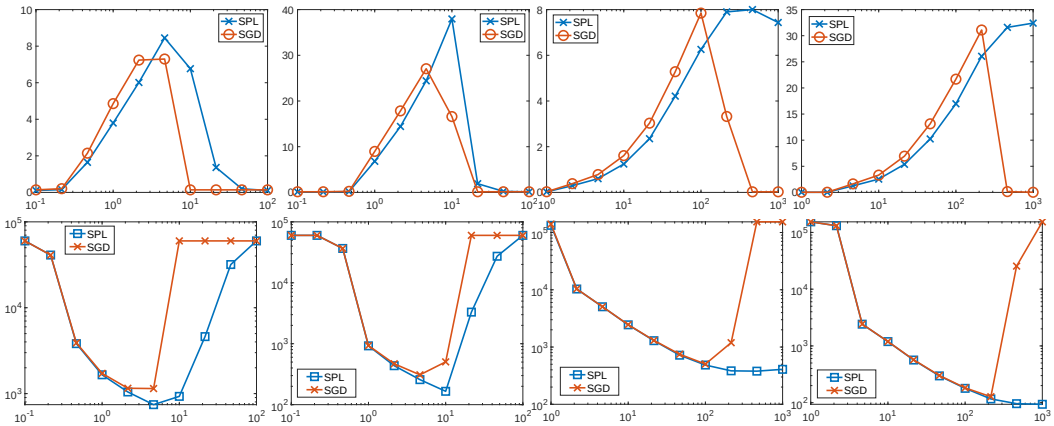

Figure 2: From left to right: synthetic datasets with $m \in \{8, 32\}$ and `zipcode` image (id=24) with $m \in \{8, 32\}$. x-axis: initial stepsize α_0. y-axis (first row): speedup over the sequential version: $T_1^*/T_m^*(\alpha_0)$ where $T_m^*(\alpha_0)$ stands for the number of iterations when using minibatchsize m and initial stepsize α_0. y-axis (second row): Total number of iterations.

Figure 1 plots the speedup of each algorithm over different values of batchsize according to the average of 20 independent runs. It can be seen that `SPL` exhibits linear acceleration over the batchsize, which confirms our theoretical analysis. Moreover, we find that `SGD` admits considerable acceleration using minibatches, and sometimes the speedup performance matches that of `SPL`. This observation seems to suggest the effectiveness of minibatch `SGD` in practice, despite the lack of theoretical support.

Next we investigate the sensitivity of minibatch acceleration to the choice of initial stepsizes. We plot the algorithm speedup over the initial stepsize α_0 in Figure 2 (1st row). It can be readily seen that, both `SGD` and `SPL` achieve considerable minibatch acceleration when choosing the initial stepsize properly. However, `SPL` has a much wider range of initial stepsizes for good speedup performance, and hence, lays more robust performance than `SGD`. To further illustrate the robustness of `SPL`, we compare the efficiency of both algorithms in the minibatch setting. In contrast to the previous comparison on the

relative scale, we directly compare the iteration complexity of the two algorithms. We plot the total iteration number over the choice of initial stepsizes in Figure 2 (2nd row) for batchsize $m = 8$ and 32. We observe that minibatch SPL exhibits promising performance for a wide range of stepsize policies while minibatch SGD quickly diverges for large stepsizes. Overall, our experiment complements the recent work [8], which shows that SPL is more robust than SGD in the sequential setting.

Our second experiment investigates the performance of the proposed momentum model-based methods. We compare three model-based methods (SGD, SPL, SPP) and extrapolated model-based methods (SEGD, SEPL, SEPP). We generate four testing cases: the synthetic datasets with $(\kappa, p_{\text{fail}}) = (10, 0.2)$ and $(10, 0.3)$; zipcode with digit images of id 2 and $p_{\text{fail}} \in \{0.2, 0.3\}$. For synthetic data we use $\alpha_0 \in [10^{-2}, 10^0], \beta = 0.6$; for zipcode dataset we use $\alpha_0 \in [10^0, 10^1], \beta = 0.9$. The rest of settings are the same as in minibatch with $m = 1$.

Figure 3 plots the number of epochs to ε-accuracy over initial stepsize a_0. It can be seen that with properly selected momentum parameters, (SEGD, SEPL, SEPP) all suggest improved convergence when stepsize is relatively small.

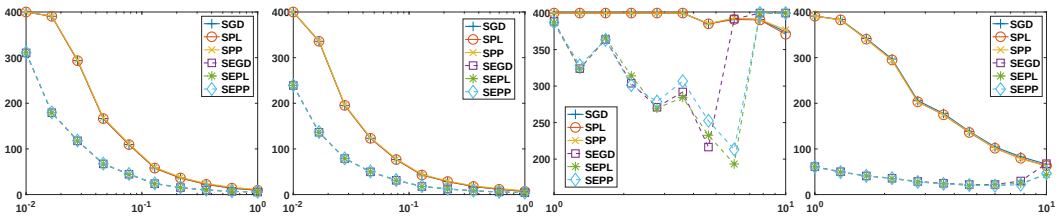

Figure 3: From left to right: synthetic datasets with $\kappa = 10, p_{\text{fail}} \in \{0.2, 0.3\}, \beta = 0.6$ and zipcode image (id=2) with $p_{\text{fail}} \in \{0.2, 0.3\}, \beta = 0.9$. x-axis: initial stepsize α_0. y-axis: number of epochs on reaching desired accuracy

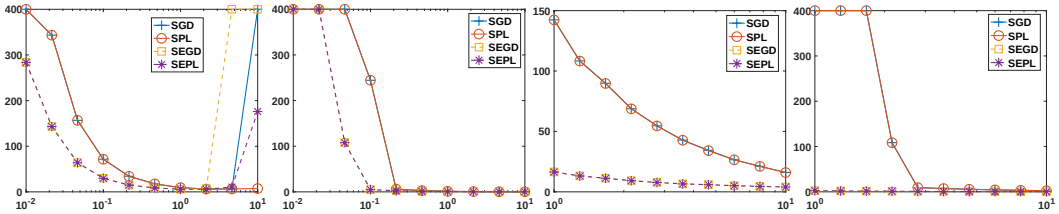

Figure 4: From left to right: synthetic datasets with $\kappa = 10, p_{\text{fail}} = 0.3, \beta = 0.6, m \in \{1, 32\}$ and zipcode image (id=24) with $p_{\text{fail}} = 0.3, \beta = 0.9, m \in \{1, 32\}$. x-axis: initial stepsize α_0. y-axis: number of epochs for reaching desired accuracy

In the last experiment, we attempt to exploit the performance of the compared algorithms when minibatching and momentum are applied simultaneously. We use the same settings as in the second experiment but with $m \in \{8, 32\}$. Results are plotted in Figure 4 and it can be seen that minibatch, when combined with momentum, exhibits even better convergence and robustness.

6 Discussion

On a broad class of non-smooth non-convex problems, we make stochastic model-based methods more competitive against SGD by leveraging two well-known techniques: minibatching and momentum. Applying algorithm stability for optimization analysis is key to our improved results, and it appears to have great potential for stochastic optimization in a much broader context. One limitation is that we are unable to show whether minibatching can accelerate SGD in the most general non-smooth case. The main difficulty can be seen from the fact that the complexity of SGD already matches the bound of full subgradient method. It would be interesting to know whether this bound is tight or improvable for SGD. It would also be interesting to study the lower bound of SGD (and other stochastic algorithms) in non-smooth setting, and to study fundamental questions on the appropriate optimality criteria in the non-smooth non-convex optimization. Some interesting recent results can be referred from [20, 29].

References

[1] H. Asi and J. C. Duchi. The importance of better models in stochastic optimization. *Proceedings of the National Academy of Sciences*, 116(46):22924–22930, 2019.

[2] H. Asi, K. Chadha, G. Cheng, and J. C. Duchi. Minibatch stochastic approximate proximal point methods. *Advances in Neural Information Processing Systems*, 33, 2020.

[3] H. H. Bauschke, P. L. Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.

[4] L. Berrada, A. Zisserman, and M. P. Kumar. Deep frank-wolfe for neural network optimization. In *ICLR 2019 : 7th International Conference on Learning Representations*, 2019.

[5] A. Botev, H. Ritter, and D. Barber. Practical gauss-newton optimisation for deep learning. In *International Conference on Machine Learning*, pages 557–565. PMLR, 2017.

[6] O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.

[7] V. Charisopoulos, Y. Chen, D. Davis, M. Díaz, L. Ding, and D. Drusvyatskiy. Low-rank matrix recovery with composite optimization: good conditioning and rapid convergence. *arXiv preprint arXiv:1904.10020*, 2019.

[8] D. Davis and D. Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *Siam Journal on Optimization*, 29(1):207–239, 2019.

[9] A. Defazio. Understanding the role of momentum in non-convex optimization: Practical insights from a lyapunov analysis. *arXiv preprint arXiv:2010.00406*, 2020.

[10] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13(1), 2012.

[11] D. Drusvyatskiy and C. Paquette. Efficiency of minimizing compositions of convex functions and smooth maps. *Mathematical Programming*, pages 1–56, 2018.

[12] J. C. Duchi and F. Ruan. Stochastic methods for composite and weakly convex optimization problems. *SIAM Journal on Optimization*, 28(4):3229–3259, 2018.

[13] J. C. Duchi and F. Ruan. Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. *Information and Inference: A Journal of the IMA*, 8(3):471–529, 2019.

[14] T. Frerix, T. Möllenhoff, M. Moeller, and D. Cremers. Proximal backpropagation. In *International Conference on Learning Representations*, 2018.

[15] S. Ghadimi and G. Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013. ISSN 1052-6234.

[16] S. Ghadimi, G. Lan, and H. Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2):267–305, 2016.

[17] I. Gitman, H. Lang, P. Zhang, and L. Xiao. Understanding the role of momentum in stochastic gradient methods. In *Advances in Neural Information Processing Systems*, volume 32, pages 9633–9643, 2019.

[18] M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pages 1225–1234. PMLR, 2016.

[19] T. Hastie, R. Tibshirani, J. Friedman, T. Hastie, J. Friedman, and R. Tibshirani. *The elements of statistical learning*, volume 2. Springer, 2009.

[20] G. Kornowski and O. Shamir. Oracle complexity in nonsmooth nonconvex optimization. *arXiv preprint arXiv:2104.06763*, 2021.

[21] Y. Liu, Y. Gao, and W. Yin. An improved analysis of stochastic gradient descent with momentum. *arXiv preprint arXiv:2007.07989*, 2020.

[22] N. Loizou and P. Richtárik. Momentum and stochastic momentum for stochastic gradient, newton, proximal point and subspace descent methods. *Computational Optimization and Applications*, 77(3):653–710, 2020.

[23] V. Mai and M. Johansson. Convergence of a stochastic gradient method with momentum for non-smooth non-convex optimization. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6630–6639, 2020.

[24] O. Sebbouh, R. M. Gower, and A. Defazio. On the convergence of the stochastic heavy ball method. *arXiv preprint arXiv:2006.07867*, 2020.

[25] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670, 2010.

[26] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*, pages 1139–1147, 2013.

[27] M. Takáč, P. Richtárik, and N. Srebro. Distributed mini-batch sdca. *arXiv preprint arXiv:1507.08322*, 2015.

[28] Y. Yan, T. Yang, Z. Li, Q. Lin, and Y. Yang. A unified analysis of stochastic momentum methods for deep learning. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pages 2955–2961, 2018.

[29] J. Zhang, H. Lin, S. Jegelka, A. Jadbabaie, and S. Sra. Complexity of finding stationary points of nonsmooth nonconvex functions. *arXiv preprint arXiv:2002.04130*, 2020.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See remarks 5, 6 in Section 3 and Discussion section
   (c) Did you discuss any potential negative societal impacts of your work? [No] The paper addresses theoretical questions on algorithm complexity, which, to the best of our knowledge, has no negative social impact
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See assumptions in Section 2 and 3
   (b) Did you include complete proofs of all theoretical results? [Yes] Proof is left in the appendix
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Code is supplied in the supplemental materials
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 5 for details of experiments

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] From the experiments the error bars are relatively thin and the results are presented by taking average.

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] The main goal of experiments is to demonstrate our theoretical foundings, thereby only showing the iteration complexity of algorithms.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] `zipcode` dataset is referenced from [19].

(b) Did you mention the license of the assets? [No] The dataset used is published on an open site without license.

(c) Did you include any new assets either in the supplemental material or as a URL? [No] The experiments do not involve new datasets.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] An open dataset is used.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] The dataset has been open for years and only involves zipcode digits.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [No] No crowdsourcing or human object is involved.

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [No] No crowdsourcing or human object is involved.

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [No] No crowdsourcing or human object is involved.