

UNIVERSAL STAGewise LEARNING FOR NON-CONVEX PROBLEMS WITH CONVERGENCE ON AVERAGED SOLUTIONS

Zaiyi Chen *

University of Science and Technology of China

Zhuoning Yuan *

University of Iowa

Jinfeng Yi

JD AI Research

Bowen Zhou

JD AI Research

Enhong Chen

University of Science and Technology of China

Tianbao Yang

University of Iowa

ABSTRACT

Although stochastic gradient descent (SGD) method and its variants (e.g., stochastic momentum methods, ADAGRAD) are algorithms of choice for solving non-convex problems (especially deep learning), big gaps still remain between the theory and the practice with many questions unresolved. For example, there is still a lack of theories of convergence for SGD and its variants that use stagewise step size and return an averaged solution in practice. In addition, theoretical insights of why adaptive step size of ADAGRAD could improve non-adaptive step size of SGD is still missing for non-convex optimization. This paper aims to address these questions and fill the gap between theory and practice. We propose a universal stagewise optimization framework for a broad family of **non-smooth non-convex** problems with the following key features: (i) at each stage any suitable stochastic convex optimization algorithms (e.g., SGD or ADAGRAD) that return an averaged solution can be employed for minimizing a regularized convex problem; (ii) the step size is decreased in a stagewise manner; (iii) an averaged solution is returned as the final solution. Our theoretical results of stagewise ADAGRAD exhibit its adaptive convergence, therefore shed insights on its faster convergence than stagewise SGD for problems with slowly growing cumulative stochastic gradients. To the best of our knowledge, this is the first work for addressing the unresolved issues of existing theories mentioned earlier. Besides theoretical contributions, our empirical studies show that our stagewise SGD and ADAGRAD improve the generalization performance of existing variants/implementations of SGD and ADAGRAD.

1 INTRODUCTION

Non-convex optimization has recently received increasing attention due to its popularity in emerging machine learning tasks, particularly for learning deep neural networks. One of the keys to the success of deep learning for big data problems is the employment of simple stochastic algorithms such as SGD or ADAGRAD (Krizhevsky et al., 2012; Dean et al., 2012). Analysis of these stochastic algorithms for non-convex optimization is an important and interesting research topic, which already attracts much attention from the community of theoreticians (Ghadimi & Lan, 2013; 2016; Yang et al., 2016; Davis & Drusvyatskiy, 2018b; Ward et al., 2018; Li & Orabona, 2018). However, one issue that has been largely ignored in existing theoretical analysis is that the employed algorithms in practice usually differ from their plain versions that are well understood in theory. Below, we will discuss several important *heuristics* used in practice for training deep neural networks and the *gap* between the practice and the theory to motivate this work.

First, a trick for setting the step size in training deep neural networks is to change it in a stagewise manner from a large value to a small value, i.e., a constant step size is used in a stage for a number of iterations and is decreased for the next stage (Krizhevsky et al., 2012; Ren et al., 2018). Although

*Equal contribution. Correspondence to: tianbao-yang@uiowa.edu

this trick has been adopted by most open-sourced libraries, e.g., Caffe (Jia et al., 2014), TensorFlow (Abadi et al., 2015), Pytorch (Paszke et al., 2017), it still lacks theoretical analysis to date for non-convex optimization. A related question that stands out is how and when to decrease the step size. On standard benchmark datasets for academic use such as CIFAR-10, CIFAR-100 (Krizhevsky et al.), people could follow the setting reported in previous studies to get a good result, which however might not work well for new datasets. Hence, a better solution to setting the stagewise step size with insights from theory would be much preferred. However, in the existing literature of theory for non-convex optimization (Ghadimi & Lan, 2013; Davis & Drusvyatskiy, 2018b), only strategies based on an iteratively decreasing step size or a small constant step size have been well analyzed. For example, the existing theory usually suggests an iteratively decreasing step size proportional to $1/t$ at the t -th iteration or a small constant step size, e.g., proportional to ϵ with $\epsilon < 1$ for finding an ϵ -stationary solution whose gradient’s magnitude (in expectation) is small than ϵ .

Second, the averaging heuristic is usually used in practice, i.e., an averaged solution is returned for prediction (Bottou, 2010), which could yield improved stability and generalization (Hardt et al., 2016). However, existing theory for many stochastic non-convex optimization algorithms only provides guarantee on a uniformly sampled solution or a non-uniformly sampled solution with *decreasing* probabilities for latest solutions (Ghadimi & Lan, 2013; Yang et al., 2016; Davis & Drusvyatskiy, 2018b). In particular, if an iteratively decreasing step size proportional to $1/t$ at the t -th iteration is employed, the convergence guarantee was provided for a random solution that is non-uniformly selected from all iterates with a sampling probability proportional to $1/t$ for the t -th iterate. This means that the latest solution always has the smallest probability to be selected as the final solution, which contradicts to the common wisdom. If a small constant step size is used, then usually a uniformly sampled solution is provided with convergence guarantee. However, both options are rarely used in practice and cannot justify the heuristic that returns the last solution.

A third common approach in practice is to use adaptive coordinate-wise step size of ADAGRAD (Dean et al., 2012). Although adaptive step size has been well analyzed for convex problems (e.g., why and when it could yield faster convergence than SGD) (Duchi et al., 2011; Chen et al., 2018b), it still remains an mystery for non-convex optimization with missing insights from theory. Several recent studies have attempted to analyze ADAGRAD for non-convex problems (Ward et al., 2018; Li & Orabona, 2018; Chen et al., 2018a; Zou & Shen, 2018). Nonetheless, none of them are able to exhibit the adaptive convergence of ADAGRAD to data as in the convex case and its potential advantage over SGD for non-convex problems.

To overcome the shortcomings of existing theories for stochastic non-convex optimization, this paper analyzes new algorithms that employ some or all of these commonly used heuristics in a systematic framework, aiming to fill the gap between theory and practice. The main results and contributions are summarized below:

We propose a universal stagewise optimization framework for solving a family of non-convex problems, i.e., weakly convex problems, which is broader than smooth non-convex problems and includes some non-smooth non-convex problems. At each stage, any suitable stochastic convex optimization algorithms (e.g., SGD, ADAGRAD) with a constant step size parameter can be employed for optimizing a regularized convex problem with a number of iterations, which usually return an averaged solution. The step size parameter is decreased in a stagewise manner following a polynomial decaying scheme.

We analyze several variants of the proposed framework by employing different basic algorithms, including SGD, ADAGRAD, stochastic heavy-ball (SHB) method, and stochastic Nesterov’s accelerated gradient (SNAG) method. We prove the convergence of their stagewise versions for an averaged solution that is randomly selected from all stagewise averaged solutions.

To justify a heuristic approach that returns the last averaged solution in stagewise learning, we present and analyze a non-uniform sampling strategy over stagewise averaged solutions with sampling probabilities increasing as the stage number.

Regarding the convergence results, for stagewise SGD, SHB, SNAG, we establish the same order of iteration complexity for finding a nearly stationary point as the existing theories of their non-stagewise variants. For stagewise ADAGRAD, we establish an adaptive convergence for finding a nearly stationary point, which is provably better than (stagewise) SGD, SHB, and SNAG when the cumulative growth of stochastic gradient is slow.

Besides theoretical contributions, we also empirically verify the effectiveness of the proposed stagewise algorithms. In particular, our empirical studies show that (i) the stagewise ADAGRAD dramatically improves the generalization performance of existing variants of ADAGRAD, (ii) stagewise SGD, SHB, SNAG also outperform their plain variants with an iteratively decreasing step size; (iii) the proposed stagewise algorithms achieve similar if not better generalization performance than their heuristic variants implemented in existing libraries on standard benchmark datasets.

2 RELATED WORK

SGD for unconstrained smooth non-convex problems was first analyzed by Ghadimi & Lan (2013), who established an $O(1/\epsilon^4)$ iteration complexity for finding an ϵ -stationary point \mathbf{x} in expectation satisfying $E[k \nabla f(\mathbf{x})^2] \leq \epsilon^2$, where $f(\cdot)$ denotes the objective function. As mentioned earlier, the returned solution is either a uniformly sampled solution or a non-uniformly sampled one with sampling probabilities proportional to the decreasing step size. Similar results were established for the stochastic momentum variants of SGD (i.e., SHB, SNAG) by Yang et al. (2016); Ghadimi & Lan (2016). Recently, SGD was also analyzed for (constrained) weakly convex problems, whose objective function is non-convex and not necessarily smooth, by Davis & Drusvyatskiy (2018b). However, none of these studies provide results for algorithms that return an averaged solution, and these analyzed algorithms differ significantly from that used in practice for achieving the state-of-the-art results (Krizhevsky et al., 2012; Ren et al., 2018; Loshchilov & Hutter, 2017).

Although adaptive variants of SGD, e.g., ADAGRAD (Duchi et al., 2011), ADAM (Kingma & Ba, 2015; Reddi et al., 2018), were widely used for training deep neural networks, there are few studies on theoretical analysis of these algorithms for non-convex problems. Several recent studies attempted to analyze ADAGRAD for non-convex problems (Ward et al., 2018; Li & Orabona, 2018; Chen et al., 2018a; Zou & Shen, 2018). Although these studies have established an iteration complexity of $O(1/\epsilon^4)$ for different variants of ADAGRAD for finding an ϵ -stationary solution of a stochastic non-convex optimization problem, none of them can exhibit the potential adaptive advantage of ADAGRAD over SGD as in the convex case. Besides that, these studies also suffer from the following shortcomings: (i) they all assume smoothness of the objective function, while we consider non-smooth and non-convex problems; (ii) their convergence is provided on a solution with minimum magnitude of gradient that is expensive to compute, though their results also imply a convergence on a random solution selected from all iterates with decreasing sampling probabilities. In contrast, these shortcomings do not exist in this paper. To the best of our knowledge, our result is the first one that explicitly shows that coordinate-wise adaptive step size could yield faster convergence than using non-adaptive step size for *non-smooth non-convex problems*, which is similar to that in the convex case and was observed in practice for deep learning (Dean et al., 2012).

The proposed stagewise algorithm is similar to several existing algorithms in design (Xu et al., 2017; Davis & Grimmer, 2017), which are originated from the proximal point algorithm (Rockafellar, 1976). I.e., at each stage a regularized convex subproblem is formed and then a stochastic algorithm is employed for optimizing the regularized subproblem inexactly with a number of iterations. Xu et al. (2017) used this idea for solving problems that satisfy a local error bound condition, aiming to achieve faster convergence than vanilla SGD. Davis & Grimmer (2017) are probably the first who analyzed this idea for solving non-smooth weakly convex problems. In these two papers SGD with decreasing step sizes for a strongly convex problem is employed at each stage for solving the regularized subproblem. Our stagewise algorithm is developed following the similar idea. The key differences from (Xu et al., 2017; Davis & Grimmer, 2017) are that (i) we focus on non-convex problems instead of convex problems considered in (Xu et al., 2017); (ii) we analyze a non-uniform sampling strategy with sampling probabilities increasing as the stage number to justify a heuristic approach that uses the last averaged solution for prediction, unlike the uniform sampling used in (Davis & Grimmer, 2017); (iii) we present a unified algorithmic framework and convergence analysis, which enable one to employ any suitable stochastic convex optimization algorithms at each stage. Importantly, it brings us several interesting variants including stagewise stochastic momentum methods and stagewise ADAGRAD. Finally, we note that a similar idea that adds a strongly convex term into the non-convex objective has been considered in several recent works (Allen-Zhu, 2017; Lan & Yang, 2018; Carmon et al., 2016). However, the key difference between these works and the present work is that they need to assume the objective function is smooth or has a smooth component. In addition, Allen-Zhu (2017); Lan & Yang (2018) consider finite-sum problems, and Carmon et al.

(2016) consider deterministic problems. In contrast, we consider more general stochastic problems without assuming the objective function is smooth.

3 PRELIMINARIES

The problem of interest in this paper is:

$$\min_{\mathbf{x} \in U} f(\mathbf{x}) = \mathbb{E} [f(\mathbf{x}; \mathbf{y})]; \quad (1)$$

where $U \subset \mathbb{R}^d$ is a closed convex set, $\mathbf{y} \in U$ is a random variable, $f(\mathbf{x})$ and $f(\mathbf{x}; \mathbf{y})$ are non-convex functions, with the basic assumptions on the problem given in Assumption 1.

To state the convergence property of an algorithm for solving the above problem. We need to introduce some definitions. These definitions can be also found in related literature, e.g., Davis & Grimmer (2017); Davis & Drusvyatskiy (2018b). In the sequel, we let $\|\cdot\|$ denote an Euclidean norm, $\mathbb{1}_S = \mathbb{1}_{\{\cdot \in S\}}$; $S \subset \mathbb{R}^d$ denote a set, and $\mathbb{1}_S(\cdot)$ denote the indicator function of the set S .

Definition 1. (Fréchet subgradient) For a non-smooth and non-convex function $f(\cdot)$,

$$\partial_F f(\mathbf{x}) = \{\mathbf{v} \in \mathbb{R}^d \mid f(\mathbf{y}) \geq f(\mathbf{x}) + \mathbf{v}^\top (\mathbf{y} - \mathbf{x}) + o(\|\mathbf{y} - \mathbf{x}\|)\} \subset \mathbb{R}^d$$

denotes the Fréchet subgradient of f .

Definition 2. (First-order stationarity) For problem (1), a point $\mathbf{x} \in U$ is a first-order stationary point if $0 \in \partial_F f(\mathbf{x})$, where $\mathbb{1}_S$ denotes the indicator function of S . Moreover, a point \mathbf{x} is said to be ϵ -stationary if $\text{dist}(0, \partial_F f(\mathbf{x})) \leq \epsilon$, where dist denotes the Euclidean distance from a point to a set.

Definition 3. (Moreau Envelope and Proximal Mapping) For any function f and $\mu > 0$, the following function is called a Moreau envelope of f

$$f_\mu(\mathbf{x}) = \min_{\mathbf{z}} f(\mathbf{z}) + \frac{\mu}{2} \|\mathbf{x} - \mathbf{z}\|^2; \quad (2)$$

Further, the optimal solution to the above problem denoted by

$$\text{prox}_\mu f(\mathbf{x}) = \arg \min_{\mathbf{z}} f(\mathbf{z}) + \frac{\mu}{2} \|\mathbf{x} - \mathbf{z}\|^2$$

is called a proximal mapping of f .

Definition 4. (Weakly convex) A function f is μ -weakly convex ($\mu > 0$), if $f(\mathbf{x}) + \frac{\mu}{2} \|\mathbf{x}\|^2$ is convex.

It is known that if $f(\mathbf{x})$ is μ -weakly convex and $\mu < L^{-1}$, then its Moreau envelope $f_\mu(\mathbf{x})$ is C^1 -smooth with the gradient given by (see e.g., Davis & Drusvyatskiy (2018b))

$$\nabla f_\mu(\mathbf{x}) = -\mathbf{1}_{\{\mathbf{x} \neq \text{prox}_\mu f(\mathbf{x})\}} \frac{\mathbf{x} - \text{prox}_\mu f(\mathbf{x})}{\mu};$$

The tool of Moreau envelope is introduced to measure the convergence for optimizing non-smooth and non-convex functions. A small norm of $\nabla f_\mu(\mathbf{x})$ has an interpretation that \mathbf{x} is close to a point that is nearly stationary. In particular for any $\mathbf{x} \in \mathbb{R}^d$, let $\mathbf{x}^* = \text{prox}_\mu f(\mathbf{x})$, then we have (Davis & Drusvyatskiy, 2018b)

$$f(\mathbf{x}^*) \leq f(\mathbf{x}); \quad \|\mathbf{x} - \mathbf{x}^*\| \leq \mu \|\nabla f_\mu(\mathbf{x})\|; \quad \text{dist}(0, \partial f(\mathbf{x}^*)) \leq \mu \|\nabla f_\mu(\mathbf{x})\|; \quad (3)$$

where the first inequality follows from the definition of $\text{prox}_\mu f(\mathbf{x})$ such that $f(\mathbf{x}^*) + \frac{\mu}{2} \|\mathbf{x}^* - \mathbf{x}\|^2 \leq f(\mathbf{x}) + \frac{\mu}{2} \|\mathbf{x} - \mathbf{x}^*\|^2$. This means that a point \mathbf{x} satisfying $\mu \|\nabla f_\mu(\mathbf{x})\| \leq \epsilon$ is close to a point in distance of $O(\epsilon)$ that is ϵ -stationary.

It is notable that for a non-smooth function $f(\cdot)$, there could exist a sequence of solutions $\{\mathbf{x}_k\}$ such that $\nabla f_\mu(\mathbf{x}_k)$ converges while $\text{dist}(0, \partial f(\mathbf{x}_k))$ may not converge (Drusvyatskiy & Paquette, 2018). A simple example is to consider $\min_{x \in \mathbb{R}} |x|$. As long as $x \neq 0$, $\text{dist}(0, \partial f(x)) = 1 \neq 0$ no matter how close is x to the stationary point 0. To handle such a challenging issue for non-smooth non-convex problems, we will follow existing works (Davis & Drusvyatskiy, 2018a; Drusvyatskiy & Paquette, 2018; Davis & Grimmer, 2017) to prove the near stationarity in terms of $\nabla f_\mu(\mathbf{x})$. In the case when f is smooth, $\mu \|\nabla f_\mu(\mathbf{x})\|$ is closely related to the magnitude of the gradient. In

particular, let us define the projected gradient below, which is used as a criterion for non-convex smooth optimization in the presence of constraints (Ghadimi & Lan, 2016; Allen-Zhu, 2017):

$$G(\mathbf{x}) = \frac{1}{\gamma}(\mathbf{x} - \text{prox}_{\gamma C}(\mathbf{x} - \gamma \nabla f(\mathbf{x}))) \quad (4)$$

Note that the projected gradient becomes the normal gradient when the constraint $\mathbf{x} \in C$ is absent. It was shown that when $f(\cdot)$ is smooth with L -Lipschitz continuous gradient (Drusvyatskiy & Lewis, 2016):

$$(1 - L\gamma)kG(\mathbf{x})k \leq k\nabla f(\mathbf{x})k \leq (1 + L\gamma)kG(\mathbf{x})k; \forall \mathbf{x} \in C \quad (5)$$

Thus, the near stationarity in terms of $\nabla f(\mathbf{x})$ implies the near stationarity in terms of $G(\mathbf{x})$ for a smooth function $f(\cdot)$ for a properly chosen $\gamma > 0$ (e.g., $\gamma = 1/(2L)$). In this work, we define $\tilde{f}(\mathbf{x})$ as the Moreau envelope of $f(\mathbf{x}) + \frac{\gamma}{2}\|\mathbf{x}\|^2$ as in (2) with $f(\mathbf{x})$ replaced by $f(\mathbf{x}) + \frac{\gamma}{2}\|\mathbf{x}\|^2$ and study the convergence in terms of $\nabla \tilde{f}(\mathbf{x})$.

Now, we are ready to state the basic assumptions of the considered problem (1).

Assumption 1. (i) There is a measurable mapping $g: U \rightarrow \mathbb{R}^d$ such that $\mathbb{E}[g(\mathbf{x}; \xi)] \in \partial_F f(\mathbf{x})$ for any $\mathbf{x} \in C$; (ii) for any $\mathbf{x} \in C$, $\mathbb{E}[k g(\mathbf{x}; \xi)k^2] \leq G^2$; (iii) there exists $\epsilon > 0$ such that $f(\mathbf{x}) \geq \min_{\mathbf{z} \in C} f(\mathbf{z}) + \epsilon$ for any $\mathbf{x} \in C$; (iv) the objective function f is ϵ -weakly convex;

Remark: Assumption 1(i), (ii) assume a stochastic subgradient is available for the objective function and its Euclidean norm square is bounded in expectation, which are standard assumptions for non-smooth optimization. Assumption 1(iii) assumes that the objective value with respect to the optimal value is bounded. Assumption 1(iv) assumes weak convexity of the objective function, which is weaker than assuming smoothness. Below, we present some examples of objective functions in machine learning that are weakly convex.

Ex. 1: Smooth Non-Convex Functions. If $f(\cdot)$ is a L -smooth function (i.e., its gradient is L -Lipschitz continuous), then it is L -weakly convex. This will include the objective function for neural networks with a smooth activation function (e.g., the sigmoid function) and a smooth loss function (e.g., softmax loss).

Ex. 2: Convex and Smooth Composition. Consider $f(\mathbf{x}; \mathbf{a}; b) = h(c(\mathbf{x}; \mathbf{a}; b))$ where $h(\cdot): \mathbb{R} \rightarrow \mathbb{R}$ is closed convex and M -Lipschitz continuous, and $c(\mathbf{x}; \mathbf{a}; b): \mathbb{R}^d \rightarrow \mathbb{R}$ is nonlinear smooth mapping with L -Lipschitz continuous gradient. This class of functions has been considered in (Drusvyatskiy & Paquette, 2018) and it was proved that $f(\mathbf{x}; \mathbf{a}; b)$ is ML -weakly convex. An interesting example is phase retrieval, where $c(\mathbf{x}; \mathbf{a}; b) = j(\mathbf{x}^\top \mathbf{a})^2 - bj$. Another example related to deep learning is that if $c(\mathbf{x}; \cdot)$ denotes the mapping function of a neural network parameterized by \mathbf{x} with a smooth activation function and h denotes a non-smooth Lipschitz continuous convex loss function (e.g., hinge loss), then the resulting loss $f(\mathbf{x}; \cdot)$ is a weakly convex function of \mathbf{x} . More examples of this class can be found in (Davis & Drusvyatskiy, 2018a).

4 STAGewise OPTIMIZATION: ALGORITHMS AND CONVERGENCE

In this section, we will present the proposed stagewise algorithms and their convergence results. We will first present a Meta algorithmic framework highlighting the key features of the proposed algorithms and then present several variants of the Meta algorithm.

The Meta algorithmic framework is described in Algorithm 1. There are several key features that differentiate Algorithm 1 from existing stochastic algorithms that come with theoretical guarantee. First, the algorithm is run with multiple stages. At each stage, a basic stochastic algorithm (SA) is called to optimize a regularized problem $f_s(\mathbf{x})$ inexactly that consists of the original objective function and a quadratic term, which is guaranteed to be convex due to the weak convexity of f and $\frac{\gamma}{2}\|\mathbf{x}\|^2$. The convexity of f_s allows one to employ any suitable existing stochastic algorithms (cf. Theorem 1) that have convergence guarantee for convex problems. The returned solution from the $(s-1)$ -th stage is used as a reference point for constructing f_s and as an initial solution for warm-start. It is notable that SA usually returns an averaged solution \mathbf{x}_s at each stage. Second, a decreasing sequence of step size parameters γ_s is used. At each stage, the SA uses a constant step size parameter γ_s and runs the updates for a number of T_s iterations. We do not initialize T_s as it might be adaptive to the data as in stagewise ADAGRAD. The setup of γ_s and T_s will depend on the specific choice of SA, which will be exhibited later for different variants.

Algorithm 1 A Meta Stagewise Algorithm: Stagewise-SA

```

1: Initialize: a sequence of decreasing step size parameters  $f_{sg}, \mathbf{x}_0, \alpha, \beta < 1$ 
2: for  $s = 1; \dots; S$  do
3:   Let  $f_s(\cdot) = f(\cdot) + \frac{1}{2}k \|\mathbf{x}_{s-1}\|^2$ 
4:    $\mathbf{x}_s = \text{SA}(f_s; \mathbf{x}_{s-1}; \alpha; \beta; T_s)$   $\mathbf{x}_s$  is usually an averaged solution
5: end for

```

To illustrate that Algorithm 1 is a universal framework such that any suitable SA algorithm can be employed, we present the following result by assuming that SA has an appropriate convergence for a convex problem.

Theorem 1. Let $f(\cdot)$ be a convex function, $\mathbf{x}_* = \arg \min_{\mathbf{x} \in \mathcal{C}} f(\mathbf{x})$ and α, β denote some problem dependent constants. Suppose for $\mathbf{x}_+ = \text{SA}(f; \mathbf{x}_0; \alpha; \beta; T)$, we have

$$\mathbb{E}[f(\mathbf{x}_+) - f(\mathbf{x}_*)] \leq \alpha_1(\alpha; \beta; T)k\|\mathbf{x}_0 - \mathbf{x}_*\|^2 + \alpha_2(\alpha; \beta; T)(f(\mathbf{x}_0) - f(\mathbf{x}_*)) + \alpha_3(\alpha; \beta; T); \quad (6)$$

Under Assumption 1(i), (iii) and (iv), by running Algorithm 1 with $\alpha = 1/(2S)$, and with $\beta; T_s$ satisfying $\alpha_1(\beta; T_s; \alpha) = 1/(48S)$; $\alpha_2(\beta; T_s; \alpha) = 1/2$; $\alpha_3(\beta; T_s; \alpha) \leq c_3/S$ for some $c_3 > 0$, we have

$$\mathbb{E} \|k\mathbf{r}(\mathbf{x})\|^2 \leq \frac{32}{(S+1)}(\alpha+1) + \frac{48c_3(\alpha+1)}{(S+1)},$$

where \mathbf{r} is randomly selected from $f_0; \dots; f_S$ with probabilities $p_i \propto (\alpha+1)^i$.

Remark: It is notable that the convergence guarantee is provided on a stagewise average solution \mathbf{x}_+ . To justify a heuristic approach that returns the final average solution for prediction, we analyze a new sampling strategy that samples a solution among all stagewise average solutions with sampling probabilities increasing as the stage number increases. This sampling strategy is better than uniform sampling strategy or a strategy with decreasing sampling probabilities in the existing literature. The convergence upper bound in (6) of SA covers the results of a broad family of stochastic convex optimization algorithms. When $\alpha_2(\beta; T_s; \alpha) = 0$ (as in SGD), the upper bound can be improved by a constant factor. Moreover, we do not optimize the value of α . Indeed, any $\alpha < 1/2$ will work, which only has an effect on constant factor in the convergence upper bound.

Next, we present several variants of the Meta algorithm by employing SGD, ADAGRAD, and stochastic momentum methods as the basic SA algorithm, to which we refer as stagewise SGD, stagewise ADAGRAD, and stagewise stochastic momentum methods, respectively. It is worth mentioning that one can follow similar analysis to analyze other stagewise algorithms by using their basic convergence for stochastic convex optimization, including RMSProp (Mukkamala & Hein, 2017), AMSGrad (Reddi et al., 2018), stochastic alternating direction methods of multipliers (ADMM) (Ouyang et al., 2013; Suzuki, 2013).

4.1 STAGewise SGD

In this subsection, we analyze the convergence of stagewise SGD, in which SGD shown in Algorithm 3 in the Appendix is employed in the Meta framework. Besides Assumption 1, we impose the following bounded domain assumption in this subsection.

Assumption 2. There exists $D > 0$ such that $\|\mathbf{x} - \mathbf{y}\| \leq D$ for any $\mathbf{x}, \mathbf{y} \in \mathcal{C}$.

It is worth mentioning that bounded domain assumption is imposed for simplicity, which is usually assumed in convex optimization. For machine learning problems, one usually imposes some bounded norm constraint to achieve a regularization. Nevertheless, the bounded domain assumption is not essential for the proposed algorithm. We present a more involved analysis in subsection 4.3 for unbounded domain $\mathcal{C} = \mathbb{R}^d$. The following is a standard basic convergence result of SGD (Zinkevich, 2003), which clearly satisfies the bound in (6).

Lemma 1. For Algorithm 3, assume that $f(\cdot)$ is convex and $\mathbb{E} \|\mathbf{g}_t\|^2 \leq G^2$, then for any $\mathbf{x} \in \mathcal{C}$

$$\mathbb{E}[f(\mathbf{x}_T) - f(\mathbf{x})] \leq \frac{k\|\mathbf{x} - \mathbf{x}_0\|^2}{2(T+1)} + \frac{G^2}{2}$$

The following theorem exhibits the convergence of stagewise SGD.

Algorithm 2 ADAGRAD($f; \mathbf{x}_0; \cdot; \cdot$)

-
- 1: **Initialize:** $\mathbf{x}_1 = \mathbf{x}_0, \mathbf{g}_{1:0} = [], H_0 \succeq \mathbb{R}^{d \times d}$
 - 2: **while** T does not satisfy the condition in Theorem 3 **do**
 - 3: Compute a stochastic subgradient \mathbf{g}_t for $f(\mathbf{x}_t)$
 - 4: Update $\mathbf{g}_{1:t} = [\mathbf{g}_{1:t-1}; \mathbf{g}(\mathbf{x}_t)], S_{t,i} = k\mathbf{g}_{1:t,i}k_2$
 - 5: Set $H_t = H_0 + \text{diag}(S_t)$ and $\mathbf{P}_t(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}_1)^\top H_t(\mathbf{x} - \mathbf{x}_1)$
 - 6: Let $\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \mathbf{x}^\top \frac{1}{t} \sum_{i=1}^t \mathbf{g}_i + \frac{1}{t} \mathbf{P}_t(\mathbf{x})$
 - 7: **end while**
 - 8: **Output:** $\bar{\mathbf{x}}_T = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t$
-

Theorem 2. Suppose Assumption 1 and 2 hold. By setting $\beta = 1/(2\epsilon)$, $\gamma_s = \epsilon/\beta$; $T_s = 12/\epsilon^2$ where $\epsilon > 0$ is a free parameter, then stagewise SGD (Algorithm 1 employing SGD) guarantees that

$$\mathbb{E} [k\mathbf{r}(\bar{\mathbf{x}})k^2] \leq \frac{16}{S+1} (\beta + 1) + \frac{24}{(S+1)} \hat{G}^2(\beta + 1);$$

where $\hat{G}^2 = 2G^2 + 2\epsilon^{-2}D^2$, and β is similarly defined as in Theorem 1.

Remark: To find a solution with $\mathbb{E} [k\mathbf{r}(\bar{\mathbf{x}})k^2] \leq \epsilon^2$, we can set $S = O(1/\epsilon^2)$ and the total iteration complexity $\sum_{s=1}^S T_s$ is in the order of $O(1/\epsilon^4)$. The above theorem is essentially a corollary of Theorem 1 by applying Lemma 1 to $f_s(\cdot)$ at each stage. We present a complete proof in the appendix.

4.2 STAGewise ADAGRAD

One of the main contributions of the present work is to develop a variant of ADAGRAD with adaptive convergence to data for stochastic non-convex optimization. In this subsection, we analyze stagewise ADAGRAD and establish its adaptive complexity. In particular, we consider the Meta algorithm that employs ADAGRAD in Algorithm 2 to optimize each f_s . The key difference of stagewise ADAGRAD from stagewise SGD is that the number of iterations T_s at each stage is adaptive to the history of learning. It is this adaptiveness that makes the proposed stagewise ADAGRAD achieve adaptive convergence. It is worth noting that such adaptive scheme has been also considered in (Chen et al., 2018b) for solving stochastic strongly convex problems. In contrast, we consider stochastic weakly convex problems. Similar to previous analysis of ADAGRAD (Duchi et al., 2011; Chen et al., 2018b), we assume $k\mathbf{g}(\mathbf{x}; \cdot)k_\infty \leq G; \beta \mathbf{x} \succeq \mathbf{0}$ in this subsection. Note that this is stronger than Assumption 1 (ii). We formally state this assumption required in this subsection below.

Assumption 3. $k\mathbf{g}(\mathbf{x}; \cdot)k_\infty \leq G$ for any $\mathbf{x} \succeq \mathbf{0}$.

The convergence analysis of stagewise ADAGRAD is build on the following lemma, which is attributed to Chen et al. (2018b) with a proof sketch provided in the Appendix.

Lemma 2. Let $f(\mathbf{x})$ be a convex function, $H_0 = G$ with $G = \max_t k\mathbf{g}_t k_\infty$, and iteration number T satisfy $T \geq M \max_i \frac{G + \max_j \|g_{1:T,i}\|}{2c}; c = \frac{d}{T} \sum_{i=1}^d k\mathbf{g}_{1:T,i}\|$ for some $c > 0$. Algorithm 2 returns an averaged solution $\bar{\mathbf{x}}_T$ such that

$$\mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}_*)] \leq \frac{c}{M} k\mathbf{x}_0 - \mathbf{x}_*k^2 + \frac{1}{Mc}; \quad (7)$$

where $\mathbf{x}_* = \arg \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$, $\mathbf{g}_{1:t} = (\mathbf{g}(\mathbf{x}_1); \dots; \mathbf{g}(\mathbf{x}_t))$ and $\mathbf{g}_{1:t,i}$ denotes the i -th row of $\mathbf{g}_{1:t}$.

The convergence property of stagewise ADAGRAD is described by following theorem.

Theorem 3. Suppose Assumption 1(i), (iii), (iv), Assumption 2 and Assumption 3 hold. By setting $\beta = 1/(2\epsilon)$, $\gamma_s = \epsilon/\beta$; $T_s = M_s \max_i (\hat{G} + \max_j k\mathbf{g}_{1:T_s,i}k) = (2c)/c = \frac{d}{T_s} \sum_{i=1}^d k\mathbf{g}_{1:T_s,i}\|$ where $\epsilon; c > 0$ are free parameters, and $M_s \leq 24/c$, then stagewise ADAGRAD guarantees that

$$\mathbb{E}[k\mathbf{r}(\bar{\mathbf{x}})k^2] \leq \frac{16}{S+1} (\beta + 1) + \frac{4}{c^2} \frac{\hat{G}^2(\beta + 1)}{(S+1)};$$

where $\hat{G} = G + \epsilon^{-1}D$, $\mathbf{g}_{1:t,i}^s$ denotes the cumulative stochastic gradient of the i -th coordinate at the s -th stage, and β is similarly defined as in Theorem 1.

Remark: Note that the free parameter c is introduced to balance the two terms in the lower bound of T_s due to that $(\hat{G} + \max_i kg_{1:T_s,i}^s) / \hat{G}^D T_s$ and $\sum_{i=1}^d kg_{1:T_s,i}^s / d\hat{G}^D T_s$ have different orders when d is very large. One way to balance these two terms is to set c in the order of $\frac{1}{1+d}$, resulting $O(d=(S+1))$ for the second term in the above convergence bound. Another way is to choose c in the s -th stage such that the two terms in the max of the lower bound of T_s match each other. One can derive a similar order of $O(d=(S+1))$ for the second term in the above convergence bound. It is obvious that the total number of iterations $\sum_{s=1}^S T_s$ is adaptive to the data. Next, let us present more discussion on the iteration complexity. Note that $M_s = O(\frac{1}{s})$ by setting c as a constant. By the boundness of stochastic gradient $kg_{1:T_s,i}^s = O(\frac{1}{T_s})$, therefore T_s in the order of $O(s)$ will satisfy the condition in Theorem 3. Thus, in the worst case the iteration complexity for finding $E[kr(\mathbf{x})k^2]^{-2}$ is in the order of $\sum_{s=1}^S O(s) = O(1=^4)$. To show the potential advantage of adaptive step size as in the convex case, let us consider a good case when the cumulative growth of stochastic gradient is slow, e.g., assuming $kg_{1:T_s,i}^s = O(T_s^{-\alpha})$ with $\alpha < 1=2$. Then $T_s = O(s^{1=(2(1-\alpha))})$ will work, and then the total number of iterations $\sum_{s=1}^S T_s = O(1=^{2+1=(1-\alpha)})$, which is better than $O(1=^4)$.

4.3 STAGewise STOCHASTIC MOMENTUM (SM) METHODS

Finally, we present stagewise stochastic momentum (SM) methods and their convergence results. In the literature, there are two popular variants of stochastic momentum methods, namely, stochastic heavy-ball method (SHB) and stochastic Nesterov’s accelerated gradient method (SNAG). Both methods have been used for training deep neural networks (Krizhevsky et al., 2012; Sutskever et al., 2013), and have been analyzed by (Yang et al., 2016) for non-convex optimization. To contrast with the results in (Yang et al., 2016), we will consider the same unified stochastic momentum methods that subsume SHB, SNAG and SGD as special cases when \mathbb{R}^d . The updates are presented in Algorithm 4 in the Appendix. There are two additional parameters: $\beta \in (0; 1)$ is the momentum parameter and γ is a parameter that can vary between $[0; 1=(1-\beta)]$. By changing the value of β , we can obtain the three variants, SHB ($\beta = 0$), SNAG ($\beta = 1$) and SGD ($\beta = 1=(1-\beta)$). Due to the limit of space, we only present the convergence of stagewise SM methods below.

Theorem 4. *Suppose Assumption 1 holds. By setting $\beta = 1=(2)$, $\gamma_s = (1-\beta) = (96s(\beta+1))$, $T_s = 2304(\beta+1)s$, with γ similarly defined as in Theorem 1 stagewise SM methods guarantee*

$$E[kr(\mathbf{x})k^2] \leq \frac{16(\beta+1)}{S+1} + \frac{G^2(\beta+48(2-\beta)(1-\beta))(S+1)}{96(\beta+1)(1-\beta)(S+1)}.$$

Remark: The bound in the above theorem is in the same order as that in Theorem 2. The total iteration complexity for finding a solution \mathbf{x} with $E[kr(\mathbf{x})k^2]^{-2}$ is $O(1=^4)$ similar to that achieved in (Yang et al., 2016).

5 EXPERIMENTS

In this section, we present some empirical results to verify the effectiveness of the proposed stagewise algorithms. We use two benchmark datasets, namely CIFAR-10 and CIFAR-100 (Krizhevsky et al.) for our experiments. We implement the proposed stagewise algorithms in TensorFlow. We compare different algorithms for learning ResNet-20 (He et al., 2016) with batch normalization (Ioffe & Szegedy, 2015) adopted after each convolution and before ReLU activation.

Baselines. We compare the proposed stagewise algorithms with their variants implemented in TensorFlow. It is notable that ADAGRAD has a step size (aka learning rate) parameter η , which is a constant in theory (Li & Orabona, 2018; Chen et al., 2018a; Zou & Shen, 2018). However, in the deep learning community a heuristic fixed frequency decay scheme for the step size parameter is commonly adopted (Ren et al., 2018; Wilson et al., 2017). We thus compare two implementations of ADAGRAD - one with a constant learning rate parameter and another one with a fixed frequency decay scheme, which are referred to as ADAGRAD (theory) and ADAGRAD (heuristic). For each baseline algorithms of SGD, SHB, SNAG, we also implement two versions - a theory version with iteratively decreasing size $\eta = \frac{\eta_0}{t}$ suggested by previous theories and a heuristic approach with fixed frequency decay scheme used in practice, using (theory) and (heuristic) to indicate them. The fixed

¹note it is not equivalent to the step size in SGD.

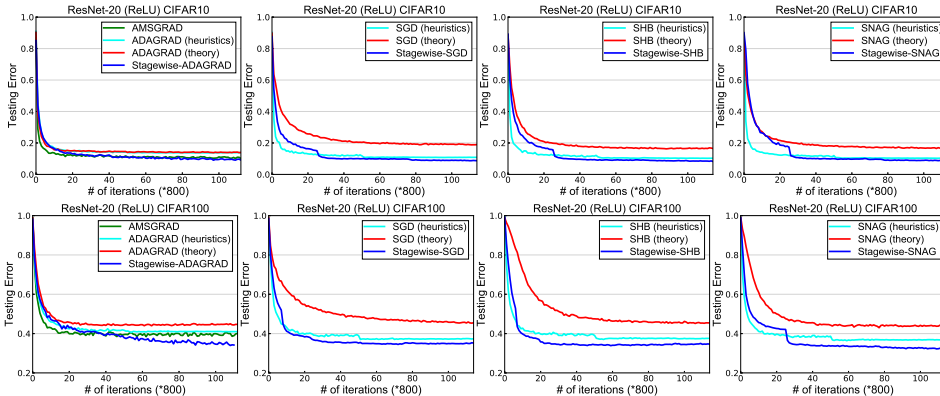


Figure 1: Comparison of Testing Error on CIFAR-10 (top) and CIFAR-100 (bottom).

frequency decay scheme used in the heuristic variants is similar as that in (He et al., 2016), i.e., the step size parameter is decreased by 10 at 40k, 60k iterations. We also compare stagewise ADAGRAD with AMSGrad (Reddi et al., 2018) - a corrected version of Adam.

Parameters. The stagewise step size $\eta_s = \frac{\eta_0}{s}$ is used in stagewise ADAGRAD, the number of iterations T_s in stagewise ADAGRAD is set according to Theorem 3 with some simplifications for dealing with unknown \hat{G} , in particular we set T_s to the smallest value larger than $T_0 \max_i k g_{1:T_s;i}^s$. For stagewise SGD, SHB, SNAG, the stagewise step size and iteration number is set to $\eta_s = \frac{\eta_0}{s}$ and $T_s = T_0 s$, respectively. The involved parameters in the compared algorithms are tuned for the best performance, including the initial step size parameter η_0 of all algorithms, the value of T_0 and η_0 for our stagewise algorithms.

Results. We consider two settings - with/without an ℓ_2 norm regularization on weights. For comparison, we evaluate the training error and testing error of obtained solutions in the process of training. For our stagewise algorithms, the evaluation is done based on the current averaged solution, and for other baselines the evaluation is done based on the current solution. Due to the limit of space, we only show the results of testing error on CIFAR-10 and CIFAR-100 for the setting without regularization in Figure 1. All results are included in the Appendix. From all results, we have several observations. (i) The proposed stagewise algorithms perform much better in terms of testing error than the existing theoretical versions reported in the literature (marked with theory in the legend). This indicates the proposed stagewise step size scheme is better than iteratively decreasing step size scheme. (ii) The proposed stagewise algorithms achieve similar and sometimes even better testing error than the heuristic approaches with a fixed frequency decay scheme used in practice. However, the heuristic approaches usually give smaller training error. This seems to indicate that the proposed algorithms are less vulnerable to the overfitting. In another word, the proposed algorithms have smaller generalization error, i.e., the difference between the testing error and the training error.

6 CONCLUSION & FUTURE WORK

In this paper, we have proposed a universal stagewise learning framework for solving stochastic non-convex optimization problems, which employs well-known tricks in practice that have not been well analyzed theoretically. We provided theoretical convergence results for the proposed algorithms for non-smooth non-convex optimization problems. We also established an adaptive convergence of a stochastic algorithm using data adaptive coordinate-wise step size of ADAGRAD, and exhibited its faster convergence than non-adaptive stepsize when the growth of cumulative stochastic gradients is slow similar to that in the convex case. For future work, one may consider developing more variants of the proposed meta algorithm, e.g., stagewise AMSGrad, stagewise RMSProp, etc. We will also consider the empirical studies on the large-scale ImageNet data set.

REFERENCES

Marín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew

- Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Zeyuan Allen-Zhu. Natasha: Faster non-convex stochastic optimization via strongly non-convex parameter. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pp. 89–97, 2017.
- Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of International Conference on Computational Statistics (COMPSTAT)*, pp. 177–187, 2010.
- Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Accelerated methods for non-convex optimization. *CoRR*, abs/1611.00756, 2016.
- Xiangyi Chen, Sijia Liu, Ruoyu Sun, and Mingyi Hong. On the convergence of a class of adam-type algorithms for non-convex optimization. *CoRR*, abs/1808.02941, 2018a.
- Zaiyi Chen, Yi Xu, Enhong Chen, and Tianbao Yang. Sadagrad: Strongly adaptive stochastic gradient methods. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018b.
- Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *CoRR*, abs/1803.06523, 2018a.
- Damek Davis and Dmitriy Drusvyatskiy. Stochastic subgradient method converges at the rate $o(k^{-1/4})$ on weakly convex functions. *CoRR*, abs/1802.02988, 2018b.
- Damek Davis and Benjamin Grimmer. Proximally guided stochastic subgradient method for nonsmooth, nonconvex problems. *arXiv preprint arXiv:1707.03505*, 2017.
- Jeffrey Dean, Greg S. Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Quoc V. Le, Mark Z. Mao, Marc’Aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, and Andrew Y. Ng. Large scale distributed deep networks. In *NIPS*, pp. 1223–1231, USA, 2012. Curran Associates Inc.
- D. Drusvyatskiy and C. Paquette. Efficiency of minimizing compositions of convex functions and smooth maps. *Mathematical Programming*, Jul 2018.
- Dmitriy Drusvyatskiy and Adrian S. Lewis. Error bounds, quadratic growth, and linear convergence of proximal methods. *arXiv:1602.06661*, 2016.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- Saeed Ghadimi and Guanghui Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Saeed Ghadimi and Guanghui Lan. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Math. Program.*, 156(1-2):59–99, 2016.
- Moritz Hardt, Ben Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pp. 1225–1234, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning, (ICML)*, pp. 448–456, 2015.

- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research).
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1106–1114, 2012.
- Guanghui Lan and Yu Yang. Accelerated stochastic algorithms for nonconvex finite-sum and multi-block optimization. *CoRR*, abs/1805.05411, 2018.
- Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. *CoRR*, abs/1805.08114, 2018.
- Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with restarts. In *ICLR*, volume abs/1608.03983, 2017.
- Mahesh Chandra Makkamala and Matthias Hein. Variants of RMSProp and Adagrad with logarithmic regret bounds. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2545–2553, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- Hua Ouyang, Niao He, Long Tran, and Alexander Gray. Stochastic alternating direction method of multipliers. In *International Conference on Machine Learning*, pp. 80–88, 2013.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018.
- Mengye Ren, Wenyan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. *arXiv preprint arXiv:1803.09050*, 2018.
- R. Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14:877–898, 1976.
- R.T. Rockafellar. *Convex Analysis*. Princeton mathematical series. Princeton University Press, 1970.
- Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pp. 1139–1147, 2013.
- Taiji Suzuki. Dual averaging and proximal gradient descent for online alternating direction multiplier method. In *International Conference on Machine Learning*, pp. 392–400, 2013.
- Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: Sharp convergence over nonconvex landscapes. *CoRR*, abs/1806.01811, 2018.
- Ashia C. Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *NIPS*, pp. 4151–4161, 2017.
- Yi Xu, Qihang Lin, and Tianbao Yang. Stochastic convex optimization: Faster local growth implies faster global convergence. In *ICML*, pp. 3821 – 3830, 2017.
- Tianbao Yang, Qihang Lin, and Zhe Li. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. volume abs/1604.03257, 2016.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pp. 928–936, 2003.

Fangyu Zou and Li Shen. On the convergence of adagrad with momentum for training deep neural networks. *CoRR*, abs/1808.03408, 2018.

Algorithm 3 SGD($f; \mathbf{x}_0; \cdot; T$)

```

for  $t = 0; \dots; T$  do
  Compute a stochastic subgradient  $\mathbf{g}_t$  for  $f(\mathbf{x}_t)$ .
   $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{g}_t$ 
end for
Output:  $\bar{\mathbf{x}}_T = \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{x}_t$ 

```

Algorithm 4 Unified Stochastic Momentum Methods: SUM($f; \mathbf{x}_0; \cdot; T$)

```

Set parameters:  $\eta, \beta \in (0, 1)$ .
for  $t = 0; \dots; T$  do
  Compute a stochastic subgradient  $\mathbf{g}_t$  for  $f(\mathbf{x}_t)$ .
   $\mathbf{y}_{t+1} = \mathbf{x}_t - \eta \mathbf{g}_t$ 
   $\mathbf{v}_{t+1} = \beta \mathbf{v}_t + (1 - \beta) \mathbf{g}_t$ 
   $\mathbf{x}_{t+1} = \mathbf{y}_{t+1} + \mathbf{v}_{t+1}$ 
end for
Output:  $\bar{\mathbf{x}}_T = \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{x}_t$ 

```

A MORE EXPERIMENTAL RESULTS

In this section, we present more experimental results. Comparison of training and testing error in the two settings (w/o regularization) on the two data sets are plotted in Figure 2, 3, 4, 5. We also report the final testing error (after running 80k iterations) of different algorithms in the two settings on the two datasets in Table 1. For parameter tuning, the initial step sizes of all algorithms are tuned in $\{0.1; 0.3; 0.5; 0.7; 0.9\}g$. The value of β of stagewise algorithms is tuned in $\{1; 10; 100; 500; 1000; 1500; 2000; 3000\}g$. The initial value T_0 for stagewise SGD, SHB, SNAG is tuned in $\{10; 100; 1k; 5k; 6k; 7k; 10k; 20k\}g$, and that for stagewise ADAGRAD is tuned in $\{1; 10; 15; 20; 25; 50; 100\}g$.

B PROOF OF THEOREM 1

Proof. Below, we use E_S to denote expectation over randomness in the S -th stage given all history before S -th stage. Define

$$\mathbf{z}_S = \arg \min_{\mathbf{x} \in \mathcal{C}} f_S(\mathbf{x}) = \text{prox}_{\mathcal{C}}(\cdot; \mathbf{x}_{S-1}) \quad (8)$$

Then $r_S(\mathbf{x}_{S-1}) = \|\mathbf{x}_{S-1} - \mathbf{z}_S\|^2$. By applying the convergence bound of SA to $f_S(\mathbf{x})$, we have

$$E_S[f_S(\mathbf{x}_S) - f_S(\mathbf{z}_S)] \leq \underbrace{\frac{1}{2} \left(\frac{\eta}{T_S} \right) k \|\mathbf{x}_{S-1} - \mathbf{z}_S\|^2 + \frac{1}{2} \left(\frac{\eta}{T_S} \right) (f_S(\mathbf{x}_{S-1}) - f_S(\mathbf{z}_S))}_{\varepsilon_S} + \frac{3}{2} \left(\frac{\eta}{T_S} \right) \varepsilon_S$$

It then follows that

$$E_S \|\mathbf{x}_S - \mathbf{z}_S\|^2 + \frac{1}{2} k \|\mathbf{x}_S - \mathbf{x}_{S-1}\|^2 \leq f_S(\mathbf{z}_S) + E_S \|\mathbf{x}_{S-1} - \mathbf{z}_S\|^2 + E_S \|\mathbf{x}_{S-1} - \mathbf{z}_S\|^2 \quad (9)$$

On the other hand, we have that

$$\begin{aligned} k \|\mathbf{x}_S - \mathbf{x}_{S-1}\|^2 &= k \|\mathbf{x}_S - \mathbf{z}_S + \mathbf{z}_S - \mathbf{x}_{S-1}\|^2 \\ &= k \|\mathbf{x}_S - \mathbf{z}_S\|^2 + k \|\mathbf{z}_S - \mathbf{x}_{S-1}\|^2 + 2\langle \mathbf{x}_S - \mathbf{z}_S, \mathbf{z}_S - \mathbf{x}_{S-1} \rangle \\ &= (1 - \beta) k \|\mathbf{x}_S - \mathbf{z}_S\|^2 + (1 + \beta) k \|\mathbf{x}_{S-1} - \mathbf{z}_S\|^2 \end{aligned}$$

Table 1: Comparison of Final Testing Error (%) on CIFAR-10 and CIFAR-100 Datasets

Algorithms	CIFAR-10		CIFAR-100	
	with reg.	without reg.	with reg.	without reg.
SGD (theory)	16.25	19.18	43.51	45.78
SGD (heuristic)	8.34	10.81	33.67	37.19
Stagewise-SGD	8.34	9.01	32.25	34.95
SHB (theory)	15.67	16.55	39.15	46.23
SHB (heuristic)	8.58	10.28	33.30	37.56
Stagewise-SHB	8.30	8.61	32.85	34.49
SNAG (theory)	17.64	16.76	39.34	44.21
SNAG (heuristic)	8.85	10.34	33.89	36.84
Stagewise-SNAG	8.00	8.93	31.42	33.29
AMSGrad	10.76	11.13	38.62	39.96
AdaGrad (theory)	12.11	13.96	39.09	44.49
AdaGrad (heuristic)	10.71	13.80	37.04	41.06
Stagewise-AdaGrad	9.09	9.51	33.95	34.62

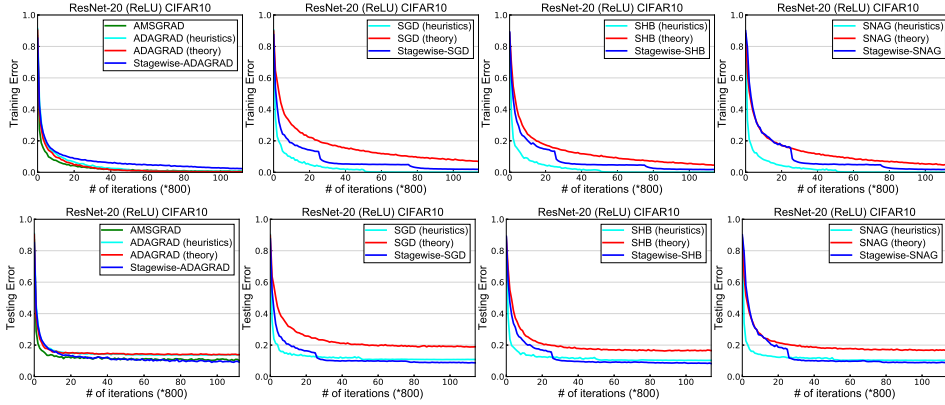
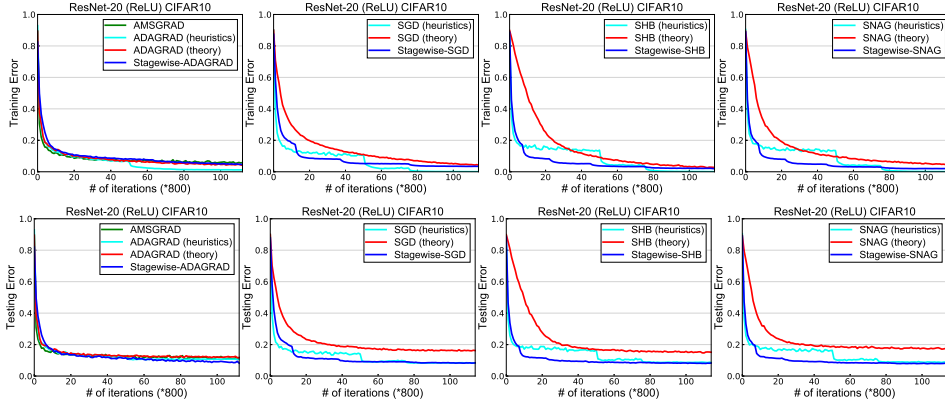


Figure 2: Comparison of Training Error (Top) and Testing Error (bottom) on CIFAR-10 without Regularization.

Figure 3: Comparison of Training Error (Top) and Testing Error (bottom) on CIFAR-10 with Regularization. The regularization parameter is set $5e^{-4}$.

where the inequality follows from the Young's inequality with $0 < s < 1$. Combining the above inequality with (9) we have that

$$\begin{aligned}
& E_s \frac{(1-s)}{2} k \mathbf{x}_{s-1} - \mathbf{z}_s k^2 \leq E_s \left(\mathbf{x}_{s-1} - \left(\mathbf{x}_s + \frac{(s^{-1}-1)}{2} k \mathbf{x}_s - \mathbf{z}_s k^2 + E_s \right) \right) \\
& E_s \left(\mathbf{x}_{s-1} - \left(\mathbf{x}_s + \frac{(s^{-1}-1)}{(-1)} (f_s(\mathbf{x}_s) - f_s(\mathbf{z}_s)) + E_s \right) \right) \\
& E_s \left(\mathbf{x}_{s-1} - \left(\mathbf{x}_s + \frac{-1}{(1)} E_s - E_s \left(\mathbf{x}_{s-1} - \left(\mathbf{x}_s \right) \right) \right) \right) \\
& + E_s \frac{-1}{(1-s)} f''_1(s; T_s) k \mathbf{x}_{s-1} - \mathbf{z}_s k^2 + \frac{1}{2} (s; T_s) (f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s)) + \frac{1}{3} (s; T_s) g ;
\end{aligned} \tag{10}$$

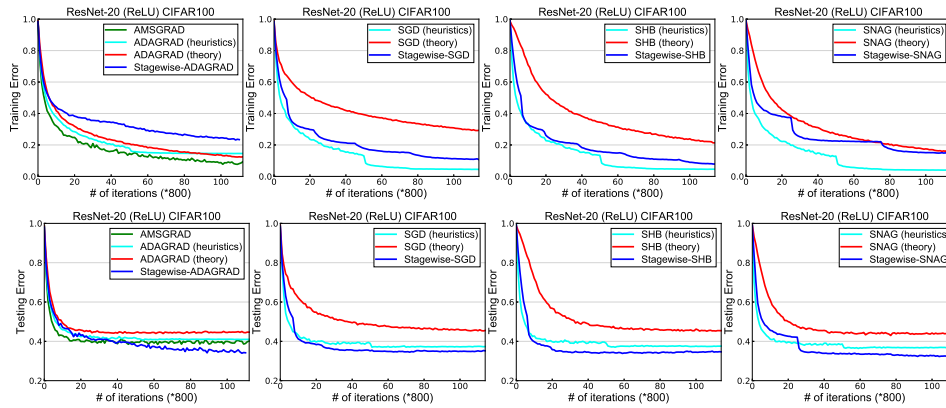


Figure 4: Comparison of Training Error (Top) and Testing Error (bottom) on CIFAR-100 without Regularization.

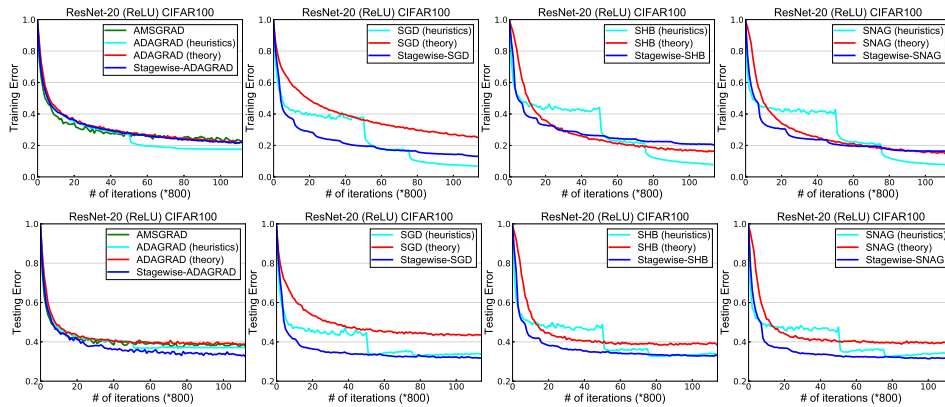


Figure 5: Comparison of Training Error (Top) and Testing Error (bottom) on CIFAR-100 with Regularization. The regularization parameter is set $5e^{-4}$.

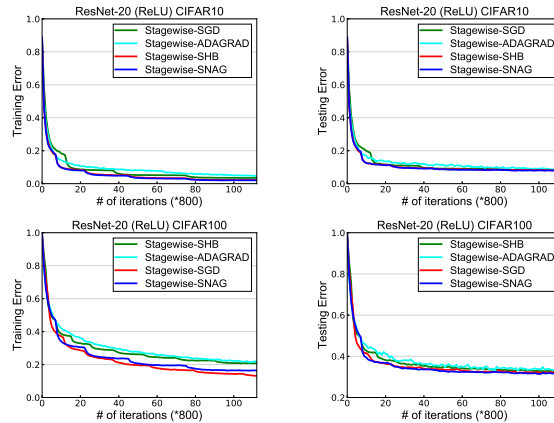


Figure 6: Comparison of different stagewise algorithms in terms of Training Error and Testing Error on CIFAR-10 (top) and CIFAR-100 (bottom) with regularization. The regularization parameter is set $5e^{-4}$.

where the second inequality uses the strong convexity of $f_S(x)$, whose strong convexity parameter is μ_S . Next, we bound $f_S(x_{S-1}) - f_S(z_S)$ given that x_{S-1} is fixed. According to the definition

of $f_s(\cdot)$, we have

$$\begin{aligned} f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s) &= (\mathbf{x}_{s-1}) - (\mathbf{z}_s) - \frac{1}{2}k\mathbf{z}_s - \mathbf{x}_{s-1}k^2 \\ &= (\mathbf{x}_{s-1}) - (\mathbf{x}_s) + (\mathbf{x}_s) - (\mathbf{z}_s) - \frac{1}{2}k\mathbf{z}_s - \mathbf{x}_{s-1}k^2 \\ &= [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] + f_s(\mathbf{x}_s) - f_s(\mathbf{z}_s) + \frac{1}{2}k\mathbf{z}_s - \mathbf{x}_{s-1}k^2 - \frac{1}{2}k\mathbf{x}_s - \mathbf{x}_{s-1}k^2 - \frac{1}{2}k\mathbf{z}_s - \mathbf{x}_{s-1}k^2 \\ &\quad [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] + [f_s(\mathbf{x}_s) - f_s(\mathbf{z}_s)]: \end{aligned}$$

Taking expectation over randomness in the s -th stage on both sides, we have

$$\begin{aligned} f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s) &= \mathbb{E}_s [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] + \mathbb{E}_s [f_s(\mathbf{x}_s) - f_s(\mathbf{z}_s)] \\ &\quad \mathbb{E} [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] + \mu_1(s; T_{s;})k\mathbf{x}_{s-1} - \mathbf{z}_sk^2 + \mu_2(s; T_{s;}) (f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s)) + \mu_3(s; T_{s;}): \end{aligned}$$

Thus,

$$(1 - \mu_2(s; T_{s;})) (f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s)) = \mathbb{E} [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] + \mu_1(s; T_{s;})k\mathbf{x}_{s-1} - \mathbf{z}_sk^2 + \mu_3(s; T_{s;}):$$

Assuming that $\mu_2(s; T_{s;}) = 1/2$, we have

$$\mu_2(s; T_{s;}) (f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s)) = \mathbb{E}_s [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] + \mu_1(s; T_{s;})k\mathbf{x}_{s-1} - \mathbf{z}_sk^2 + \mu_3(s; T_{s;}):$$

Plugging this upper bound into (10), we have

$$\begin{aligned} \mathbb{E}_s \left(\frac{1-s}{2} \right) k\mathbf{x}_{s-1} - \mathbf{z}_sk^2 &= \mathbb{E}_s [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] \\ &\quad + \mathbb{E}_s \left(\frac{s}{1-s} \right) \mu_1(s; T_{s;})k\mathbf{x}_{s-1} - \mathbf{z}_sk^2 + (\mathbf{x}_{s-1}) - (\mathbf{x}_s) + 2\mu_3(s; T_{s;})g \quad (11) \end{aligned}$$

By setting $s = 1/2$; $\mu_1(s; T_{s;}) = 1/48$ and assuming $\mu_3(s; T_{s;}) = 1/48$, we have

$$\mathbb{E}_s \frac{1}{8} k\mathbf{x}_{s-1} - \mathbf{z}_sk^2 = 4\mathbb{E}_s [(\mathbf{x}_{s-1}) - (\mathbf{x}_s)] + 6\mu_3(s; T_{s;})g$$

Define $w_s = 8$. Multiplying both sides by w_s , we have that

$$w_s \mathbb{E}_s [kr - (\mathbf{x}_{s-1})k^2] = \mathbb{E}_s [32w_s - s + 48\mu_3(s; T_{s;})w_s]$$

By summing over $s = 1; \dots; S + 1$, we have

$$\sum_{s=1}^{S+1} w_s \mathbb{E} [kr - (\mathbf{x}_{s-1})k^2] = \mathbb{E} \left[\frac{32}{S+1} \sum_{s=1}^{S+1} w_s - s + \frac{48}{S+1} \sum_{s=1}^{S+1} w_s \mu_3(s; T_{s;}) \right]$$

Taking the expectation w.r.t. $\mathcal{F}_0; \dots; \mathcal{F}_S$, we have that

$$\mathbb{E} [kr - (\mathbf{x})k^2] = \mathbb{E} \left[\frac{32}{\prod_{s=1}^{S+1} w_s} \sum_{s=1}^{S+1} w_s - s + \frac{48}{\prod_{s=1}^{S+1} w_s} \sum_{s=1}^{S+1} w_s \mu_3(s; T_{s;}) \right]$$

For the first term on the R.H.S, we have that

$$\begin{aligned} \sum_{s=1}^{S+1} \frac{w_s}{s} &= \sum_{s=1}^{S+1} \frac{w_s}{s} (\mathbf{x}_{s-1}) - (\mathbf{x}_s) = \sum_{s=1}^{S+1} \frac{w_{s-1}}{s-1} (\mathbf{x}_{s-1}) - \frac{w_s}{s} (\mathbf{x}_s) + \sum_{s=1}^{S+1} \frac{w_s}{s} (\mathbf{x}_{s-1}) \\ &= w_0 (\mathbf{x}_0) - w_{S+1} (\mathbf{x}_{S+1}) + \sum_{s=1}^{S+1} (w_s - w_{s-1}) (\mathbf{x}_{s-1}) \\ &= \sum_{s=1}^{S+1} (w_s - w_{s-1}) (\mathbf{x}_{s-1}) - (\mathbf{x}_{S+1}) + \sum_{s=1}^{S+1} (w_s - w_{s-1}) = w_{S+1} \end{aligned}$$

Then,

$$E[kr(\mathbf{x})k^2] = \frac{32}{S+1} \frac{w_{S+1}}{w_s} + \frac{48}{S+1} \frac{w_s}{w_s} \mathbb{P}^{S+1}(\cdot; T_s; \cdot)$$

The standard calculus tells that

$$\int_0^z x dx = \frac{1}{2} z^2$$

$$\int_0^z x^{-1} dx = \ln z; \quad \int_0^z x^{-2} dx = \frac{1}{1} z^{-1}$$

Combining these facts and the assumption $\mathbb{P}^{S+1}(\cdot; T_s; \cdot) \leq c$, we have that

$$E[kr(\mathbf{x})k^2] \leq \frac{32}{S+1} c + \frac{48c}{S+1} = \frac{80c}{S+1}$$

$$\leq \frac{80c}{S+1} < c$$

In order to have $E[kr(\mathbf{x})k^2] \leq \epsilon$, we can set $S = O(1/\epsilon)$. The total number of iterations is

$$\sum_{s=1}^S T_s = \sum_{s=1}^S 12s = 6S(S+1) = O(1/\epsilon^2)$$

□

C PROOF OF THEOREM 2

Proof. The proof is almost a duplicate to that of Theorem 1. Define $w_s = 1/s$. We apply Lemma 1 to each call of SGD in stagewise SGD,

$$E[f_s(\mathbf{x}_s) - f_s(\mathbf{z}_s)] \leq \frac{k\mathbf{z}_s - \mathbf{x}_{s-1}k^2}{2sT_s} + \frac{s\hat{G}^2}{2\epsilon_s}$$

where \hat{G}^2 is the upper bound of $E[kg(\mathbf{x}; \cdot) + g(\mathbf{x}; \cdot) - g(\mathbf{x}_{s-1}; \cdot)]^2$, which exists and can be set to $2G^2 + 2L^2D^2$ due to the Assumption 1-(ii) and the bounded assumption of the domain. Then following the same analysis as that in the proof of Theorem 1, we have

$$E_s \left[\frac{(1-s)}{2} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 \right] \leq E_s \left[\frac{(1-s)}{2} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 + E_s \right]$$

$$E \left[\frac{(1-s)}{2} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 \right] \leq E \left[\frac{(1-s)}{2} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 + E_s \right] + E_s$$

$$E \left[\frac{(1-s)}{2} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 \right] \leq E \left[\frac{(1-s)}{2} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 + E_s \right] \quad (12)$$

Combining the above inequalities, we have that

$$(1-s) \frac{2(1-s)}{(1-s)sT_s} E_s[kr(\mathbf{x}_{s-1})k^2] \leq E_s \left[2s + \frac{(1-s)}{(1-s)} s\hat{G}^2 \right]$$

Multiplying both sides by w_s , we have that

$$w_s (1-s) \frac{2(1-s)}{(1-s)sT_s} E_s[kr(\mathbf{x}_{s-1})k^2] \leq E_s \left[2w_s s + \frac{(1-s)}{(1-s)} w_s s\hat{G}^2 \right]$$

By setting $s = 1/2$ and $T_s = 12$, we have

$$\frac{1}{4} w_s E_s[kr(\mathbf{x}_{s-1})k^2] \leq E_s[2w_s s + 3w_s s\hat{G}^2]$$

By summing over $s = 1; \dots; S + 1$, we have

$$\sum_{s=1}^{S+1} w_s E[kr(\mathbf{x}_{s-1})k^2] \leq E \sum_{s=1}^{S+1} 16 w_s s + 24 \sum_{s=1}^{S+1} w_s s \hat{G}^2$$

Taking the expectation w.r.t. $\mathcal{Z} = \{f_0; \dots; S\}$, we have that

$$E[kr(\mathbf{x})k^2] \leq E \left[\frac{16 \sum_{s=1}^{S+1} w_s s}{\sum_{s=1}^{S+1} w_s} + \frac{24 \sum_{s=1}^{S+1} w_s s \hat{G}^2}{\sum_{s=1}^{S+1} w_s} \right]$$

By similar analysis, we have that

$$\begin{aligned} E[kr(\mathbf{x})k^2] &\leq \frac{16}{S+1} (\hat{G}^2 + 1) + \frac{24}{S+1} \hat{G}^2 (\hat{G}^2 + 1) \\ &\leq \frac{16}{S+1} (\hat{G}^2 + 1) + \frac{24}{(S+1)} \hat{G}^2 (\hat{G}^2 + 1) \quad 0 < \hat{G} < 1 \end{aligned}$$

In order to have $E[kr(\mathbf{x})k^2] \leq \epsilon$, we can set $S = O(1/\epsilon^2)$. The total number of iterations is

$$\sum_{s=1}^S T_s \leq \sum_{s=1}^S 12 s \leq 6 S(S+1) = O(1/\epsilon^4)$$

□

D PROOF OF THEOREM 4

We need the following lemma for the convergence bound of stochastic momentum methods for a strongly convex problem, whose proof is postponed to Section F.

Lemma 3. For Algorithm 4, assume $f(\mathbf{x}) = \frac{1}{2} k\mathbf{x} - \mathbf{x}_0k^2$ is a μ -strongly convex function, $\mathbf{g}_t = \mathbf{g}(\mathbf{x}_t) + \frac{1}{2}(\mathbf{x}_t - \mathbf{x}_0)$ where $\mathbf{g}(\mathbf{x}; \cdot) \geq \mathcal{G}_F(\mathbf{x}_t)$ such that $E[k\mathbf{g}(\mathbf{x}; \cdot)k^2] \leq \hat{G}^2$, and $(1 - \mu)^2 \geq (8 + 4)\epsilon$, then we have that

$$\begin{aligned} E[f(\mathbf{x}_T) - f(\mathbf{x}_*)] &\leq \frac{(1 - \mu)k\mathbf{x}_0 - \mathbf{x}_*k^2}{2(T+1)} + \frac{(f(\mathbf{x}_0) - f(\mathbf{x}_*))}{(1 - \mu)(T+1)} + \frac{2\hat{G}^2(2 - \mu)}{1} + \frac{4 - \mu}{(1 - \mu)^2} k\mathbf{x}_0 - \mathbf{x}_*k^2 \quad (13) \end{aligned}$$

where $\mathbf{x}_T = \sum_{t=0}^T \mathbf{x}_t = (1 + T)\mathbf{x}_*$ and $\mathbf{x}_* \in \arg \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$.

Remark: It is notable that in the above result, we do not use the bounded domain assumption since we consider \mathbb{R}^d for the unified momentum methods in this subsection. The key to get rid of bounded domain assumption is by exploring the strong convexity of $f(\mathbf{x}) = \frac{1}{2} k\mathbf{x} - \mathbf{x}_0k^2$.

Proof. of Theorem 4 According to the definition of \mathbf{z}_s in (8) and Lemma 3, we have that

$$\begin{aligned} E_s \left(\mathbf{x}_s \right) + \frac{1}{2} k\mathbf{x}_s - \mathbf{x}_{s-1}k^2 &\leq f_s(\mathbf{z}_s) + \frac{(f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s))}{(1 - \mu)(T_s + 1)} + \frac{(1 - \mu)k\mathbf{x}_{s-1} - \mathbf{z}_s k^2}{2 - \mu(T_s + 1)} + \frac{2\hat{G}^2(2 - \mu)}{1} + \frac{1}{24} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 \\ &\leq \underbrace{\left(\frac{(1 - \mu)k\mathbf{x}_{s-1} - \mathbf{z}_s k^2}{2 - \mu(T_s + 1)} + \frac{2\hat{G}^2(2 - \mu)}{1} + \frac{1}{24} k\mathbf{x}_{s-1} - \mathbf{z}_s k^2 \right)}_{\mathcal{E}_s} \\ &\quad + E_s \left(\mathbf{x}_{s-1} \right) \end{aligned}$$

where the last inequality uses the value of $\mu = (1 - \mu) = (96(2 - \mu))^{-1}$, $(1 - \mu) = (96(2 - \mu))^{-1}$, which also satisfies the condition in Lemma 3 by noting that $\mu = (1 - \mu) = 1/(2 - \mu)$. Similar to the proof of Theorem 1, we have

$$\frac{(1 - \mu)k\mathbf{x}_{s-1} - \mathbf{z}_s k^2}{2} \leq E_s \left(\mathbf{x}_{s-1} - \mathbf{z}_s \right) + \frac{1}{(1 - \mu)} E_s \quad (14)$$

Plugging the expression of E_s and rearranging above inequality, we have that

$$(1-s) \frac{2 \binom{-1}{s} (1)}{(1) s(T_s+1)} \frac{-1}{(1)} \frac{1}{24} k r (\mathbf{x}_{s-1}) k^2$$

$$2E_s[\] + \frac{2 \binom{-1}{s}}{(1)} \frac{(f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s))}{(1)(T_s+1)} + \frac{2_s G^2 (2 + 1)}{1}$$

The definition of f_s gives that

$$f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s) = (\mathbf{x}_{s-1}) - (\mathbf{z}_s) \frac{1}{2} k \mathbf{z}_s - \mathbf{x}_{s-1} k^2$$

On the other hand, the μ -weakly convexity of g gives that

$$(\mathbf{z}_s) - (\mathbf{x}_{s-1}) + \mu g(\mathbf{x}_{s-1}); \mathbf{z}_s - \mathbf{x}_{s-1} \geq \frac{\mu}{2} k \mathbf{z}_s - \mathbf{x}_{s-1} k^2;$$

where $g(\mathbf{x}_{s-1}) \geq \mu_F (\mathbf{x}_{s-1})$. Combing these two inequalities we have that

$$f_s(\mathbf{x}_{s-1}) - f_s(\mathbf{z}_s) \geq \mu g(\mathbf{x}_{s-1}); \mathbf{x}_{s-1} - \mathbf{z}_s \geq \frac{\mu}{2} k \mathbf{z}_s - \mathbf{x}_{s-1} k^2$$

$$\frac{G^2}{2} + \frac{\mu}{2} k \mathbf{z}_s - \mathbf{x}_{s-1} k^2 = \frac{G^2}{2}$$

where the second inequality follows from Jensen's inequality for k and Young's inequality. Combining above inequalities and multiplying both side by w_s , we have that

$$w_s (1-s) \frac{2 \binom{-1}{s} (1)}{(1) s(T_s+1)} \frac{-1}{(1)} \frac{1}{24} k r (\mathbf{x}_{s-1}) k^2$$

$$2w_s E_s[\] + \frac{2w_s \binom{-1}{s}}{(1)} \frac{G^2}{2 (1)(T_s+1)} + \frac{2_s G^2 (2 + 1)}{1} \quad (15)$$

By setting $s = 1, 2, \dots, S+1$, we have that

$$\frac{w_s}{4} k r (\mathbf{x}_{s-1}) k^2 \geq 2w_s E_s[\] + \frac{w_s s G^2}{4(1-s)^2} + \frac{12w_s s G^2 (2 + 1)}{1}$$

Summing over $s = 1, \dots, S+1$ and rearranging, we have

$$\sum_{s=1}^{S+1} w_s k r (\mathbf{x}_{s-1}) k^2 = E \sum_{s=1}^{S+1} \frac{8}{w_s s} + \frac{w_s s G^2 (+ 48(2 + 1)(1 - s))}{(1-s)^2}$$

Following similar analysis as in the proof of Theorem 2, we can finish the proof. \square

E PROOF OF THEOREM 3

Proof. Applying Lemma 2 with $T_s = M_s \max_i \frac{\hat{G} + \max_i \|g_{1:T_s,i}^s\|}{2c}; c \prod_{i=1}^d k g_{1:T_s,i}^s \| k g M_s > 0$, and the fact that $(\mathbf{x}_{s-1}) - (\mathbf{z}_s) + \frac{1}{2} k \mathbf{x}_{s-1} - \mathbf{z}_s k^2$ in s th stage, we have that

$$E_s (\mathbf{x}_s) + \frac{1}{2} k \mathbf{x}_{s-1} - \mathbf{z}_s k^2 \geq f_s(\mathbf{z}_s) + \frac{c}{M_s s} k \mathbf{x}_{s-1} - \mathbf{z}_s k^2 + \frac{s}{M_s c}$$

$$(\mathbf{x}_s) + E_s$$

According to (14), we have that

$$\frac{(1-s)}{2} E_s [k \mathbf{x}_{s-1} - \mathbf{z}_s k^2] \geq (\mathbf{x}_{s-1}) - (\mathbf{x}_s) + \frac{\binom{-1}{s} (1)}{2} k \mathbf{x}_{s-1} - \mathbf{z}_s k^2 + E_s$$

$$(\mathbf{x}_{s-1}) - (\mathbf{x}_s) + \frac{-1}{(1-s)} \frac{c}{M_s s} k \mathbf{x}_{s-1} - \mathbf{z}_s k^2 + \frac{s}{M_s c}$$

Rearranging above inequality then multiplying both side by w_s , we have that

$$w_s (1 - \alpha) \frac{2c(\alpha^{-1})}{(1 - \alpha)M_s} k r (\mathbf{x}_{s-1})k^2 \\ 2w_s E_s [- s] + \frac{2w_s s(\alpha^{-1})}{cM_s(1 - \alpha)}$$

By using $M_s \leq 24/c$ and summing over $s = 1; \dots; S + 1$, we have that

$$\sum_{s=1}^{S+1} w_s k r (\mathbf{x}_{s-1})k^2 \leq E \sum_{s=1}^{S+1} \frac{8w_s}{c^2} + \frac{w_s}{c^2}$$

By the definition of α in the theorem, taking expectation of $k r (\mathbf{x}_s)k^2$ w.r.t. $\mathcal{F}_0; \dots; \mathcal{F}_S$ we have that

$$E[k r (\mathbf{x}_s)k^2] = E \left[\frac{8}{c^2} \sum_{s=1}^{S+1} \prod_{i=1}^s w_i + \frac{2}{c^2} \sum_{s=1}^{S+1} \prod_{i=1}^{s-1} w_i \right] \\ = \frac{8}{c^2} \frac{(1 - \alpha)^{S+1}}{(S+1)} + \frac{2}{c^2} \frac{(1 - \alpha)^{S+1}}{(S+1)}$$

where $\mathbb{1}(\alpha < 1)$ is 1 if $\alpha < 1$ and 0 otherwise. □

F PROOF OF LEMMA 3

Proof. Following the analysis in Yang et al. (2016), we directly have the following inequality,

$$E[k\mathbf{x}_{k+1} + \mathbf{p}_{k+1} - \mathbf{x}_*k^2] = \\ = E[k\mathbf{x}_k + \mathbf{p}_k - \mathbf{x}_*k^2] - \frac{2}{1} E[(\mathbf{x}_k - \mathbf{x}_*)^\top @f(\mathbf{x}_k)] - \frac{2}{(1 - \alpha)^2} E[(\mathbf{x}_k - \mathbf{x}_{k-1})^\top @f(\mathbf{x}_k)] \\ - \frac{2}{(1 - \alpha)^2} E[\mathbf{g}_{k-1}^\top @f(\mathbf{x}_k)] + \frac{2}{1} E[k\mathbf{g}_k k^2]$$

We also note that

$$f(\mathbf{x}_k) - f(\mathbf{x}_*) = (\mathbf{x}_k - \mathbf{x}_*)^\top @f(\mathbf{x}_k) - \frac{1}{2} k\mathbf{x}_k - \mathbf{x}_*k^2 \\ f(\mathbf{x}_k) - f(\mathbf{x}_{k-1}) = (\mathbf{x}_k - \mathbf{x}_{k-1})^\top @f(\mathbf{x}_k) - \frac{1}{2} k\mathbf{x}_k - \mathbf{x}_{k-1}k^2 \\ E[\mathbf{g}_{k-1}^\top @f(\mathbf{x}_k)] = \frac{E[k\mathbf{g}_{k-1}k^2 + k@f(\mathbf{x}_k)k^2]}{2} = \frac{1}{2} k\mathbf{x}_{k-1} - \mathbf{x}_0k^2 + \frac{1}{2} k\mathbf{x}_k - \mathbf{x}_0k^2 + 2G^2 \\ E_k[k\mathbf{g}_k k^2] = \frac{2}{2} k\mathbf{x}_k - \mathbf{x}_0k^2 + 2G^2$$

where the first two inequalities are due to the strong convexity of $f(\cdot)$ and the last three inequalities are due to the boundness assumption. Thus

$$E[k\mathbf{x}_{k+1} + \mathbf{p}_{k+1} - \mathbf{x}_*k^2] = E[k\mathbf{x}_k + \mathbf{p}_k - \mathbf{x}_*k^2] - \frac{2}{1} E[(f(\mathbf{x}_k) - f(\mathbf{x}_*))] \\ - \frac{2}{(1 - \alpha)^2} E[(f(\mathbf{x}_k) - f(\mathbf{x}_{k-1}))] + \frac{2}{1} (2 - \alpha) 4G^2 \\ - \frac{1}{1} k\mathbf{x}_k - \mathbf{x}_*k^2 - \frac{1}{(1 - \alpha)^2} k\mathbf{x}_k - \mathbf{x}_{k-1}k^2 \\ + \frac{2}{(1 - \alpha)^2} \frac{2}{2} k\mathbf{x}_{k-1} - \mathbf{x}_0k^2 + \frac{2}{(1 - \alpha)^2} \frac{2}{2} k\mathbf{x}_k - \mathbf{x}_0k^2$$

By summarizing the above inequality over $k = 0; \dots; T$, we have

$$\begin{aligned} & \frac{2}{1} \mathbb{E} \sum_{k=0}^T (f(\mathbf{x}_k) - f(\mathbf{x}_*)) \leq \mathbb{E}[k\mathbf{x}_0 - \mathbf{x}_*k^2] + \frac{2}{(1-\mu)^2} \mathbb{E}[f(\mathbf{x}_0) - f(\mathbf{x}_*)] \\ & + \frac{1}{1-\mu} (2\mu + 1)4G^2(T+1) \\ & + \frac{1}{1-\mu} \sum_{k=0}^T k\mathbf{x}_k - \mathbf{x}_*k^2 + \frac{4}{(1-\mu)^2} \frac{1}{2} \sum_{k=0}^T k\mathbf{x}_{k-1} - \mathbf{x}_*k^2 + \frac{4}{(1-\mu)^2} \frac{1}{2} \sum_{k=0}^T k\mathbf{x}_k - \mathbf{x}_*k^2 \\ & + \frac{4}{(1-\mu)^2} \frac{1}{2} (T+1)k\mathbf{x}_0 - \mathbf{x}_*k^2 \end{aligned}$$

When $(1-\mu)^2 = (8\mu + 4)$, we have

$$\begin{aligned} \mathbb{E} (f(\mathbf{x}_T) - f(\mathbf{x}_*)) & \leq \frac{(1-\mu)k\mathbf{x}_0 - \mathbf{x}_*k^2}{2(T+1)} + \frac{f(\mathbf{x}_0) - f(\mathbf{x}_*)}{(1-\mu)(T+1)} + \frac{1}{1-\mu} (2\mu + 1)2G^2 \\ & + \frac{4}{(1-\mu)^2} \frac{1}{2} k\mathbf{x}_0 - \mathbf{x}_*k^2 \end{aligned}$$

□

G PROOF OF LEMMA 2

The proof is almost a duplicate of the proof of Proposition 1 in Chen et al. (2018b). For completeness, we present a proof here.

Proof. Let $\mathbf{x}_0(\mathbf{x}) \in \mathbb{P}_0$ and $k\mathbf{x}k_H = \frac{D}{\mathbf{x}^\top H \mathbf{x}}$. First, we can see that $\mathbf{x}_{t+1}(\mathbf{x}) \in \mathbb{P}_t(\mathbf{x})$ for any $t \geq 0$. Define $\mathbf{g}_t = \frac{1}{t} \sum_{s=1}^t \mathbf{g}_s$ and $\mathbf{g}_t = (\nabla F(\mathbf{x}_t) - \mathbf{g}_t)^\top (\mathbf{x}_t - \mathbf{x})$. Let \mathbf{g}_t^* be defined by

$$\mathbf{g}_t^*(\mathbf{g}) = \sup_{\mathbf{x} \in \mathbb{P}_t} \mathbf{g}^\top \mathbf{x} - \frac{1}{t} \sum_{s=1}^t \mathbf{g}_s^\top \mathbf{x}$$

Taking the summation of objective gap in all iterations, we have

$$\begin{aligned} & \sum_{t=1}^T (f(\mathbf{x}_t) - f(\mathbf{x})) \leq \sum_{t=1}^T (\nabla F(\mathbf{x}_t) - \mathbf{g}_t)^\top (\mathbf{x}_t - \mathbf{x}) = \sum_{t=1}^T \mathbf{g}_t^\top (\mathbf{x}_t - \mathbf{x}) + \sum_{t=1}^T \frac{1}{t} \sum_{s=1}^t \mathbf{g}_s^\top (\mathbf{x}_t - \mathbf{x}) \\ & = \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t - \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} - \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} + \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} + \sum_{t=1}^T \left(\sum_{s=1}^{t-1} \mathbf{g}_s^\top \mathbf{x}_t - \sum_{s=1}^{t-1} \mathbf{g}_s^\top \mathbf{x} \right) \\ & = \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t + \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t + \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} + \sup_{\mathbf{x} \in \mathbb{P}_T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} - \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} \end{aligned}$$

Note that

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_{T+1} - \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_{T+1} & = \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_{T+1} - \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_{T+1} \\ \sup_{\mathbf{x} \in \mathbb{P}_T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} - \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} & = \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} - \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x} \\ \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_{T+1} - \frac{1}{T} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_{T+1} & = \frac{1}{2} k\mathbf{g}_T k^2_{T+1} \end{aligned}$$

where the last inequality uses the fact that $f_t(\mathbf{x})$ is 1-strongly convex w.r.t $k_t = k_{H_t}$ and consequentially $f_t^*(\mathbf{x})$ is k_t -smooth w.r.t. $k_t = k_{H_t}$. Thus, we have

$$\begin{aligned} & \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t + f_t^*(\mathbf{x}_T) - \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t + f_{T-1}^*(\mathbf{x}_{T-1}) - \mathbf{g}_T^\top \mathbf{x}_{T-1} + f_{T-1}^*(\mathbf{x}_{T-1}) + \frac{1}{2} k_T k_{T-1}^2 \\ & = \sum_{t=1}^{T-1} \mathbf{g}_t^\top \mathbf{x}_t + f_{T-1}^*(\mathbf{x}_{T-1}) + \frac{1}{2} k_T k_{T-1}^2 \end{aligned}$$

By repeating this process, we have

$$\sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t + f_t^*(\mathbf{x}_T) = f_0^*(\mathbf{x}_0) + \frac{1}{2} \sum_{t=1}^T k_t k_{t-1}^2 = \frac{1}{2} \sum_{t=1}^T k_t k_{t-1}^2$$

Then

$$\sum_{t=1}^T (f(\mathbf{x}_t) - f(\mathbf{x})) \leq \frac{1}{2} \sum_{t=1}^T k_t k_{t-1}^2 + \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t \quad (16)$$

Following the analysis in Duchi et al. (2011), we have

$$\sum_{t=1}^T k_t k_{t-1}^2 \leq \frac{1}{2} \sum_{i=1}^d k_{g_{1:T};i} k_2$$

Thus

$$\begin{aligned} \sum_{t=1}^T (f(\mathbf{x}_t) - f(\mathbf{x})) & \leq \frac{G \|\mathbf{x} - \mathbf{x}_1\|_2^2}{2} + \frac{(\mathbf{x} - \mathbf{x}_1)^\top \text{diag}(S_T) (\mathbf{x} - \mathbf{x}_1)}{2} + \sum_{i=1}^d k_{g_{1:T};i} k_2 + \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t \\ & \leq \frac{G + \max_i k_{g_{1:T};i} k_2}{2} \|\mathbf{x} - \mathbf{x}_1\|_2^2 + \sum_{i=1}^d k_{g_{1:T};i} k_2 + \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t \end{aligned}$$

Now by the value of $T = M \max_{i=1}^d \frac{(G + \max_i k_{g_{1:T};i} k_2)}{2c} \cdot c \sum_{i=1}^d k_{g_{1:T};i} k_2$, we have

$$\frac{(G + \max_i k_{g_{1:T};i} k_2)}{2T} \leq \frac{c}{M} \cdot \frac{\sum_{i=1}^d k_{g_{1:T};i} k_2}{T} \leq \frac{c}{M}$$

Dividing by T on both sides and setting $\mathbf{x} = \mathbf{x}_*$, following the inequality (3) and the convexity of $f(\mathbf{x})$ we have

$$f(\mathbf{x}_*) - f_* \leq \frac{c}{M} \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + \frac{1}{M} \sum_{t=1}^T \mathbf{g}_t^\top \mathbf{x}_t$$

Let \mathcal{F}_t be the filtration associated with Algorithm 1 in the paper. Noticing that T is a random variable with respect to \mathcal{F}_t , we cannot get rid of the last term directly. Define the Sequence $\mathcal{X}_t \in \mathbb{N}_+$ as

$$\mathcal{X}_t = \frac{1}{t} \sum_{i=1}^t \mathbf{g}_i^\top \mathbf{x}_i - \mathbf{g}_i^\top \mathbf{x}_* \quad (17)$$

where $\mathbf{g}_i \in \partial f(\mathbf{x}_i)$. Since $\mathbb{E}[\mathbf{g}_{t+1} | \mathcal{F}_t] = 0$ and $\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \frac{1}{t} \sum_{i=1}^t \mathbf{g}_i + \frac{1}{t} f_t(\mathbf{x})$, which is measurable with respect to $\mathbf{g}_1, \dots, \mathbf{g}_t$ and $\mathbf{x}_1, \dots, \mathbf{x}_t$, it is easy to see $\mathcal{X}_t \in \mathcal{F}_t$ is a martingale difference sequence with respect to \mathcal{F}_t , e.g. $\mathbb{E}[\mathcal{X}_t | \mathcal{F}_{t-1}] = 0$. On the other hand, since $\|\mathbf{g}_t\|_2$ is upper bounded by G , following the statement of T in the theorem, $T \leq N = M^2 \max_{i=1}^d \frac{(G+1)^2}{4} \cdot d^2 G^2 < 1$ always holds. Then following Lemma 1 in (Chen et al., 2018b) we have that $\mathbb{E}[\mathcal{X}_T] = 0$.

Now taking the expectation we have that

$$\mathbb{E}[f(\mathbf{x}_*) - f_*] \leq \frac{c}{M} \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + \frac{c}{M}$$

Then we finish the proof. \square

H PROOF OF INEQUALITY (3)

Let us recall the definition of $\mathbf{x} = \text{prox}_{f(\mathbf{x})}$.

$$\text{prox}_{f(\mathbf{x})} = \arg \min_{\mathbf{z}} f(\mathbf{z}) + \frac{1}{2} k \|\mathbf{z} - \mathbf{x}\|^2 \quad (18)$$

First, note that when $f(\mathbf{z})$ is μ -weakly convex and $k > \mu$, the above problem is strongly convex and $\mathbf{x} = \text{prox}_{f(\mathbf{x})}$ is well-defined and unique.

By the optimality condition of \mathbf{x} , it is clear that $f(\mathbf{x}) + \frac{1}{2} k \|\mathbf{x} - \mathbf{x}\|^2 = f(\mathbf{x})$, which proves the first inequality $f(\mathbf{x}) \leq f(\mathbf{x})$. Note that when $\mathbf{x} = \mathbf{x}_* \in \arg \min_{\mathbf{z}} f(\mathbf{z})$, we can prove that $\mathbf{x} = \mathbf{x} = \mathbf{x}_*$. This is because that by definition of \mathbf{x} , it is unique and satisfies $0 \in \partial f(\mathbf{x}) + \frac{1}{2} k(\mathbf{x} - \mathbf{x})$. When $\mathbf{x} = \mathbf{x}_*$, we have $0 \in \partial f(\mathbf{x}_*)$ and \mathbf{x}_* clearly satisfies the first-order condition $0 \in \partial f(\mathbf{x}_*) + \frac{1}{2} k(\mathbf{x}_* - \mathbf{x}_*)$. By the uniqueness of \mathbf{x} , it follows that $\mathbf{x} = \mathbf{x}_*$.

The first-order condition $0 \in \partial f(\mathbf{x}) + \frac{1}{2} k(\mathbf{x} - \mathbf{x})$ also gives that $\text{dist}(0; \partial f(\mathbf{x})) \leq \frac{1}{2} k \|\mathbf{x} - \mathbf{x}\|$. Also, we have $r f(\mathbf{x}) = \frac{1}{2} k \|\mathbf{x} - \mathbf{x}\|$ (Rockafellar, 1970), which implies the second and the third inequality in (3).