
Revisiting Convergence of AdaGrad with Relaxed Assumptions

Yusu Hong¹

Junhong Lin¹

¹Zhejiang University

Abstract

In this study, we revisit the convergence of AdaGrad with momentum (covering AdaGrad as a special case) on non-convex smooth optimization problems. We consider a general noise model where the noise magnitude is controlled by the function value gap together with the gradient magnitude. This model encompasses a broad range of noises including bounded noise, sub-Gaussian noise, affine variance noise and the expected smoothness, and it has been shown to be more realistic in many practical applications. Our analysis yields a probabilistic convergence rate which, under the general noise, could reach at $\tilde{O}(1/\sqrt{T})$. This rate does not rely on prior knowledge of problem-parameters and could accelerate to $\tilde{O}(1/T)$ where T denotes the total number iterations, when the noise parameters related to the function value gap and noise level are sufficiently small. The convergence rate thus matches the lower rate for stochastic first-order methods over non-convex smooth landscape up to logarithm terms [Arjevani et al., 2023]. We further derive a convergence bound for AdaGrad with momentum, considering the generalized smoothness where the local smoothness is controlled by a first-order function of the gradient norm.

1 INTRODUCTION

In recent years, AdaGrad [Duchi et al., 2011] and its variants have witnessed a large success in solving the following stochastic optimization problems:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}), \quad \text{where } f(\mathbf{x}) = \mathbb{E}_{\zeta} [f_{\zeta}(\mathbf{x}; \zeta)].$$

Distinct from vanilla Stochastic Gradient Descent (SGD) [Robbins and Monro, 1951], which typically requires

smoothness or Lipschitz constants for tuning step-sizes, AdaGrad applies an adaptive step-size with each coordinate satisfying that

$$\eta_{t,i} = \frac{\eta}{\sqrt{\sum_{s=1}^t g_{s,i}^2 + \epsilon}}, \quad \forall t \in \mathbb{N}, i \in [d],$$

where $\epsilon > 0$ is a constant and $g_{s,i}$ denotes the i -th coordinate of the stochastic gradient \mathbf{g}_s . This approach assigns larger step-sizes for infrequent features whose corresponding gradients are small, reminding learners of taking notice of those infrequent features. It also liberates the algorithm from the need for problem-parameters, which may be challenging to obtain in practical applications. Moreover, AdaGrad's efficiency has been empirically validated, especially in scenarios with sparse gradients [Duchi et al., 2011].

Numerous works have studied the convergence of AdaGrad and its scalar version, AdaGrad-Norm [Duchi et al., 2011, Streeter and McMahan, 2010]. Duchi et al. [2011] first provided the convergence bound of AdaGrad on online convex optimization. In non-convex smooth scenario, Ward et al. [2020] first obtained a convergence bound for AdaGrad-Norm without pre-tuning step-sizes, assuming bounded gradients and noises. Liu et al. [2023b] proved the convergence for AdaGrad under coordinate-wise sub-Gaussian noise, discarding the bounded gradient assumption.

Recently, several studies have proven AdaGrad-Norm's convergence under the affine variance noise, both in expectation [Faw et al., 2022, Wang et al., 2023] and in high probability [Attia and Koren, 2023]. The noise model assumes that the stochastic gradient $g(\mathbf{x}), \forall \mathbf{x} \in \mathbb{R}^d$ satisfies that for some constants $B, C > 0$,

$$\begin{aligned} \mathbb{E} \|g(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 &\leq B \|\nabla f(\mathbf{x})\|^2 + C \\ \text{or } \|g(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 &\leq B \|\nabla f(\mathbf{x})\|^2 + C. \end{aligned} \quad (1)$$

This noise model, verified in machine learning applications with feature noise [Fuller, 2009, Khani and Liang, 2020], and in robust linear regression [Xu et al., 2008], offers a

more realistic portrayal by allowing the noise norm to increase with the gradient norm, covering both bounded noise and sub-Gaussian noise. These studies not only provided a convergence rate of $\tilde{O}(1/\sqrt{T})$, but also addressed challenges posed by the entanglement of adaptive step-sizes and stochastic gradients, and the additional variance in (1). However, to the best of our knowledge, none of existing works have proved the convergence of vanilla AdaGrad under (1) without assuming bounded gradients. Moreover, the distinct step-size for each coordinate in AdaGrad, as opposed to an unified step-size for all coordinates in AdaGrad-Norm, brings more challenges when considering (1).

In this paper, we provide a deep analysis framework and establish a probabilistic convergence bound for AdaGrad with heavy-ball style momentum, covering AdaGrad as a special case. More importantly, we consider a general noise model such that for some constants $A, B, C \geq 0$,

$$\|g(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq A(f(\mathbf{x}) - f^*) + B\|\nabla f(\mathbf{x})\|^2 + C, \quad (2)$$

where $f(\mathbf{x}) \geq f^*, \forall \mathbf{x} \in \mathbb{R}^d$. It's obvious to verify that (2) is strictly weaker than the almost surely affine variance noise in (1), and thus than the bounded noise or sub-Gaussian noise condition¹. Indeed, (2) could be regarded as an extension of (1) and the following expected smoothness condition [Gower et al., 2019, Grimmer, 2019, Wang and Yuan, 2023],

$$\begin{aligned} \mathbb{E}\|g(\mathbf{x})\|^2 &\leq A(f(\mathbf{x}) - f^*) + B \\ \text{or } \|g(\mathbf{x})\|^2 &\leq A(f(\mathbf{x}) - f^*) + B. \end{aligned} \quad (3)$$

Existing researches have studied SGD's convergence behavior under (2) with smooth objective functions, both in asymptotic [Poljak and Tsytkin, 1973] and non-asymptotic view [Khaled and Richtárik, 2023]. More importantly, it has been shown that numerous of practical stochastic gradient settings satisfy (2) but out of the range of (1), including commonly used perturbation, sub-sampling and compression [Khaled and Richtárik, 2023]. However, the analysis for SGD could not be directly extended to AdaGrad due to the correlation of adaptive-sizes and stochastic gradients, and the coordinate-wise performance in AdaGrad.

Finally, we apply our analysis framework to the (L_0, L_1) -smoothness where the local smoothness of f satisfies that when $\|\mathbf{y} - \mathbf{x}\| \leq 1/L_1$,

$$\|\nabla f(\mathbf{y}) - \nabla f(\mathbf{x})\| \leq (L_0 + L_1\|\nabla f(\mathbf{x})\|)\|\mathbf{y} - \mathbf{x}\|. \quad (4)$$

This assumption was proposed by [Zhang et al., 2020b] through empirical studies on language models and later verified in large language models, e.g., [Zhang et al., 2020a,

¹For conciseness, we mainly consider the almost-sure version of this general noise model. Extending our high probability analysis from an almost-sure version to a sub-Gaussian version is easy, which will be included in Appendix.

Crawshaw et al., 2022]. (4) generalizes the standard global smoothness condition and allows unbounded smooth parameter, bringing more challenges for the convergence analysis of adaptive methods. Previous works [Faw et al., 2023, Wang et al., 2023] have derived convergence bounds for AdaGrad-Norm with (1) and (4). Also, prior knowledge of problem-parameters is necessary as pointed out by the counter examples in [Wang et al., 2023]. However, the analysis for coordinate-wise AdaGrad is non-trivial and requires more delicate constructions, particularly when considering the weaker noise assumption in (2).

In the following, we will summarize our main contributions as follows. We also refer readers to see the comparison of our results with existing works in Table 1 from the appendix.

Contribution

- We demonstrate the probabilistic convergence of AdaGrad with momentum on non-convex smooth optimization under a general noise assumption in (2). For an L -smooth function f , we demonstrate that after T iterations of the algorithm, with probability at least $1 - \delta$, $\sum_{s=1}^T \|\nabla f(\mathbf{x}_s)\|^2/T$ is bounded by

$$\mathcal{O}\left(\frac{\text{poly}\left(\log \frac{T}{\delta}\right)}{T} + \sqrt{\frac{(A+C)\text{poly}\left(\log \frac{T}{\delta}\right)}{T}}\right),$$

which could also accelerate to $\tilde{O}(1/T)$ rate when the noise parameters A, C are sufficiently low.

- As direct corollaries, we also derive similar probabilistic convergence results of AdaGrad on non-convex smooth optimization with affine variance noise. More importantly, the convergence rate is optimal and adaptive to the noise level C in (1).
- We derive a convergence bound for AdaGrad with momentum considering (2) and (4). The rate is similar to the smooth case and adaptive to the noise level as well, and necessitating problem-parameters to tune step-sizes.

Our analysis relies on the descent lemma with telescoping, and the novel decomposition and estimations over the first-order term related to new proxy step-sizes that are used to decorrelate stochastic gradients and adaptive step-sizes in this new noise regime. We also prove that the function value gap as well as the gradient norm are controlled by the polynomial of $\log T$ along the optimization process.

The rest of the paper are organized as follows. The next section introduces some extra related works. Section 3 provides the problem setup and basic assumptions, and the introduction of AdaGrad with momentum. Section 4 provides high probability convergence bounds for AdaGrad with momentum, and also for AdaGrad as direct corollaries. Section 5 provides proof details for the main results. Section

6 presents necessary introduction of the generalized smooth condition and the subsequent convergence result. All missing proofs for some of the lemmas and convergence results under generalized smoothness are given in Appendix.

2 RELATED WORKS

SGD and its adaptive variants have been a target of intense interest in the last decade. We refer to [Bottou et al., 2018, Ruder, 2016] for an overview. We limit our discussions to the most relevant literature in the sequel.

Convergence of AdaGrad Numerous of works mainly studied the convergence of AdaGrad-Norm over non-convex smooth landscape. Li and Orabona [2019] first proved the convergence for AdaGrad-Norm. However, they studied a variant with a delayed step-size that is independent from the current stochastic gradient and required knowledge of the smoothness parameter for tuning step-sizes. Getting rid of prior-knowledge on problem-parameters, Ward et al. [2020] relied on a novel proxy step-size technique and showed the convergence with an uniform bound of stochastic gradients for vanilla AdaGrad-Norm. Kavis et al. [2022] and Liu et al. [2023b] proved probabilistic convergence under the sub-Gaussian noise without relying on the bounded gradient assumption. In case of the affine variance noise, Faw et al. [2022] provided the convergence bound. However, their rate is adaptive to the noise level only when $B \sim \mathcal{O}(1/T)$. Wang et al. [2023] relied on a distinct framework to improve the dependency on T in the convergence rate [Faw et al., 2022] and achieved the adaptivity on noise level without any restriction over B . Concurrently, Attia and Koren [2023] deduced a probabilistic bound, using a novel induction argument to control the function value gap. Their result also adapted to the noise level without further requirement on B . Liu et al. [2023a] formulated a convergence bound for AdaGrad-Norm and its acceleration version in quasar-convex smooth setting.

The element-wise version of AdaGrad was first studied in [Duchi et al., 2011] on online convex optimization. In non-convex smooth case, a line of works investigated AdaGrad with bounded gradients. Zou et al. [2019] explored the convergence of AdaGrad with a heavy-ball or Nesterov style momentum. Zhou et al. [2020] also covered AdaGrad in their analysis, but they deduced a bound under requiring bounded gradients’ summation. Défossez et al. [2022] studied AdaGrad with Adam-type momentum and improved the dependency on the momentum parameter β to $\mathcal{O}((1-\beta)^{-1})$. Shen et al. [2023] introduced a weighted AdaGrad with unified momentum covering both heavy-ball and Nesterov’s acceleration. Recently, Liu et al. [2023b], a work mentioned before, derived a convergence bound under coordinate-wise sub-Gaussian noise, i.e., $g(\mathbf{x})_i - \nabla f(\mathbf{x})_i$ is sub-Gaussian for each $i \in [d]$, without requiring bounded gradients.

Convergence with affine variance noise We briefly summarize some works on the convergence of SGD or AdaGrad with (1) and its variants under the non-convex smooth landscape. Bertsekas and Tsitsiklis [2000] provided an almost-surely convergence bound for SGD. In non-asymptotic view, Bottou et al. [2018] derived a convergence bound for SGD of the form $\mathcal{O}(1/T + \sqrt{C/T})$ when step-sizes are well tuned by the smooth parameter, B and C . They also pointed out that the extension is immediate from the bounded noise case [Ghadimi and Lan, 2013]. The convergence of AdaGrad-Norm under (1) has been well studied by [Faw et al., 2022, Wang et al., 2023, Attia and Koren, 2023] as mentioned before. Faw et al. [2023] further extended the analysis considering a generalized smooth condition. However, none of these existing works could prove the convergence of coordinate-wise version of AdaGrad under (1) or a weaker noise condition.

Convergence with the expected smoothness The expected smoothness condition was once applied for convex optimization such that for some constant $A > 0$,

$$\mathbb{E}[\|g(\mathbf{x}) - g(\mathbf{x}^*)\|^2] \leq A(f(\mathbf{x}) - f^*), \quad (5)$$

where \mathbf{x}^* denotes the global minimizer. Based on (5), Richtárik and Takác [2020] relied on matrix analysis to bound the identities of expected iterates of SGD in the setting of stochastic reformulations of linear systems. Gower et al. [2021] applied (5) to analyze the JacSketch method (a general form of SAGA) over strongly convex optimization.

Since \mathbf{x}^* is ill-defined for non-convex optimization, Gower et al. [2019] then directly set $\mathbb{E}[\|g(\mathbf{x}^*)\|^2] = B$ and deduced the non-convex version of the expected smoothness in (3), which aligns with the weak growth condition [Vaswani et al., 2019] when $B = 0$. Gower et al. [2019] relied on (3) to analyze SGD over quasi-strongly convex optimization. Independently, Grimmer [2019] relied on (3) and developed a general framework for SGD equipped with projection operators over convex non-smooth functions. Wang and Yuan [2023] also used (3) to derive a convergence bound for SGD using bandwidth-based step-sizes.

Regarding the noise model in (2), Poljak and Tsyppkin [1973] provided an asymptotic convergence bound for SGD with smooth objective functions. Very recently, Khaled and Richtárik [2023] derived a non-asymptotic convergence rate of $\mathcal{O}(1/\sqrt{T})$ for SGD with non-convex smooth functions.

In conclusion, it’s clear to see that (2) is weaker than the above conditions including (1) (when assuming the existence of f^*), (3) and (5)². Our result then shows that AdaGrad could find a stationary point under this mild noise assumption without prior knowledge of problem-parameters.

²We make the comparison when assuming all conditions are in almost-surely form.

Convergence with generalized smoothness The generalized smooth condition [Zhang et al., 2020b] has been well studied under different algorithms, e.g., [Qian et al., 2021, Zhao et al., 2021, Reisizadeh et al., 2023, Zhang et al., 2020a, Crawshaw et al., 2022]. Considering AdaGrad and its variants, Faw et al. [2023] established a convergence bound for AdaGrad-Norm considering (1). However, their result required $B < 1$. Wang et al. [2023] further tightened the dependency to the iteration number T and got rid of restriction on B .

3 PROBLEM SETTING AND ALGORITHM

We consider unconstrained stochastic optimization over the Euclidean space \mathbb{R}^d with l_2 -norm. The objective function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is L -smooth satisfying that for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$,

$$f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq \frac{L}{2} \|\mathbf{x} - \mathbf{y}\|^2.$$

Given $\mathbf{x} \in \mathbb{R}^d$, we assume a gradient oracle that returns a random vector $g(\mathbf{x}, \mathbf{z}) \in \mathbb{R}^d$, where \mathbf{z} denotes a random sample. The deterministic gradient of f at \mathbf{x} is denoted by $\nabla f(\mathbf{x}) \in \mathbb{R}^d$.

Notations We denote the set $\{1, 2, \dots, T\}$ as $[T]$, and use $\|\cdot\|$, $\|\cdot\|_1$, and $\|\cdot\|_\infty$ to represent the l_2 -norm, l_1 -norm, and l_∞ -norm, respectively. The notations $a \sim \mathcal{O}(b)$ and $a \leq \mathcal{O}(b)$ refer to $a = c_1 b$ and $a \leq c_2 b$ with c_1, c_2 being positive universal constants, and $a \leq \tilde{\mathcal{O}}(b)$ indicates $a \leq \mathcal{O}(b) \text{poly}(\log b)$. For any vector $\mathbf{x} \in \mathbb{R}^d$, the expressions \mathbf{x}^2 and $\sqrt{\mathbf{x}}$ refer to the coordinate-wise square and square root. For two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $\mathbf{x} \odot \mathbf{y}$ and \mathbf{x}/\mathbf{y} denote the coordinate-wise product and quotient. $\mathbf{0}_d$ and $\mathbf{1}_d$ signify zero and one vectors in d dimensions. Further, we write $\mathbf{1}_d/\mathbf{x}$ as $1/\mathbf{x}$, whenever there is no any confusion.

Assumption We make the following assumptions.

- **(A1) Bounded below:** The objective function is bounded below, i.e., there exists $f^* > -\infty$ such that $f(\mathbf{x}) \geq f^*, \forall \mathbf{x} \in \mathbb{R}^d$;
- **(A2) Unbiased estimator:** The gradient oracle provides an unbiased estimator of $\nabla f(\mathbf{x})$, i.e., $\forall \mathbf{x} \in \mathbb{R}^d$, $\mathbb{E}_{\mathbf{z}} [g(\mathbf{x}, \mathbf{z})] = \nabla f(\mathbf{x})$;
- **(A3) Relaxed affine variance noise:** The gradient oracle satisfies that for some constants $A, B, C > 0$, $\|g(\mathbf{x}, \mathbf{z}) - \nabla f(\mathbf{x})\|^2 \leq A(f(\mathbf{x}) - f^*) + B\|\nabla f(\mathbf{x})\|^2 + C$, a.s., $\forall \mathbf{x} \in \mathbb{R}^d$.

The first two assumptions are standard in the analysis of algorithm's convergence. With a simple calculation, it's easy to verify that Assumption (A3) is equivalent to

$$\|g(\mathbf{x}, \mathbf{z})\|^2 \leq A'(f(\mathbf{x}) - f^*) + B'\|\nabla f(\mathbf{x})\|^2 + C'$$

for another three positive constants A', B', C' . Therefore, (A3) is a generalization of (1) and (3). For more detailed examples of stochastic gradient settings satisfying Assumption (A3), we refer interested readers to see [Khaled and Richtárik, 2023, Proposition 2,3].

Algorithm 1 AdaGrad with momentum

Input: Horizon T , $\mathbf{x}_1 \in \mathbb{R}^d$, $\beta \in [0, 1)$, $\mathbf{m}_0 = \mathbf{v}_0 = \mathbf{0}_d$, $\eta, \epsilon > 0$, $\epsilon = \epsilon \mathbf{1}_d$
for $s = 1, \dots, T$ **do**
 Draw a random sample \mathbf{z}_s and generate $\mathbf{g}_s = g(\mathbf{x}_s, \mathbf{z}_s)$;
 $\mathbf{v}_s = \mathbf{v}_{s-1} + \mathbf{g}_s^2$;
 $\mathbf{m}_s = \beta \mathbf{m}_{s-1} - \eta \mathbf{g}_s / (\sqrt{\mathbf{v}_s} + \epsilon)$;
 $\mathbf{x}_{s+1} = \mathbf{x}_s + \mathbf{m}_s$;
end for

AdaGrad with momentum Throughout the paper, we study AdaGrad with momentum given in Algorithm 1. We can transform Algorithm 1 into the classical Polyak's heavy-ball method [Polyak, 1964] with an adaptive step-size:

$$\mathbf{x}_{s+1} = \mathbf{x}_s - \eta \frac{\mathbf{g}_s}{\sqrt{\mathbf{v}_s} + \epsilon} + \beta(\mathbf{x}_s - \mathbf{x}_{s-1}), \quad \forall s \in [T], \quad (6)$$

where we set $\mathbf{x}_0 = \mathbf{x}_1$.

AdaGrad AdaGrad is Algorithm 1 with $\beta = 0$.

4 MAIN CONVERGENCE RESULT

In this section, we provide the probabilistic convergence result for Algorithm 1 under Assumption (A3) and smooth objective functions.

Theorem 1. *Given $T \geq 1$, let $\{\mathbf{x}_s\}_{s \in [T]}$ be generated by Algorithm 1. If Assumptions (A1), (A2), (A3) hold, then for any $\beta \in [0, 1)$, $\eta, \epsilon > 0$ and $\delta \in (0, 1)$, it holds that with probability at least $1 - \delta$,*

$$\begin{aligned} & \frac{1}{T} \sum_{s=1}^T \|\nabla f(\mathbf{x}_s)\|^2 \\ & \leq \mathcal{O} \left[\Delta_1 \left(\frac{B_1 \Delta_1 + \sqrt{B_1 L \Delta} + \epsilon}{T} + \sqrt{\frac{A \Delta + C}{T}} \right) \right], \end{aligned}$$

where $B_1 = B + 1$, $\Delta_1 = \Delta(1 - \beta)/\eta$, and Δ is given by³

$$\begin{aligned} \Delta \sim \mathcal{O} & \left[f(\mathbf{x}_1) - f^* + \frac{\sqrt{C} \eta d}{1 - \beta} \log \left(\frac{T}{\delta} + \frac{T}{\epsilon^2} \right) \right. \\ & \left. + \frac{(A + B_1 L) \eta^2 d^2}{(1 - \beta)^3} \log^2 \left(\frac{T}{\delta} + \frac{T}{\epsilon^2} \right) \right]. \end{aligned}$$

³The detailed expression of Δ could be found in (20).

Remark 4.1. With a simple calculation, when $\eta = c_1(1 - \beta)^{3/2}$ for some constant $c_1 > 0$, the above upper bound has a minimum order of $\mathcal{O}((1 - \beta)^{-1})$ with respect to $(1 - \beta)^{-1}$. The comparison of existing results with our convergence bound could be found in Table 1 from the appendix.

Convergence of AdaGrad with affine variance noise As a direct consequence of Theorem 1, it's worthy to mention the following convergence bound for AdaGrad with affine variance noise considering their empirical significance.

Corollary 1. Under the assumptions and notations of Theorem 1, let $\beta = 0$ and $A = 0$. Then for any $\eta, \epsilon > 0$ and $\delta \in (0, 1)$, it holds that with probability at least $1 - \delta$,

$$\begin{aligned} & \frac{1}{T} \sum_{s=1}^T \|\nabla f(\mathbf{x}_s)\|^2 \\ & \leq \mathcal{O} \left[\Delta_1 \left(\frac{B_1 \Delta_1 + \sqrt{B_1 L \Delta} + \epsilon}{T} + \sqrt{\frac{C}{T}} \right) \right], \end{aligned}$$

where $B_1 = B + 1$, $\Delta_1 = \Delta/\eta$, and Δ is defined as follows

$$\begin{aligned} \Delta \sim \mathcal{O} & \left[f(\mathbf{x}_1) - f^* + \sqrt{C} \eta d \log \left(\frac{T}{\delta} + \frac{T}{\epsilon^2} \right) \right. \\ & \left. + B_1 L \eta^2 d^2 \log^2 \left(\frac{T}{\delta} + \frac{T}{\epsilon^2} \right) \right]. \end{aligned}$$

Remark 4.2. 1) Setting $\eta \sim 1/(d \log(\frac{T}{\delta}))$, then $\Delta \sim 1$, and the above derived upper bound is of order $\mathcal{O}(d^2 \log^2(\frac{T}{\delta})/T + d \log(\frac{T}{\delta}) \sqrt{C/T})$, matching the lower rate in [Arjevani et al., 2023] up to logarithm factors.

2) The convergence rate is of order $\tilde{\mathcal{O}}(1/T + \sqrt{C/T})$, and when the noise level C is sufficiently low, the convergence rate could be $\tilde{\mathcal{O}}(1/T)$, which aligns with the result for non-adaptive SGD under the same conditions [Ghadimi and Lan, 2013, Bottou et al., 2018] up to logarithmic terms.

3) As in standard probability theory, the derived high-probability convergence can ensure expected convergence.

4) Assumption (A3) can be replaced by its sub-Gaussian form where $\mathbb{E}_{\mathbf{z}} \left[\exp \left(\frac{\|g(\mathbf{x}, \mathbf{z}) - \nabla f(\mathbf{x})\|^2}{A(f(\mathbf{x}) - f^*) + B \|\nabla f(\mathbf{x})\|^2 + C} \right) \right] \leq e$, and our results still hold true, as shown in Appendix.

5 PROOF DETAIL

To start with, we let $\mathbf{g}_s = (g_{s,i})_i$ be as in Algorithm 1 and let $\nabla f(\mathbf{x}_s) = \bar{\mathbf{g}}_s = (\bar{g}_{s,i})_i$, $\boldsymbol{\xi}_s = (\xi_{s,i})_i = \mathbf{g}_s - \bar{\mathbf{g}}_s$ and $\Delta_s^{(x)} = f(\mathbf{x}_s) - f^*$.

During the proof, we will introduce several key lemmas to deduce the final results. All the missing proofs could be found in Appendix.

5.1 PRELIMINARY

Before proving the main result, we shall introduce several useful auxiliary sequences. The first sequence $\{\mathbf{y}_s\}_{s \geq 1}$ is defined as

$$\mathbf{y}_1 = \mathbf{x}_1, \mathbf{y}_s = \frac{\beta}{1 - \beta} (\mathbf{x}_s - \mathbf{x}_{s-1}) + \mathbf{x}_s, \quad \forall s \geq 2, \quad (7)$$

following from [Ghadimi et al., 2015, Yang et al., 2016] which was used to prove the convergence of SGD with momentum and later applied to handle with many variants of momentum-based algorithms. When \mathbf{x}_s is generated by Algorithm 1, we reveal that \mathbf{y}_s satisfies that for any $s \geq 1$,

$$\mathbf{y}_{s+1} = \mathbf{y}_s - \frac{\eta}{1 - \beta} \frac{\mathbf{g}_s}{\mathbf{b}_s}, \quad \mathbf{b}_s = \sqrt{\mathbf{v}_s} + \epsilon. \quad (8)$$

We let the function value gap $\Delta_s^{(y)} = f(\mathbf{y}_s) - f^*$. In addition, we introduce $\{\mathcal{G}_s\}_{s \geq 1}$ and the value \mathcal{G} ,

$$\begin{aligned} \mathcal{G}_s &= \sqrt{X \Delta_s^{(x)} + 2C}, \\ \mathcal{G} &= \sqrt{X \Delta + 2C}, \quad X = 2A + 4LB + 4L, \end{aligned} \quad (9)$$

where Δ is as in Theorem 1.

5.2 ROUGH ESTIMATIONS

Motivated by [Faw et al., 2022], we provide some rough estimations for several key algorithm-dependent terms in this section. These estimations are not delicate, but they play vital roles in further deducing the final convergence rate.

Lemma 5.1. For any $s \geq 1$ and $\beta \in [0, 1)$,

$$\|\mathbf{m}_s\| \leq \frac{\eta \sqrt{d}}{1 - \beta}, \quad \|\bar{\mathbf{g}}_s\| \leq \|\bar{\mathbf{g}}_1\| + \frac{L \eta s \sqrt{d}}{1 - \beta}.$$

Lemma 5.2. Suppose that $\beta \in [0, 1)$. Then for any $T \geq 1$,

$$\begin{aligned} \sum_{t=1}^T \Delta_t^{(x)} &\leq \Delta_1^{(x)} T \\ &+ \left(\frac{\eta \|\bar{\mathbf{g}}_1\| \sqrt{d}}{1 - \beta} + \frac{L \eta^2 d}{2(1 - \beta)^2} \right) T^2 + \frac{L \eta^2 d T^3}{(1 - \beta)^2}. \end{aligned}$$

5.3 START POINT AND DECOMPOSITION

We now proceed the proof for the main result. We fix the horizon T . Following [Ward et al., 2020], we start from the descent lemma of smoothness over \mathbf{y}_s with both sides subtracting with f^* ,

$$\Delta_{s+1}^{(y)} \leq \Delta_s^{(y)} + \langle \nabla f(\mathbf{y}_s), \mathbf{y}_{s+1} - \mathbf{y}_s \rangle + \frac{L}{2} \|\mathbf{y}_{s+1} - \mathbf{y}_s\|^2.$$

Combining with (8), and summing over $s \in [t]$,

$$\begin{aligned} \Delta_{t+1}^{(y)} &\leq \Delta_1^{(x)} + \frac{\eta}{1-\beta} \underbrace{\left(- \sum_{s=1}^t \left\langle \nabla f(\mathbf{y}_s), \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\rangle \right)}_{\mathbf{A}} \\ &\quad + \frac{L\eta^2}{2(1-\beta)^2} \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2, \end{aligned} \quad (10)$$

where we apply $\mathbf{y}_1 = \mathbf{x}_1$. We subsequently further make a decomposition over \mathbf{A} as

$$\mathbf{A} = - \underbrace{\sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s, \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\rangle}_{\mathbf{A.1}} + \underbrace{\sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s - \nabla f(\mathbf{y}_s), \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\rangle}_{\mathbf{A.2}}. \quad (11)$$

5.4 ESTIMATING \mathbf{A}

The first main challenge comes from the entanglement of \mathbf{g}_s and \mathbf{b}_s emerging in \mathbf{A} , which is a key problem distinct from the analysis for SGD.

Estimating A.1 We adopt the so-called proxy step-size technique which is a commonly used technique for breaking the correlation of \mathbf{b}_s and \mathbf{g}_s in the analysis of adaptive methods. This technique relies on introducing appropriate proxy step-sizes. It has been first introduced in [Ward et al., 2020] for AdaGrad-Norm with bounded stochastic gradients and variants of proxy step-sizes have been developed in the related literature, e.g., [Défossez et al., 2022, Faw et al., 2022, Attia and Koren, 2023, Liu et al., 2023b]. However, none of these proxy step-sizes could be potentially applied for AdaGrad with potential unbounded gradients under the mild noise model in Assumption (A3).

We thus provide a construction of proxy step-sizes that is general enough to handle with Assumption (A3). The proxy step-sizes rely on \mathcal{G}_s given in (9), specifically defined in terms of

$$\mathbf{a}_s = \sqrt{\mathbf{v}_{s-1} + (\mathcal{G}_s \mathbf{1}_d)^2} + \epsilon, \quad \forall s \in [T]. \quad (12)$$

Based on the proxy step-size η/\mathbf{a}_s , we further have

$$\begin{aligned} \mathbf{A.1} &= - \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 \\ &\quad - \underbrace{\sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s, \frac{\boldsymbol{\xi}_s}{\mathbf{a}_s} \right\rangle}_{\mathbf{A.1.1}} + \underbrace{\sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s, \left(\frac{1}{\mathbf{a}_s} - \frac{1}{\mathbf{b}_s} \right) \mathbf{g}_s \right\rangle}_{\mathbf{A.1.2}}. \end{aligned} \quad (13)$$

The estimation for **A.1.1** relies on a probabilistic analysis over a summation of martingale difference sequence.

Lemma 5.3. *Given $T \geq 1$ and $\delta \in (0, 1)$, if Assumptions (A2) and (A3) hold, then with probability at least $1 - \delta$,*

$$\mathbf{A.1.1} \leq \frac{1}{4} \sum_{s=1}^t \frac{\mathcal{G}_s}{\mathcal{G}} \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 + 3\mathcal{G} \log \left(\frac{T}{\delta} \right), \quad \forall t \in [T], \quad (14)$$

where $\mathcal{G}_s, \mathcal{G}$ are as in (9).

The **A.1.2** serves as an error term for introducing \mathbf{a}_s . However, due to the delicate construction of \mathbf{a}_s , we could estimate the gap as follows,

Lemma 5.4. *Under Assumption (A3), let $\mathbf{b}_s = (b_{s,i})_i, \mathbf{a}_s = (a_{s,i})_i$ be defined in (8) and (12). Then*

$$\left| \frac{1}{a_{s,i}} - \frac{1}{b_{s,i}} \right| \leq \frac{\mathcal{G}_s}{a_{s,i} b_{s,i}}, \quad \forall s \in [T], \forall i \in [d].$$

Based on this lemma, it's then shown in the following lemma that **A.1.2** could be controlled.

Lemma 5.5. *Under Assumption (A3), for any $t \geq 1$, if $\beta \in [0, 1)$, it holds that*

$$\mathbf{A.1.2} \leq \frac{1}{4} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 + \sum_{s=1}^t \mathcal{G}_s \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2. \quad (15)$$

Finally, we rely on the smoothness to estimate **A.2**.

Lemma 5.6. *For any $t \geq 1$, if $\beta \in [0, 1)$, it holds that*

$$\mathbf{A.2} \leq \frac{L}{2\eta} \sum_{s=1}^t \|\mathbf{m}_{s-1}\|^2 + \frac{L\eta}{2(1-\beta)^2} \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2. \quad (16)$$

5.5 BOUNDING THE FUNCTION VALUE GAP

Based on the above estimations, we could use an induction argument to deduce an upper bound for function value gaps. The induction technique is motivated by [Attia and Koren, 2023] where AdaGrad-Norm with affine variance noise was studied. As we study a more relaxed assumption on AdaGrad, it's required to provide some new estimations.

Proposition 5.1. *Under the same conditions of Theorem 1, the following two inequalities hold with probability at least $1 - \delta$,*

$$\Delta_t^{(x)} \leq \Delta, \quad \mathcal{G}_t \leq \mathcal{G}, \quad \forall t \in [T+1], \quad (17)$$

and

$$\Delta_{t+1}^{(x)} \leq \Delta - \frac{\eta}{1-\beta} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2, \quad \forall t \in [T], \quad (18)$$

where Δ is as in Theorem 1 and $\mathcal{G}_t, \mathcal{G}$ are as in (9).

In what follows, we prove Proposition 5.1. We assume that (14) always happens and then deduce (17) and (18). Recall that (14) holds with probability at least $1 - \delta$. We therefore obtain that both (17) and (18) would hold with probability at least $1 - \delta$. We first plug (13), (14) and (15) into (11), and then combine with (16) and (10) to get that

$$\begin{aligned} \Delta_{t+1}^{(y)} &\leq \Delta_1^{(x)} + \frac{\eta}{1-\beta} \sum_{s=1}^t \left(\frac{\mathcal{G}_s}{4\mathcal{G}} - \frac{3}{4} \right) \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 \\ &\quad + \frac{3\mathcal{G}\eta}{1-\beta} \log\left(\frac{T}{\delta}\right) + \frac{\eta}{1-\beta} \sum_{s=1}^t \mathcal{G}_s \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 \\ &\quad + \frac{L}{2(1-\beta)} \sum_{s=1}^t \|\mathbf{m}_{s-1}\|^2 + \tilde{L} \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2, \end{aligned} \quad (19)$$

where we let $\tilde{L} = \frac{L\eta^2}{2(1-\beta)^3} + \frac{L\eta^2}{2(1-\beta)^2}$. Then, we present the specific definition of Δ as

$$\begin{aligned} \Delta &:= 4\Delta_1^{(x)} + \frac{12\sqrt{2C}\eta}{1-\beta} \log\left(\frac{T}{\delta}\right) \\ &\quad + 4 \left(\frac{\sqrt{2C}\eta}{1-\beta} + \frac{\eta^2 L}{(1-\beta)^3} + \tilde{L} \right) d \log \mathcal{F}_T \\ &\quad + \frac{72X\eta^2}{(1-\beta)^2} \log^2\left(\frac{T}{\delta}\right) + \frac{8X\eta^2}{(1-\beta)^2} d^2 \log^2 \mathcal{F}_T. \end{aligned} \quad (20)$$

Here, \mathcal{F}_T is a polynomial with respect to T with the detailed expression in (41) from Appendix. Then, it's easy to verify that $\Delta_1^{(x)} \leq \Delta$. Suppose that for some $t \in [T]$,

$$\Delta_s^{(x)} \leq \Delta, \forall s \in [t], \quad \text{thus,} \quad \mathcal{G}_s \leq \mathcal{G}, \forall s \in [t]. \quad (21)$$

In order to apply (19) to control $\Delta_{t+1}^{(x)}$, we introduce the following lemma to lower bound the LHS of (19).

Lemma 5.7. *Let \mathbf{y}_s be defined in (7) and $\beta \in [0, 1)$. Then for any $s \geq 1$,*

$$\Delta_s^{(y)} \geq \frac{\Delta_s^{(x)}}{2} - \frac{L\|\mathbf{m}_{s-1}\|^2}{2(1-\beta)^2}.$$

Based on Lemma 5.7, the LHS of (19) could be lower bounded in terms of $\Delta_{t+1}^{(x)}$. We use (21) to upper bound the RHS of (19), which leads to

$$\begin{aligned} \frac{\Delta_{t+1}^{(x)}}{2} &\leq \Delta_1^{(x)} - \frac{\eta}{2(1-\beta)} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 + \frac{L\|\mathbf{m}_t\|^2}{2(1-\beta)^2} \\ &\quad + \frac{3(\sqrt{X\Delta} + \sqrt{2C})\eta}{1-\beta} \log\left(\frac{T}{\delta}\right) + \tilde{L} \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 \\ &\quad + \frac{(\sqrt{X\Delta} + \sqrt{2C})\eta}{1-\beta} \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 + \sum_{s=1}^t \frac{L\|\mathbf{m}_{s-1}\|^2}{2(1-\beta)}, \end{aligned}$$

where we use $\mathcal{G} \leq \sqrt{X\Delta} + \sqrt{2C}$. Further, using Young's inequality twice for the terms related to $\sqrt{X\Delta}$, and $\beta < 1$,

$$\begin{aligned} \frac{\Delta_{t+1}^{(x)}}{2} &\leq \frac{\Delta}{4} + \Delta_1^{(x)} - \frac{\eta}{2(1-\beta)} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 \\ &\quad + \frac{L\|\mathbf{m}_t\|^2}{2(1-\beta)^2} + \frac{3\sqrt{2C}\eta}{1-\beta} \log\left(\frac{T}{\delta}\right) \\ &\quad + \left(\frac{\sqrt{2C}\eta}{1-\beta} + \tilde{L} \right) \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 + \sum_{s=1}^t \frac{L\|\mathbf{m}_{s-1}\|^2}{2(1-\beta)} \\ &\quad + \frac{18X\eta^2}{(1-\beta)^2} \log^2\left(\frac{T}{\delta}\right) + \frac{2X\eta^2}{(1-\beta)^2} \left(\sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 \right)^2. \end{aligned} \quad (22)$$

Finally, we shall use the following lemma to further estimate $\|\mathbf{m}_t\|$ and the other two summations related to $\mathbf{g}_s, \mathbf{b}_s, \mathbf{m}_s$.

Lemma 5.8. *Given $T \geq 1$ and $\beta \in [0, 1)$, then for any $t \in [T]$,*

$$\begin{aligned} \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 &\leq d \log \mathcal{F}_T, \quad \|\mathbf{m}_t\|^2 \leq \frac{\eta^2 d}{1-\beta} \log \mathcal{F}_T, \\ \sum_{s=1}^t \|\mathbf{m}_s\|^2 &\leq \frac{\eta^2 d}{(1-\beta)^2} \log \mathcal{F}_T, \end{aligned}$$

where \mathcal{F}_T is a polynomial with respect to T with the detailed expression in (41) from Appendix.

Compared with Lemma 5.1, Lemma 5.8 improves the dependency to $1 - \beta$ for estimating $\|\mathbf{m}_t\|^2$, which leads to the $\mathcal{O}((1 - \beta)^{-1})$ order for the final convergence as in Remark 4.1. Thus, applying Lemma 5.8 over (22), and then combining with Δ in (20),

$$\frac{\Delta_{t+1}^{(x)}}{2} \leq \frac{\Delta}{2} - \frac{\eta}{2(1-\beta)} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 \leq \frac{\Delta}{2}.$$

The induction is complete and the desired result in (17) is proved. Finally, as an intermediate result, we verify (18).

5.6 PROOF OF THE MAIN RESULT

Based on Proposition 5.1, we are able to prove Theorem 1.

Proof of Theorem 1. In what follows, we will obtain the final convergence result based on (17) and (18). Since (17) and (18) hold with probability at least $1 - \delta$, the final convergence result then holds with probability at least $1 - \delta$. Let us first set $t = T$ in (18), and we get

$$\frac{\eta}{1-\beta} \sum_{s=1}^T \frac{\|\bar{\mathbf{g}}_s\|^2}{\|\mathbf{a}_s\|_\infty} \leq \frac{\eta}{1-\beta} \sum_{s=1}^T \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}_s}} \right\|^2 \leq \Delta. \quad (23)$$

Using (12), the basic inequality and Assumption (A3), we have that for any $s \in [T]$, with $B_1 = B + 1$,

$$\begin{aligned} \|\mathbf{a}_s\|_\infty - \epsilon &\leq \max_{i \in [d]} \sqrt{v_{s-1,i} + \mathcal{G}_s^2} = \max_{i \in [d]} \sqrt{\sum_{j=1}^{s-1} g_{j,i}^2 + \mathcal{G}_s^2} \\ &\leq \sqrt{\sum_{j=1}^{s-1} \|\mathbf{g}_j\|^2 + \mathcal{G}_s^2} \leq \sqrt{2 \sum_{j=1}^{s-1} (\|\bar{\mathbf{g}}_j\|^2 + \|\boldsymbol{\xi}_j\|^2) + \mathcal{G}_s^2} \\ &\leq \sqrt{2 \sum_{j=1}^{s-1} (A\Delta_j^{(x)} + B_1 \|\bar{\mathbf{g}}_j\|^2 + C) + \mathcal{G}_s^2}. \end{aligned} \quad (24)$$

Further applying (17) where $\Delta_s^{(x)} \leq \Delta$, $\mathcal{G}_s \leq \mathcal{G}$, $\forall s \in [T]$,

$$\|\mathbf{a}_s\|_\infty - \epsilon \leq \sqrt{2B_1 \sum_{s=1}^T \|\bar{\mathbf{g}}_s\|^2 + 2(A\Delta + C)T + \mathcal{G}^2}.$$

Combining with (23), using $\Delta_1 = \Delta(1 - \beta)/\eta$, then applying Young's inequality,

$$\begin{aligned} &\sum_{s=1}^T \|\bar{\mathbf{g}}_s\|^2 - \Delta_1 \epsilon \\ &\leq \Delta_1 \left(\sqrt{2B_1 \sum_{s=1}^T \|\bar{\mathbf{g}}_s\|^2} + \sqrt{2(A\Delta + C)T} + \mathcal{G} \right) \\ &\leq \sum_{s=1}^T \frac{\|\bar{\mathbf{g}}_s\|^2}{2} + \Delta_1^2 B_1 + \Delta_1 \left(\sqrt{2(A\Delta + C)T} + \mathcal{G} \right). \end{aligned}$$

We then re-arrange the order and divide T on both sides, leading to a desired convergence result

$$\frac{1}{T} \sum_{s=1}^T \|\bar{\mathbf{g}}_s\|^2 \leq 2\Delta_1 \left[\frac{\Delta_1 B_1 + \mathcal{G} + \epsilon}{T} + \sqrt{\frac{2(A\Delta + C)}{T}} \right].$$

The proof is complete. \square

6 CONVERGENCE UNDER GENERALIZED SMOOTHNESS

In this section, we present the convergence of Algorithm 1 in the generalized smooth case.

6.1 GENERALIZED SMOOTHNESS

For a differentiable objective function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we consider the following (L_0, L_1) -smoothness condition: there exist constants $L_0, L_1 > 0$, satisfying that for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ with $\|\mathbf{x} - \mathbf{y}\| \leq 1/L_1$,

$$\|\nabla f(\mathbf{y}) - \nabla f(\mathbf{x})\| \leq (L_0 + L_1 \|\nabla f(\mathbf{x})\|) \|\mathbf{x} - \mathbf{y}\|. \quad (25)$$

The generalized smooth condition was originally put forward by [Zhang et al., 2020b] for any twice-order differentiable function f satisfying that

$$\|\nabla^2 f(\mathbf{x})\| \leq L_0 + L_1 \|\nabla f(\mathbf{x})\|. \quad (26)$$

They revealed the superior of SGD with gradient-clipping in convergence over the vanilla SGD when considering (26). Moreover, empirical evidence has demonstrated that numerous objective functions satisfy (26) while deviating from the global smoothness, particularly in large language models, see e.g., [Zhang et al., 2020b, Figure 1] and [Crawshaw et al., 2022]. To better explain the convergence of gradient-clipping algorithms, Zhang et al. [2020a] provided an alternative form in (25), only requiring f to be first-order differentiable.

The condition in (25) is selected in our paper for three main reasons. First, it's easy to verify that (25) is strictly weaker than L -smoothness. A concrete example is that the simple function $f(x) = x^3$, $x \in \mathbb{R}$ does not satisfy any global smoothness but (25). Second, (25) aligns with the practical limitation to first-order stochastic gradients in our setting, making it more reasonable to assume that f is only first-order differentiable. Finally, both (25) and (26) are shown to be equivalent up to constant factors for twice-order differentiable functions, see [Zhang et al., 2020a, Lemma A.2] and [Faw et al., 2023, Proposition 1]. Thus, (25) includes more functions than (26). We refer interested readers to see [Zhang et al., 2020b,a, Faw et al., 2023] for more discussions and concrete examples of the generalized smoothness.

6.2 CONVERGENCE RESULT

In the following, we establish the convergence bound for AdaGrad with momentum under the generalized smooth condition.

Theorem 2. *Let $T \geq 1$ and $\delta \in (0, 1)$. Suppose that $\{\mathbf{x}_s\}_{s \in [T]}$ is a sequence generated by Algorithm 1, f is (L_0, L_1) -smooth satisfying (25), Assumptions (A1), (A2), (A3) hold, and the parameters satisfy that $\beta \in [0, 1)$,*

$$\epsilon > 0, \quad \eta \leq \min \left\{ C_0, \frac{C_0}{\mathcal{H}}, \frac{C_0}{\mathcal{L}}, \frac{(1 - \beta)^2}{L_1 \sqrt{d}} \right\}, \quad (27)$$

where $C_0 > 0$ is a constant, \mathcal{H}, \mathcal{L} are defined as

$$\begin{aligned} \mathcal{H} &= \sqrt{2A\Lambda_x + 2(B + 1) \left(4L_1\Lambda_x + \sqrt{4L_0\Lambda_x} \right)^2} + 2C, \\ \mathcal{L} &= 2L_0 + 2L_1 \left(4L_1\Lambda_x + \sqrt{4L_0\Lambda_x} \right), \end{aligned} \quad (28)$$

and $\Lambda_y, \tilde{\Lambda}_y, \Lambda_x$ are given with the following order,⁴

$$\Lambda_y \sim \mathcal{O}\left(\Delta_1^{(x)} + \frac{C_0^2 d + C_0 d}{(1-\beta)^3} \log\left(\frac{T}{\delta} + \frac{T}{\epsilon^2}\right)\right),$$

$$\tilde{\Lambda}_y = \frac{2\Lambda_y(1-\beta)}{\eta}, \quad \Lambda_x \sim \mathcal{O}(\Lambda_y^2).$$

Then, with $B_1 = B + 1$, it holds that with probability at least $1 - \delta$,

$$\frac{1}{T} \sum_{s=1}^T \|\nabla f(\mathbf{x}_s)\|^2$$

$$\leq \mathcal{O}\left(\tilde{\Lambda}_y \left(\frac{B_1 \tilde{\Lambda}_y + \sqrt{B_1 L \Lambda_x} + \epsilon}{T} + \sqrt{\frac{A \Lambda_x + C}{T}}\right)\right).$$

It's easy to verify that $\Lambda_y \sim \mathcal{O}(\log(T/\delta))$ and thereby $\Lambda_x, \mathcal{H}, \mathcal{L} \sim \mathcal{O}(\log^2(T/\delta))$. Then, from (27), when $T \gg d$, a typical setting is $\eta \sim \mathcal{O}(\log^2(T/\delta))$. Moreover, the convergence rate is still adaptive to the noise parameters A, C and requires problem-parameters to tune step-sizes potentially due to the relaxation of smoothness. The subsequent result for AdaGrad under (25) could be directly deduced from Theorem 2 and will be presented in Appendix.

7 CONCLUSION

In this paper, we provide high probability convergence bounds for AdaGrad and its momentum variant under the non-convex smooth optimization. In particular, we consider a mild noise model incorporating affine variance noise and the expected smoothness. We rely on a new proxy step-size and some delicate estimations to derive the bound. Our findings reveal that without problem-parameters dependent step-sizes, AdaGrad can find a stationary point with a rate of $\tilde{\mathcal{O}}(1/\sqrt{T})$, particularly accelerating to $\tilde{\mathcal{O}}(1/T)$ when specific noise parameters are sufficiently small. Furthermore, we extend our framework to the generalized smooth case that allows for unbounded smooth parameters, showing the same convergence rate, albeit that problem-parameters dependent step-sizes are required in the latter.

Limitation Although AdaGrad plays an important role in the adaptive method field, several other adaptive methods including RMSProp, Adam and AdamW, may be preferred in some real applications. Therefore, it is also pertinent to study these algorithms under relaxed assumptions. In addition, it is still unknown whether similar convergence result could be also achieved under an expected version of Assumption (A3). Finally, as we study a new assumption over AdaGrad, it would be more beneficial to provide more experimental results to support the theoretical results.

⁴The detailed expressions of Λ_y, Λ_x could be found in (72) and (73) respectively from Appendix.

ACKNOWLEDGMENTS

The authors would like to thank the reviewers and area chairs for their constructive comments. This work was supported in part by the National Key Research and Development Program of China under grant number 2021YFA1003500, and NSFC under grant numbers 11971427. We also thank Chenhao Yu very much on pointing out one error in the generalized-smooth analysis, and Mouxiang Chen for his great help with the experimental results. The corresponding author is Junhong Lin.

References

- Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, 199(1-2):165–214, 2023.
- Amit Attia and Tomer Koren. SGD with AdaGrad step-sizes: full adaptivity with high probability to unknown parameters, unbounded gradients and affine variance. In *International Conference on Machine Learning*, 2023.
- Dimitri P Bertsekas and John N Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- Michael Crawshaw, Mingrui Liu, Francesco Orabona, Wei Zhang, and Zhenxun Zhuang. Robustness to unbounded smoothness of generalized signSGD. In *Advances in Neural Information Processing Systems*, 2022.
- Alexandre Défossez, Leon Bottou, Francis Bach, and Nicolas Usunier. A simple convergence proof of Adam and Adagrad. *Transactions on Machine Learning Research*, 2022.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive sub-gradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7):2121–2159, 2011.
- Matthew Faw, Isidoros Tziotis, Constantine Caramanis, Aryan Mokhtari, Sanjay Shakkottai, and Rachel Ward. The power of adaptivity in SGD: self-tuning step sizes with unbounded gradients and affine variance. In *Conference on Learning Theory*, 2022.
- Matthew Faw, Litu Rout, Constantine Caramanis, and Sanjay Shakkottai. Beyond uniform smoothness: a stopped analysis of adaptive SGD. In *Conference on Learning Theory*, 2023.

- Wayne A Fuller. *Measurement error models*. John Wiley & Sons, 2009.
- Euhanna Ghadimi, Hamid Reza Feyzmahdavian, and Mikael Johansson. Global convergence of the heavy-ball method for convex optimization. In *European Control Conference*, 2015.
- Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Robert M Gower, Peter Richtárik, and Francis Bach. Stochastic quasi-gradient methods: Variance reduction via jacobian sketching. *Mathematical Programming*, 188: 135–192, 2021.
- Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. SGD: General analysis and improved rates. In *International Conference on Machine Learning*, pages 5200–5209. PMLR, 2019.
- Benjamin Grimmer. Convergence rates for deterministic and stochastic subgradient methods without Lipschitz continuity. *SIAM Journal on Optimization*, 29(2):1350–1365, 2019.
- Yusu Hong and Junhong Lin. High probability convergence of Adam under unbounded gradients and affine variance noise. *arXiv preprint arXiv:2311.02000*, 2023.
- Ali Kavis, Kfir Yehuda Levy, and Volkan Cevher. High probability bounds for a class of nonconvex algorithms with AdaGrad stepsize. In *International Conference on Learning Representations*, 2022.
- Ahmed Khaled and Peter Richtárik. Better theory for SGD in the nonconvex world. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.
- Fereshte Khani and Percy Liang. Feature noise induces loss discrepancy across groups. In *International Conference on Machine Learning*, pages 5209–5219. PMLR, 2020.
- Kfir Y Levy, Alp Yurtsever, and Volkan Cevher. Online adaptive methods, universality and acceleration. In *Advances in Neural Information Processing Systems*, 2018.
- Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- Xiaoyu Li and Francesco Orabona. A high probability analysis of adaptive SGD with momentum. In *Workshop on International Conference on Machine Learning*, 2020.
- Zijian Liu, Ta Duy Nguyen, Alina Ene, and Huy Nguyen. On the convergence of AdaGrad(Norm) on \mathbb{R}^d : beyond convexity, non-asymptotic rate and acceleration. In *International Conference on Learning Representations*, 2023a.
- Zijian Liu, Ta Duy Nguyen, Thien Hang Nguyen, Alina Ene, and Huy Nguyen. High probability convergence of stochastic gradient methods. In *International Conference on Machine Learning*, 2023b.
- BT Poljak and Ya Z Tsyppkin. Pseudogradient adaptation and training algorithms. *Automation and Remote Control*, 34:45–67, 1973.
- Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *Ussr Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- Jiang Qian, Yuren Wu, Bojin Zhuang, Shaojun Wang, and Jing Xiao. Understanding gradient clipping in incremental gradient methods. In *International Conference on Artificial Intelligence and Statistics*, 2021.
- Amirhossein Reisizadeh, Haochuan Li, Subhro Das, and Ali Jadbabaie. Variance-reduced clipping for non-convex optimization. *arXiv preprint arXiv:2303.00883*, 2023.
- Peter Richtárik and Martin Takáč. Stochastic reformulations of linear systems: algorithms and convergence theory. *SIAM Journal on Matrix Analysis and Applications*, 41 (2):487–524, 2020.
- Herbert Robbins and Sutton Monroe. A stochastic approximation method. *Annals of Mathematical Statistics*, pages 400–407, 1951.
- Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- Li Shen, Congliang Chen, Fangyu Zou, Zequn Jie, Ju Sun, and Wei Liu. A unified analysis of AdaGrad with weighted aggregation and momentum acceleration. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- Matthew Streeter and H Brendan McMahan. Less regret via online conditioning. *arXiv preprint arXiv:1002.4862*, 2010.
- Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of SGD for over-parameterized models and an accelerated perceptron. In *International Conference on Artificial Intelligence and Statistics*, pages 1195–1204. PMLR, 2019.
- Bohan Wang, Huishuai Zhang, Zhiming Ma, and Wei Chen. Convergence of AdaGrad for non-convex objectives: simple proofs and relaxed assumptions. In *Conference on Learning Theory*, 2023.

- Xiaoyu Wang and Ya-Xiang Yuan. On the convergence of stochastic gradient descent with bandwidth-based step size. *Journal of Machine Learning Research*, 24(48): 1–49, 2023.
- Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: sharp convergence over nonconvex landscapes. *Journal of Machine Learning Research*, 21(1):9047–9076, 2020.
- Huan Xu, Constantine Caramanis, and Shie Mannor. Robust regression and Lasso. In *Advances in Neural Information Processing Systems*, 2008.
- Tianbao Yang, Qihang Lin, and Zhe Li. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. *arXiv preprint arXiv:1604.03257*, 2016.
- Bohang Zhang, Jikai Jin, Cong Fang, and Liwei Wang. Improved analysis of clipping algorithms for non-convex optimization. In *Advances in Neural Information Processing Systems*, 2020a.
- Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: a theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2020b.
- Shen-Yi Zhao, Yin-Peng Xie, and Wu-Jun Li. On the convergence and improvement of stochastic normalized gradient descent. *Science China Information Sciences*, 64:1–13, 2021.
- Dongruo Zhou, Jinghui Chen, Yuan Cao, Yiqi Tang, Ziyang Yang, and Quanquan Gu. On the convergence of adaptive gradient methods for nonconvex optimization. In *Annual Workshop on Optimization for Machine Learning*, 2020.
- Fangyu Zou, Li Shen, Zequn Jie, Weizhong Zhang, and Wei Liu. A sufficient condition for convergences of Adam and RMSProp. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

Revisiting Convergence of AdaGrad with Relaxed Assumptions (Supplementary Material)

Yusu Hong¹

Junhong Lin¹

¹Zhejiang University

Table 1: Comparison of existing results with ours for AdaGrad/AdaGrad-Norm on non-convex smooth case

	Alg. type	Smooth	Noise	Unbounded Gradients	Conv. type
[Li and Orabona, 2019]	Both ^a	L	Sub-Gaussian	-	w.h.p.
[Ward et al., 2020]	Scalar	L	Bounded	-	\mathbb{E}
[Kavis et al., 2022]	Scalar	L	Sub-Gaussian	✓	w.h.p.
[Faw et al., 2022]	Scalar	L	Affine	✓	\mathbb{E}
[Wang et al., 2023]	Both	$L/(L_0, L_1)$	Affine	✓	\mathbb{E}
[Liu et al., 2023b]	Both	L	Coordinate-wise Sub-Gaussian	✓	w.h.p.
[Attia and Koren, 2023]	Scalar	L	Affine	✓	w.h.p.
[Faw et al., 2023]	Scalar	(L_0, L_1)	Affine	✓	\mathbb{E}
This paper, Thm. 1	Coordinate	L	Relaxed Affine	✓	w.h.p.
This paper, Thm. 2	Coordinate	(L_0, L_1)	Relaxed Affine	✓	w.h.p.

In the "Alg. type" column, "Scalar" refers to AdaGrad-Norm, "Coordinate" refers to AdaGrad, and "Both" refers to both algorithms. "Relaxed Affine" corresponds to Assumption (A3) in this paper. In the "Conv. type" column, "w.h.p." stands for the high probability convergence bound, and " \mathbb{E} " represents the expected convergence bound.

^a Li and Orabona [2019] studied a variant of AdaGrad/AdaGrad-Norm using a delayed step-size which is independent from the current stochastic gradient.

A COMPLEMENTARY LEMMAS

Following [Li and Orabona, 2020, Attia and Koren, 2023], we will first present several important technical lemmas. The first lemma is a standard result in the smooth-based optimization which will be used in our analysis motivated also by [Attia and Koren, 2023, Hong and Lin, 2023].

Lemma A.1. *Suppose that f is L -smooth and Assumption (A1) holds. Then for any $\mathbf{x} \in \mathbb{R}^d$,*

$$\|\nabla f(\mathbf{x})\|^2 \leq 2L(f(\mathbf{x}) - f^*).$$

We introduce a concentration inequality for the martingale difference sequence, see [Li and Orabona, 2020] for a proof.

Lemma A.2. *Suppose that $\{Z_s\}_{s \in [T]}$ is a martingale difference sequence with respect to ζ_1, \dots, ζ_T . Assume that for each $s \in [T]$, σ_s is a random variable dependent on $\zeta_1, \dots, \zeta_{s-1}$ and satisfies that*

$$\mathbb{E} \left[\exp \left(\frac{Z_s^2}{\sigma_s^2} \right) \mid \zeta_1, \dots, \zeta_{s-1} \right] \leq e.$$

Then, for any $\lambda > 0$, and for any $\delta \in (0, 1)$, it holds that

$$\mathbb{P} \left(\sum_{s=1}^T Z_s > \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right) + \frac{3}{4} \lambda \sum_{s=1}^T \sigma_s^2 \right) \leq \delta.$$

The following lemma is a commonly used result in the analysis of adaptive methods. See [Levy et al., 2018] for a proof.

Lemma A.3. *Let $\{\alpha_s\}_{s \geq 1}$ be a non-negative sequence. For any $\varepsilon > 0$ and positive integer t ,*

$$\sum_{s=1}^t \frac{\alpha_s}{\varepsilon + \sum_{k=1}^s \alpha_k} \leq \log \left(1 + \frac{1}{\varepsilon} \sum_{s=1}^t \alpha_s \right).$$

B OMITTED PROOFS UNDER SMOOTH CASE

In this section, we provide the missing detailed proofs for some results and lemmas used in the proof of Theorem 1.

Proof of Remark 4.2. Here, we prove the fourth point in Remark 4.2. Let us fix horizon T and denote $\gamma_s = \frac{\|\mathbf{g}_s - \bar{\mathbf{g}}_s\|^2}{A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C}$, $\forall s \in [T]$. Then from the new assumption, we have

$$\mathbb{E}_{\mathbf{z}_s} [\exp(\gamma_s)] \leq e, \quad \text{thus,} \quad \mathbb{E} [\exp(\gamma_s)] \leq e.$$

By Markov's inequality, for any $E \in \mathbb{R}$,

$$\mathbb{P} \left(\max_{s \in [T]} \gamma_s \geq E \right) = \mathbb{P} \left(\exp \left(\max_{s \in [T]} \gamma_s \right) \geq e^E \right) \leq e^{-E} \mathbb{E} \left[\exp \left(\max_{s \in [T]} \gamma_s \right) \right] \leq e^{-E} \mathbb{E} \left[\sum_{s=1}^T \exp(\gamma_s) \right] \leq e^{-E} T e,$$

which leads to that with probability at least $1 - \delta$,

$$\|\boldsymbol{\xi}_s\|^2 = \|\mathbf{g}_s - \bar{\mathbf{g}}_s\|^2 \leq \log \left(\frac{eT}{\delta} \right) \left(A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C \right), \quad \forall s \in [T]. \quad (29)$$

Compared (29) with Assumption (A3), an additional $\log T$ factor emerges. Hence, \mathcal{G}_s and \mathcal{G} defined in (9) should be revised as

$$\mathcal{G}_s = \sqrt{\left(X\Delta_s^{(x)} + 2C \right) \log \left(\frac{eT}{\delta} \right)}, \quad \mathcal{G} = \sqrt{\left(X\Delta + 2C \right) \log \left(\frac{eT}{\delta} \right)}. \quad (30)$$

Consequently, using the similar analysis in Lemma 5.3, Lemma 5.5 and Lemma 5.6, we could reach (19) with a new \mathcal{G}_s and \mathcal{G} in (30). Then, using a similar induction argument, we could deduce the final convergence rate. The additional logarithm factor will make no essential difference to the order of Δ , Δ_1 and the convergence rate in Theorem 1 up to logarithm factors. \square

Proof of Lemma 5.1. Let us denote $\boldsymbol{\eta}_s = \boldsymbol{\eta} / (\sqrt{v_s} + \epsilon)$. Recalling $\mathbf{m}_0 = 0$ and \mathbf{m}_s in Algorithm 1, we have

$$\mathbf{m}_s = \beta \mathbf{m}_{s-1} - \boldsymbol{\eta}_s \odot \mathbf{g}_s = \beta^2 \mathbf{m}_{s-2} - \beta \boldsymbol{\eta}_{s-1} \odot \mathbf{g}_{s-1} - \boldsymbol{\eta}_s \odot \mathbf{g}_s = \cdots = - \sum_{j=1}^s \beta^{s-j} \boldsymbol{\eta}_j \odot \mathbf{g}_j. \quad (31)$$

Note that $|g_{s,i}| \leq \sqrt{v_{s,i}}$, $\forall i \in [d]$. We therefore verify that $\|\mathbf{g}_s / \sqrt{v_s}\| \leq \sqrt{d} \max_{i \in [d]} |g_{s,i} / v_{s,i}| \leq \sqrt{d}$. Moreover,

$$\|\mathbf{m}_s\| \leq \sum_{j=1}^s \beta^{s-j} \|\boldsymbol{\eta}_j \odot \mathbf{g}_j\| \leq \eta \sqrt{d} \sum_{j=1}^s \beta^{s-j} \left\| \frac{\mathbf{g}_j}{\sqrt{v_j}} \right\|_{\infty} \leq \frac{\eta \sqrt{d}}{1 - \beta}. \quad (32)$$

Using the smoothness of f ,

$$\|\bar{\mathbf{g}}_s\| \leq \|\bar{\mathbf{g}}_{s-1}\| + \|\bar{\mathbf{g}}_s - \bar{\mathbf{g}}_{s-1}\| \leq \|\bar{\mathbf{g}}_{s-1}\| + L \|\mathbf{x}_s - \mathbf{x}_{s-1}\| = \|\bar{\mathbf{g}}_{s-1}\| + L \|\mathbf{m}_{s-1}\|.$$

Further using (32),

$$\|\bar{\mathbf{g}}_s\| \leq \|\bar{\mathbf{g}}_{s-1}\| + \frac{L\eta\sqrt{d}}{1-\beta} \leq \|\bar{\mathbf{g}}_1\| + \frac{L\eta s\sqrt{d}}{1-\beta}.$$

□

Proof of Lemma 5.2. Recalling the iteration in Algorithm 1 and then using the descent lemma,

$$f(\mathbf{x}_{s+1}) \leq f(\mathbf{x}_s) + \langle \bar{\mathbf{g}}_s, \mathbf{x}_{s+1} - \mathbf{x}_s \rangle + \frac{L}{2} \|\mathbf{x}_{s+1} - \mathbf{x}_s\|^2 = f(\mathbf{x}_s) + \langle \bar{\mathbf{g}}_s, \mathbf{m}_s \rangle + \frac{L}{2} \|\mathbf{m}_s\|^2.$$

Using Cauchy-Schwarz inequality and Lemma 5.1,

$$\langle \bar{\mathbf{g}}_s, \mathbf{m}_s \rangle \leq \|\bar{\mathbf{g}}_s\| \cdot \|\mathbf{m}_s\| \leq \frac{\eta\sqrt{d}}{1-\beta} \left(\|\bar{\mathbf{g}}_1\| + \frac{L\eta\sqrt{d}s}{1-\beta} \right), \quad \frac{L}{2} \|\mathbf{m}_s\|^2 \leq \frac{L\eta^2 d}{2(1-\beta)^2}.$$

Combining with the above, we obtain that

$$f(\mathbf{x}_{s+1}) \leq f(\mathbf{x}_s) + \frac{\eta\sqrt{d}}{1-\beta} \left(\|\bar{\mathbf{g}}_1\| + \frac{L\eta\sqrt{d}s}{1-\beta} \right) + \frac{L\eta^2 d}{2(1-\beta)^2}.$$

With both sides subtracting f^* and summing up over $s \in [t]$, we obtain that

$$\Delta_{t+1}^{(x)} \leq \Delta_1^{(x)} + \frac{\eta\sqrt{d}}{1-\beta} \sum_{s=1}^t \left(\|\bar{\mathbf{g}}_1\| + \frac{L\eta\sqrt{d}s}{1-\beta} \right) + \frac{L\eta^2 dt}{2(1-\beta)^2}. \quad (33)$$

We define $\sum_a^b = 0$ when $a < b$. Then, we sum up both sides of (33) over $t \in [0, 1, \dots, T-1]$ to obtain that

$$\begin{aligned} \sum_{t=1}^T \Delta_t^{(x)} &\leq \sum_{t=0}^{T-1} \Delta_1^{(x)} + \frac{\eta\sqrt{d}}{1-\beta} \sum_{t=0}^{T-1} \sum_{s=1}^t \left(\|\bar{\mathbf{g}}_1\| + \frac{L\eta\sqrt{d}s}{1-\beta} \right) + \sum_{t=0}^{T-1} \frac{L\eta^2 dt}{2(1-\beta)^2} \\ &\leq \Delta_1^{(x)} T + \frac{\eta\sqrt{d}}{1-\beta} \sum_{t=1}^T \sum_{s=1}^t \left(\|\bar{\mathbf{g}}_1\| + \frac{L\eta\sqrt{d}s}{1-\beta} \right) + \sum_{t=1}^T \frac{L\eta^2 dt}{2(1-\beta)^2} \\ &\leq \Delta_1^{(x)} T + \left(\frac{\eta\|\bar{\mathbf{g}}_1\|\sqrt{d}}{1-\beta} + \frac{L\eta^2 d}{2(1-\beta)^2} \right) T^2 + \frac{L\eta^2 d T^3}{(1-\beta)^2}. \end{aligned}$$

□

Proof of Lemma 5.3. Let $X_s = -\left\langle \bar{\mathbf{g}}_s, \frac{\boldsymbol{\xi}_s}{\mathbf{a}_s} \right\rangle$ for any $s \in [T]$. Note that $\bar{\mathbf{g}}_s$ and \mathbf{a}_s are random variables dependent on $\mathbf{z}_1, \dots, \mathbf{z}_{s-1}$ and $\boldsymbol{\xi}_s$ is dependent on $\mathbf{z}_1, \dots, \mathbf{z}_{s-1}, \mathbf{z}_s$. It is easy to prove that X_s is a martingale difference sequence since

$$\mathbb{E}[X_s \mid \mathbf{z}_1, \dots, \mathbf{z}_{s-1}] = -\left\langle \bar{\mathbf{g}}_s, \frac{\mathbb{E}_{\mathbf{z}_s}[\boldsymbol{\xi}_s]}{\mathbf{a}_s} \right\rangle = 0,$$

where the last equality follows from Assumption (A2). Let

$$\zeta_s = \left\| \frac{\bar{\mathbf{g}}_s}{\mathbf{a}_s} \right\| \sqrt{A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C}, \quad \forall s \in [T].$$

Similarly, ζ_s is a random variable dependent on $\mathbf{z}_1, \dots, \mathbf{z}_{s-1}$. Using Cauchy-Schwarz inequality and Assumption (A3), we have

$$\mathbb{E} \left[\exp \left(\frac{X_s^2}{\zeta_s^2} \right) \mid \mathbf{z}_1, \dots, \mathbf{z}_{s-1} \right] \leq \mathbb{E} \left[\exp \left(\frac{\|\boldsymbol{\xi}_s\|^2}{A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C} \right) \mid \mathbf{z}_1, \dots, \mathbf{z}_{s-1} \right] \leq e.$$

Therefore, given any fixed $t \in [T]$, applying Lemma A.2, we have that for any $\lambda > 0$, with probability at least $1 - \delta$,

$$\begin{aligned} \sum_{s=1}^t X_s^2 &\leq \frac{3\lambda}{4} \sum_{s=1}^t \zeta_s^2 + \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right) = \frac{3\lambda}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2}{a_{s,i}^2} \left(A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C \right) + \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right) \\ &\leq \frac{3\lambda}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2 \mathcal{G}_s^2}{a_{s,i}^2} + \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right) \leq \frac{3\lambda}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2 \mathcal{G}_s}{a_{s,i}} + \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right), \end{aligned} \quad (34)$$

where the second inequality follows from Lemma A.1 and (9). The last inequality follows from using $1/a_{s,i} \leq 1/\mathcal{G}_s$, implied by (12). We then obtain that for any $t \in [T]$, (34) holds with probability at least $1 - \delta$. Therefore, for any fixed $\lambda > 0$, we can re-scale over δ and have that for all $t \in [T]$, with probability at least $1 - \delta$,

$$\sum_{s=1}^t X_s^2 \leq \frac{3\lambda}{4} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\mathbf{a}}_s} \right\|^2 \mathcal{G}_s + \frac{1}{\lambda} \log \left(\frac{T}{\delta} \right). \quad (35)$$

Finally setting $\lambda = 1/(3\mathcal{G})$, we obtain the desired result. \square

Proof of Lemma 5.4. Using the basic inequality, Assumption (A3) and Lemma A.1, we have

$$\|\mathbf{g}_s\|^2 \leq 2\|\bar{\mathbf{g}}_s\|^2 + 2\|\boldsymbol{\xi}_s\|^2 \leq 2A\Delta_s^{(x)} + 2(B+1)\|\bar{\mathbf{g}}_s\|^2 + 2C \leq (2A+4LB+4L)\Delta_s^{(x)} + 2C.$$

Thus, $\|\mathbf{g}_s\| \leq \mathcal{G}_s, \forall s \geq 1$. Let $\mathbf{a}_s = \sqrt{\tilde{\mathbf{v}}_s} + \boldsymbol{\epsilon}$ where $\tilde{\mathbf{v}}_s = (\tilde{v}_{s,i})_i$. Then, for any $i \in [d]$,

$$\begin{aligned} \left| \frac{1}{a_{s,i}} - \frac{1}{b_{s,i}} \right| &= \frac{|\sqrt{\tilde{v}_{s,i}} - \sqrt{v_{s,i}}|}{a_{s,i}b_{s,i}} = \frac{|\tilde{v}_{s,i} - v_{s,i}|}{a_{s,i}b_{s,i}(\sqrt{\tilde{v}_{s,i}} + \sqrt{v_{s,i}})} = \frac{1}{a_{s,i}b_{s,i}} \frac{|\mathcal{G}_s^2 - g_{s,i}^2|}{\sqrt{v_{s,i}} + \sqrt{\tilde{v}_{s,i}}} \\ &\leq \frac{1}{a_{s,i}b_{s,i}} \frac{\mathcal{G}_s^2}{\sqrt{v_{s,i}} + \sqrt{\tilde{v}_{s,i}}} \leq \frac{\mathcal{G}_s}{a_{s,i}b_{s,i}}. \end{aligned} \quad (36)$$

\square

Proof of Lemma 5.5. Applying Lemma 5.4 and Young's inequality,

$$\begin{aligned} \mathbf{A.1.2} &\leq \sum_{s=1}^t \sum_{i=1}^d \left| \frac{1}{a_{s,i}} - \frac{1}{b_{s,i}} \right| |\bar{g}_{s,i}| |g_{s,i}| \leq \sum_{s=1}^t \sum_{i=1}^d \frac{\mathcal{G}_s}{a_{s,i}b_{s,i}} |\bar{g}_{s,i}| |g_{s,i}| \leq \frac{1}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2}{a_{s,i}} + \sum_{s=1}^t \sum_{i=1}^d \frac{\mathcal{G}_s^2}{a_{s,i}} \frac{g_{s,i}^2}{b_{s,i}^2} \\ &\leq \frac{1}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2}{a_{s,i}} + \sum_{s=1}^t \sum_{i=1}^d \mathcal{G}_s \frac{g_{s,i}^2}{b_{s,i}^2}, \end{aligned} \quad (37)$$

where we use $1/a_{s,i} \leq 1/\mathcal{G}_s$ for the last inequality. The proof is complete. \square

Proof of Lemma 5.7. Using the descent lemma of smoothness and (7),

$$\begin{aligned} f(\mathbf{x}_s) &\leq f(\mathbf{y}_s) + \langle \nabla f(\mathbf{y}_s), \mathbf{x}_s - \mathbf{y}_s \rangle + \frac{L}{2} \|\mathbf{x}_s - \mathbf{y}_s\|^2 \\ &= f(\mathbf{y}_s) - \frac{\beta}{1-\beta} \langle \nabla f(\mathbf{y}_s), \mathbf{x}_s - \mathbf{x}_{s-1} \rangle + \frac{L\beta^2}{2(1-\beta)^2} \|\mathbf{x}_s - \mathbf{x}_{s-1}\|^2. \end{aligned}$$

Using Young's inequality and subtracting f^* on both sides,

$$\Delta_s^{(x)} \leq \Delta_s^{(y)} + \frac{1}{2L} \|\nabla f(\mathbf{y}_s)\|^2 + \frac{(L+L)\beta^2}{2(1-\beta)^2} \|\mathbf{x}_s - \mathbf{x}_{s-1}\|^2 \leq 2\Delta_s^{(y)} + \frac{L\|\mathbf{m}_{s-1}\|^2}{(1-\beta)^2},$$

where we apply Lemma A.1 and $\beta \in [0, 1)$. Finally dividing 2 on both sides and re-arranging the order, we obtain the desired result. \square

Proof of Lemma 5.8. Recalling the definition of \mathbf{b}_s and \mathbf{v}_s , and then using Lemma A.3,

$$\sum_{s=1}^t \frac{g_{s,i}^2}{b_{s,i}^2} = \sum_{s=1}^t \frac{g_{s,i}^2}{(\sqrt{v_{s,i}} + \epsilon)^2} \leq \sum_{s=1}^t \frac{g_{s,i}^2}{v_{s,i} + \epsilon^2} \leq \log \left(1 + \frac{1}{\epsilon^2} \sum_{s=1}^t g_{s,i}^2 \right). \quad (38)$$

Using the basic inequality, Assumption (A3), and Lemma A.1,

$$\begin{aligned} \sum_{j=1}^t g_{j,i}^2 &\leq \sum_{j=1}^T g_{j,i}^2 \leq 2 \sum_{j=1}^T (\bar{g}_{j,i}^2 + \xi_{j,i}^2) \leq 2 \sum_{j=1}^T (\|\bar{\mathbf{g}}_j\|^2 + \|\boldsymbol{\xi}_j\|^2) \\ &\leq 2A \sum_{j=1}^T \Delta_j^{(x)} + 2(B+1) \sum_{j=1}^T \|\bar{\mathbf{g}}_j\|^2 + 2CT \leq X \sum_{j=1}^T \Delta_j^{(x)} + 2CT. \end{aligned} \quad (39)$$

Combining with (38) and Lemma 5.2, we obtain that

$$\sum_{s=1}^t \frac{g_{s,i}^2}{b_{s,i}^2} \leq \log \left(1 + \frac{1}{\epsilon^2} \sum_{j=1}^T g_{j,i}^2 \right) \leq \log \left(1 + \frac{1}{\epsilon^2} \left(X \sum_{j=1}^T \Delta_j^{(x)} + 2CT \right) \right) \leq \log \mathcal{F}_T, \quad (40)$$

where \mathcal{F}_T is defined as

$$\mathcal{F}_T := 1 + \frac{1}{\epsilon^2} \left[(\Delta_1^{(x)} X + 2C)T + \left(\frac{\eta \|\bar{\mathbf{g}}_1\| \sqrt{d}}{1-\beta} + \frac{L\eta^2 d}{2(1-\beta)^2} \right) XT^2 + \frac{L\eta^2 d XT^3}{(1-\beta)^2} \right]. \quad (41)$$

Summing up (40) over $i \in [d]$, we prove the first desired result.

Then we move to estimate terms related to \mathbf{m}_s . Let $M_s = \sum_{j=1}^s \beta^{s-j}$. For any $i \in [d]$, recalling (31) and using the convexity of the square function,

$$m_{s,i}^2 = \eta^2 \left(\sum_{j=1}^s \frac{\beta^{s-j} g_{j,i}}{\sqrt{v_{j,i}} + \epsilon} \right)^2 \leq \eta^2 M_s \sum_{j=1}^s \frac{\beta^{s-j} g_{j,i}^2}{(\sqrt{v_{j,i}} + \epsilon)^2} \leq \eta^2 M_s \sum_{j=1}^s \frac{\beta^{s-j} g_{j,i}^2}{v_{j,i} + \epsilon^2}. \quad (42)$$

Summing over $i \in [d]$, using $\beta < 1$ with $M_s \leq \frac{1}{1-\beta}$, Lemma A.3 and $s \leq T$,

$$\|\mathbf{m}_s\|^2 \leq \eta^2 M_s \sum_{i=1}^d \sum_{j=1}^s \frac{g_{j,i}^2}{\sqrt{\sum_{k=1}^j g_{k,i}^2} + \epsilon^2} \leq \frac{\eta^2}{1-\beta} \sum_{i=1}^d \log \left(1 + \frac{1}{\epsilon^2} \sum_{j=1}^s g_{j,i}^2 \right). \quad (43)$$

We also sum up (42) over $s \in [t]$ and apply Lemma A.3 to obtain that

$$\begin{aligned} \sum_{s=1}^t m_{s,i}^2 &\leq \eta^2 \sum_{s=1}^t M_s \sum_{j=1}^s \frac{\beta^{s-j} g_{j,i}^2}{v_{j,i} + \epsilon^2} = \eta^2 \sum_{j=1}^t \frac{g_{j,i}^2}{v_{j,i} + \epsilon^2} \sum_{s=j}^t M_s \beta^{s-j} \\ &\leq \frac{\eta^2}{(1-\beta)^2} \sum_{j=1}^t \frac{g_{j,i}^2}{v_{j,i} + \epsilon^2} \leq \frac{\eta^2}{(1-\beta)^2} \log \left(1 + \frac{1}{\epsilon^2} \sum_{j=1}^t g_{j,i}^2 \right). \end{aligned}$$

Summing over $i \in [d]$, we obtain that

$$\sum_{s=1}^t \|\mathbf{m}_s\|^2 \leq \frac{\eta^2}{(1-\beta)^2} \sum_{i=1}^d \log \left(1 + \frac{1}{\epsilon^2} \sum_{j=1}^t g_{j,i}^2 \right). \quad (44)$$

Finally, we could follow the similar analysis for (39) and (40) to deduce that the terms inside the logarithm operator in (44) could be further bounded by \mathcal{F}_T and thereby verify the target results. \square

C PROOF OF CONVERGENCE UNDER THE GENERALIZED SMOOTHNESS

In this section, we provide the detailed proof of Theorem 2. We still follow all the notations defined in Section 5. We shall first introduce two sequences $\{\mathcal{H}_s\}_{s \geq 1}$ and $\{\mathcal{L}_s\}_{s \geq 1}$ as follows:

$$\mathcal{H}_s = \sqrt{2A\Delta_s^{(x)} + 2(B+1) \left(4L_1\Delta_s^{(x)} + \sqrt{4L_0\Delta_s^{(x)}}\right)^2} + 2C, \quad \mathcal{L}_s = 2L_0 + 2L_1 \left(4L_1\Delta_s^{(x)} + \sqrt{4L_0\Delta_s^{(x)}}\right). \quad (45)$$

As a consequence, we slightly change the proxy step-size \mathbf{a}_s in (12) as

$$\tilde{\mathbf{a}}_s = \sqrt{\mathbf{v}_{s-1} + (\mathcal{H}_s \mathbf{1}_d)^2} + \epsilon, \quad \forall s \in [T]. \quad (46)$$

C.1 CONVERGENCE OF ADAGRAD

As a consequence of Theorem 2, we obtain the following convergence bound for AdaGrad considering affine variance noise and the generalized smoothness.

Corollary 2. *Let $T \geq 1$ and $\delta \in (0, 1)$. Suppose that $\{\mathbf{x}_s\}_{s \in [T]}$ is a sequence generated by Algorithm 1 with $\beta = 0$, f is (L_0, L_1) -smooth satisfying (25), Assumptions (A1), (A2) hold and Assumption (A3) holds with $A = 0$, and the parameters follow the condition in (27) with $\beta = 0$, \mathcal{H}, \mathcal{L} follow the definitions in (28),*

$$\Lambda_y \sim \mathcal{O} \left(\Delta_1^{(x)} + C_0^2 d \log \left(\frac{T}{\delta} + \frac{T}{\epsilon^2} \right) \right), \quad \Lambda_x \sim \mathcal{O} (\Lambda_y^2).$$

Then it holds that with probability at least $1 - \delta$,

$$\frac{1}{T} \sum_{s=1}^T \|\nabla f(\mathbf{x}_s)\|^2 \leq \frac{4\Lambda_y}{\eta} \left(\frac{2\Lambda_y(B+1)/\eta + \mathcal{H} + \epsilon}{T} + \sqrt{\frac{2C}{T}} \right).$$

Remark C.1. Wang et al. [2023] provided a convergence rate for AdaGrad-Norm under the generalized smoothness with the expected affine variance noise, specifically when $\eta < \frac{1}{L_1} \min \left\{ \frac{1}{64B}, \frac{1}{8\sqrt{B}} \right\}$, with probability at least $1 - \delta$,

$$\min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|^2 = \mathcal{O} \left(\frac{\log(\sqrt{CT})}{T\delta^2} + \frac{\sqrt{C} \log(\sqrt{CT})}{\sqrt{T}\delta^2} \right). \quad (47)$$

Thus, our convergence bound in Corollary 2 could reduce to the AdaGrad-Norm case and match the rate in (47) up to logarithm factors, while with a better dependency on the probability margin δ .

C.2 TECHNICAL LEMMAS

We provide an equivalent result in [Zhang et al., 2020a, Lemma A.2], which establishes a different relationship of the gradient norm and the function value gap. We refer to the proof of [Zhang et al., 2020a, Lemma A.2].

Lemma C.1. *Suppose that f is (L_0, L_1) -smooth and Assumption (A1) holds. Then, for any $\mathbf{x} \in \mathbb{R}^d$,*

$$\|\nabla f(\mathbf{x})\| \leq \max \left\{ 4L_1(f(\mathbf{x}) - f^*), \sqrt{4L_0(f(\mathbf{x}) - f^*)} \right\}.$$

We also have the following lemma to ensure the distance of \mathbf{y}_s and \mathbf{x}_s within $1/L_1$ in order to ensure the generalized smoothness in (25).

Lemma C.2. *Let $\mathbf{x}_s, \mathbf{y}_s$ be as in Algorithm 1 and (7). If $\beta \in [0, 1)$, then for any $s \geq 1$,*

$$\max\{\|\mathbf{x}_{s+1} - \mathbf{x}_s\|, \|\mathbf{y}_s - \mathbf{x}_s\|, \|\mathbf{y}_{s+1} - \mathbf{y}_s\|\} \leq \frac{\eta\sqrt{d}}{(1-\beta)^2}. \quad (48)$$

As a consequence, when

$$\eta \leq \frac{(1-\beta)^2}{L_1\sqrt{d}}, \quad \text{then,} \quad \max\{\|\mathbf{x}_{s+1} - \mathbf{x}_s\|, \|\mathbf{y}_s - \mathbf{x}_s\|, \|\mathbf{y}_{s+1} - \mathbf{y}_s\|\} \leq \frac{1}{L_1}, \quad \forall s \geq 1. \quad (49)$$

Proof. Recalling in Lemma 5.1, we have already bounded $\|\mathbf{m}_s\|$ that is independent from smooth-related conditions as follows:

$$\|\mathbf{x}_{s+1} - \mathbf{x}_s\| = \|\mathbf{m}_s\| \leq \frac{\eta\sqrt{d}}{1-\beta}, \quad \forall s \geq 1. \quad (50)$$

Applying the definition of \mathbf{y}_s in (7), (50) and $\beta \in [0, 1)$,¹

$$\|\mathbf{y}_s - \mathbf{x}_s\| = \frac{\beta}{1-\beta} \|\mathbf{x}_s - \mathbf{x}_{s-1}\| \leq \frac{\eta\sqrt{d}}{(1-\beta)^2}, \quad \forall s \geq 1. \quad (51)$$

Using the iteration of \mathbf{y}_s in (8) and Young's inequality

$$\|\mathbf{y}_{s+1} - \mathbf{y}_s\| = \frac{\eta}{1-\beta} \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\| \leq \frac{\eta\sqrt{d}}{1-\beta} \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|_{\infty} \leq \frac{\eta\sqrt{d}}{1-\beta}, \quad \forall s \geq 1, \quad (52)$$

where we apply $|g_{s,i}/b_{s,i}| \leq 1, \forall i \in [d]$ in the last inequality. Combining with (50), (51) and (52), and using $\beta \in [0, 1)$, we then deduce an uniform bound for all three gaps. Finally, letting $\frac{\eta\sqrt{d}}{(1-\beta)^2} \leq \frac{1}{L_1}$, we then prove that (49) holds. \square

Lemma C.3. *Suppose that (49) holds. If f is (L_0, L_1) -smooth, then for any $s \geq 1$,*

$$\begin{aligned} \|\nabla f(\mathbf{x}_{s+1}) - \nabla f(\mathbf{x}_s)\| &\leq \mathcal{L}_s \|\mathbf{x}_{s+1} - \mathbf{x}_s\|, \\ \|\nabla f(\mathbf{y}_s) - \nabla f(\mathbf{x}_s)\| &\leq \mathcal{L}_s \|\mathbf{y}_s - \mathbf{x}_s\|, \\ \|\nabla f(\mathbf{y}_{s+1}) - \nabla f(\mathbf{y}_s)\| &\leq \mathcal{L}_s \|\mathbf{y}_{s+1} - \mathbf{y}_s\|. \end{aligned} \quad (53)$$

As a consequence, for any $s \geq 1$,

$$\begin{aligned} f(\mathbf{y}_{s+1}) - f(\mathbf{y}_s) - \langle \nabla f(\mathbf{y}_s), \mathbf{y}_{s+1} - \mathbf{y}_s \rangle &\leq \frac{\mathcal{L}_s}{2} \|\mathbf{y}_{s+1} - \mathbf{y}_s\|^2, \\ f(\mathbf{x}_s) - f(\mathbf{y}_s) - \langle \nabla f(\mathbf{y}_s), \mathbf{x}_s - \mathbf{y}_s \rangle &\leq \frac{\mathcal{L}_s}{2} \|\mathbf{x}_s - \mathbf{y}_s\|^2, \\ f(\mathbf{x}_{s+1}) - f(\mathbf{x}_s) - \langle \nabla f(\mathbf{x}_s), \mathbf{x}_{s+1} - \mathbf{x}_s \rangle &\leq \frac{\mathcal{L}_s}{2} \|\mathbf{x}_{s+1} - \mathbf{x}_s\|^2. \end{aligned} \quad (54)$$

Proof. Noting that when (49) holds, we could use $\|\mathbf{y}_s - \mathbf{x}_s\| \leq 1/L_1$ and the generalized smoothness in (25) to get that

$$\begin{aligned} \|\nabla f(\mathbf{y}_s)\| &\leq \|\nabla f(\mathbf{x}_s)\| + \|\nabla f(\mathbf{y}_s) - \nabla f(\mathbf{x}_s)\| \\ &\leq \|\nabla f(\mathbf{x}_s)\| + (L_0 + L_1 \|\nabla f(\mathbf{x}_s)\|) \|\mathbf{y}_s - \mathbf{x}_s\| \leq 2\|\nabla f(\mathbf{x}_s)\| + L_0/L_1. \end{aligned}$$

Using Lemma C.1 and combining with \mathcal{L}_s in (45), we have

$$\begin{aligned} L_0 + L_1 \|\nabla f(\mathbf{x}_s)\| &\leq L_0 + L_1 \left(4L_1 \Delta_s^{(x)} + \sqrt{4L_0 \Delta_s^{(x)}} \right) \leq \mathcal{L}_s, \\ L_0 + L_1 \|\nabla f(\mathbf{y}_s)\| &\leq 2L_0 + 2L_1 \|\nabla f(\mathbf{x}_s)\| \leq 2L_0 + 2L_1 \left(4L_1 \Delta_s^{(x)} + \sqrt{4L_0 \Delta_s^{(x)}} \right) \leq \mathcal{L}_s. \end{aligned} \quad (55)$$

Thus, combining with (25), we prove (53). Now based on (53), we could deduce (54). We refer to the proof of [Zhang et al., 2020a, Lemma A.3]. \square

C.3 ROUGH ESTIMATIONS

In generalized smooth cases, we revise the estimations in Lemmas 5.1 and 5.2 as follows.

Lemma C.4. *Suppose that f is (L_0, L_1) -smooth, $\beta \in [0, 1)$ and (49) holds. Then, for any $s \geq 1$,*

$$\|\mathbf{m}_s\| \leq \frac{\eta\sqrt{d}}{1-\beta}, \quad \|\bar{\mathbf{g}}_s\| \leq \|\bar{\mathbf{g}}_1\| + \frac{\eta\sqrt{d}}{1-\beta} \sum_{j=1}^s \mathcal{L}_j.$$

¹The inequality still holds for $s = 1$ since $\mathbf{x}_1 = \mathbf{y}_1$.

Proof. First, the estimation for $\|\mathbf{m}_s\|$ remains unchanged as in Lemma 5.1 since it does not rely on smooth-related conditions. Note that (49) holds. Then using (53), for any $s \geq 2$,

$$\|\bar{\mathbf{g}}_s\| \leq \|\bar{\mathbf{g}}_{s-1}\| + \|\bar{\mathbf{g}}_s - \bar{\mathbf{g}}_{s-1}\| \leq \|\bar{\mathbf{g}}_{s-1}\| + \mathcal{L}_{s-1}\|\mathbf{x}_s - \mathbf{x}_{s-1}\| = \|\bar{\mathbf{g}}_{s-1}\| + \mathcal{L}_{s-1}\|\mathbf{m}_{s-1}\|. \quad (56)$$

Further using (56) and the estimation for $\|\mathbf{m}_s\|$, for any $s \geq 2$,

$$\|\bar{\mathbf{g}}_s\| \leq \|\bar{\mathbf{g}}_{s-1}\| + \frac{\mathcal{L}_{s-1}\eta\sqrt{d}}{1-\beta} \leq \|\bar{\mathbf{g}}_1\| + \frac{\eta\sqrt{d}}{1-\beta} \sum_{j=1}^{s-1} \mathcal{L}_j.$$

Note that the above inequality also holds when $s = 1$. Thus, the proof is complete. \square

Lemma C.5. *Under the same conditions of Lemma C.4, for any $T \geq 1$,*

$$\sum_{t=1}^T \Delta_t^{(x)} \leq \mathcal{I}_T, \quad \mathcal{I}_T := \Delta_1^{(x)}T + \frac{\eta\sqrt{d}}{1-\beta} \sum_{t=1}^T \sum_{s=1}^t \left(\|\bar{\mathbf{g}}_1\| + \frac{\eta\sqrt{d}}{1-\beta} \sum_{j=1}^s \mathcal{L}_j \right) + \frac{\eta^2 d}{2(1-\beta)^2} \sum_{t=1}^T \sum_{s=1}^t \mathcal{L}_s. \quad (57)$$

Proof. Since (49) holds, we could rely on the updated rule in Algorithm 1 and (54) to obtain that

$$f(\mathbf{x}_{s+1}) \leq f(\mathbf{x}_s) + \langle \bar{\mathbf{g}}_s, \mathbf{x}_{s+1} - \mathbf{x}_s \rangle + \frac{\mathcal{L}_s}{2} \|\mathbf{x}_{s+1} - \mathbf{x}_s\|^2 = f(\mathbf{x}_s) + \langle \bar{\mathbf{g}}_s, \mathbf{m}_s \rangle + \frac{\mathcal{L}_s}{2} \|\mathbf{m}_s\|^2. \quad (58)$$

Using Cauchy-Schwarz inequality and Lemma C.4, for any $s \geq 1$,

$$\langle \bar{\mathbf{g}}_s, \mathbf{m}_s \rangle \leq \|\bar{\mathbf{g}}_s\| \|\mathbf{m}_s\| \leq \frac{\eta\sqrt{d}}{1-\beta} \left(\|\bar{\mathbf{g}}_1\| + \frac{\eta\sqrt{d}}{1-\beta} \sum_{j=1}^s \mathcal{L}_j \right), \quad \frac{\mathcal{L}_s}{2} \|\mathbf{m}_s\|^2 \leq \frac{\mathcal{L}_s \eta^2 d}{2(1-\beta)^2}.$$

Combining the above, subtracting f^* on both sides of (58) and summing up over $s \in [t]$, we obtain that for any $t \geq 1$,

$$\Delta_{t+1}^{(x)} \leq \Delta_1^{(x)} + \frac{\eta\sqrt{d}}{1-\beta} \sum_{s=1}^t \left(\|\bar{\mathbf{g}}_1\| + \frac{\eta\sqrt{d}}{1-\beta} \sum_{j=1}^s \mathcal{L}_j \right) + \frac{\eta^2 d}{2(1-\beta)^2} \sum_{s=1}^t \mathcal{L}_s.$$

We define $\sum_a^b = 0$ when $a < b$. Then, we sum up over $t \in [0, 1, \dots, T-1]$ and obtain the desired result. \square

C.4 START POINT AND DECOMPOSITION

To start with, we recall η in (27) and verify that (49) always holds. Then, we could use the descent lemma (54) and apply (8), and sum up both sides over $s \in [t]$ to get that

$$\begin{aligned} f(\mathbf{y}_{t+1}) &\leq f(\mathbf{x}_1) + \sum_{s=1}^t \langle \nabla f(\mathbf{y}_s), \mathbf{y}_{s+1} - \mathbf{y}_s \rangle + \sum_{s=1}^t \frac{\mathcal{L}_s}{2} \|\mathbf{y}_{s+1} - \mathbf{y}_s\|^2 \\ &= f(\mathbf{x}_1) + \frac{\eta}{1-\beta} \cdot \mathbf{A} + \frac{\eta^2}{2(1-\beta)^2} \sum_{s=1}^t \mathcal{L}_s \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2, \end{aligned} \quad (59)$$

where we use $\mathbf{x}_1 = \mathbf{y}_1$ and the definition of \mathbf{A} in (10). We follow the decomposition in (60) and restate as follows,

$$\mathbf{A} = \underbrace{- \sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s, \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\rangle}_{\mathbf{A.1}} + \underbrace{\sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s - \nabla f(\mathbf{y}_s), \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\rangle}_{\mathbf{A.2}}. \quad (60)$$

C.5 ESTIMATING A

We then introduce $\tilde{\mathbf{a}}$ in (46) into (60) to derive that

$$\mathbf{A}.1 = - \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\tilde{\mathbf{a}}_s}} \right\|^2 - \underbrace{\sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s, \frac{\boldsymbol{\xi}_s}{\tilde{\mathbf{a}}_s} \right\rangle}_{\mathbf{A}'\mathbf{.1.1}} + \underbrace{\sum_{s=1}^t \left\langle \bar{\mathbf{g}}_s, \left(\frac{1}{\tilde{\mathbf{a}}_s} - \frac{1}{\mathbf{b}_s} \right) \mathbf{g}_s \right\rangle}_{\mathbf{A}'\mathbf{.1.2}}. \quad (61)$$

The technique for estimating $\mathbf{A}'\mathbf{.1.1}$ is similar to Lemma 5.3. Let $X'_s = - \left\langle \bar{\mathbf{g}}_s, \frac{\boldsymbol{\xi}_s}{\tilde{\mathbf{a}}_s} \right\rangle$. We could still verify that X'_s is a martingale difference sequence and define $\zeta'_s = \left\| \frac{\bar{\mathbf{g}}_s}{\tilde{\mathbf{a}}_s} \right\| \sqrt{A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C}$. Similarly, ζ'_s is a random variable dependent on $\mathbf{z}_1, \dots, \mathbf{z}_{s-1}$. Using Cauchy-Schwarz inequality and Assumption (A3), we have

$$\mathbb{E} \left[\exp \left[\left(\frac{X'_s}{\zeta'_s} \right)^2 \mid \mathbf{z}_1, \dots, \mathbf{z}_{s-1} \right] \leq \mathbb{E} \left[\exp \left(\frac{\|\boldsymbol{\xi}_s\|^2}{A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C} \right) \mid \mathbf{z}_1, \dots, \mathbf{z}_{s-1} \right] \leq e.$$

However, using Lemma C.1 and \mathcal{H}_s in (45), we derive that

$$A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C \leq A\Delta_s^{(x)} + B \left(4L_1\Delta_s^{(x)} + \sqrt{4L_0\Delta_s^{(x)}} \right)^2 + C \leq \mathcal{H}_s^2.$$

Hence, we obtain a different inequality from (34) that

$$\begin{aligned} \mathbf{A}'\mathbf{.1.1} &\leq \frac{3\lambda}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2}{\tilde{a}_{s,i}^2} \left(A\Delta_s^{(x)} + B\|\bar{\mathbf{g}}_s\|^2 + C \right) + \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right) \\ &\leq \frac{3\lambda}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2 \mathcal{H}_s^2}{\tilde{a}_{s,i}^2} + \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right) \leq \frac{3\lambda}{4} \sum_{s=1}^t \sum_{i=1}^d \frac{\bar{g}_{s,i}^2 \mathcal{H}_s}{\tilde{a}_{s,i}} + \frac{1}{\lambda} \log \left(\frac{1}{\delta} \right), \end{aligned} \quad (62)$$

where the last inequality applies $1/\tilde{a}_{s,i} \leq 1/\mathcal{H}_s$ from (46). Then, we can re-scale δ and take $\lambda = 1/(3\mathcal{H})$, leading to the new estimation as follows: with probability at least $1 - \delta$,

$$\mathbf{A}'\mathbf{.1.1} \leq \frac{1}{4} \sum_{s=1}^t \frac{\mathcal{H}_s}{\mathcal{H}} \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\tilde{\mathbf{a}}_s}} \right\|^2 + 3\mathcal{H} \log \left(\frac{T}{\delta} \right), \quad \forall t \in [T]. \quad (63)$$

The estimation for $\mathbf{A}'\mathbf{.1.2}$ remains similar to (15). We first derive from the basic inequality and Assumption (A3) that

$$\|\mathbf{g}_s\|^2 \leq 2\|\bar{\mathbf{g}}_s\|^2 + 2\|\boldsymbol{\xi}_s\|^2 \leq 2A\Delta_s^{(x)} + 2(B+1)\|\bar{\mathbf{g}}_s\|^2 + 2C \leq \mathcal{H}_s^2.$$

Then, based on $\|\mathbf{g}_s\|^2 \leq \mathcal{H}_s^2$, we derive a similar result as in (36) where

$$\left| \frac{1}{\tilde{a}_{s,i}} - \frac{1}{\mathbf{b}_{s,i}} \right| \leq \frac{\mathcal{H}_s}{\tilde{a}_{s,i}\mathbf{b}_{s,i}}, \quad \forall s \in [T], \forall i \in [d].$$

Then, using a similar deduction in (15), we derive that

$$\mathbf{A}'\mathbf{.1.2} \leq \frac{1}{4} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\tilde{\mathbf{a}}_s}} \right\|^2 + \sum_{s=1}^t \mathcal{H}_s \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2. \quad (64)$$

The estimation for $\mathbf{A}.2$ in (60) is revised by the following lemma.

Lemma C.6. *Suppose that f is (L_0, L_1) -smooth and (49) holds. For any $t \geq 1$, if $\beta \in [0, 1)$, it holds that*

$$\mathbf{A}.2 \leq \sum_{s=1}^t \frac{\mathcal{L}_s}{2\eta} \|\mathbf{m}_{s-1}\|^2 + \sum_{s=1}^t \frac{\mathcal{L}_s \eta}{2(1-\beta)^2} \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2. \quad (65)$$

Proof. Noting that when (49) holds, we could rely on (53) and $\beta \in [0, 1)$ to obtain that

$$\|\bar{\mathbf{g}}_s - \nabla f(\mathbf{y}_s)\| \leq \mathcal{L}_s \|\mathbf{y}_s - \mathbf{x}_s\| = \frac{\mathcal{L}_s \beta}{1 - \beta} \|\mathbf{x}_s - \mathbf{x}_{s-1}\| \leq \frac{\mathcal{L}_s}{1 - \beta} \|\mathbf{m}_{s-1}\|. \quad (66)$$

Applying Cauchy-Schwarz inequality and using (66),

$$\mathbf{A.2} \leq \sum_{s=1}^t \|\bar{\mathbf{g}}_s - \nabla f(\mathbf{y}_s)\| \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\| \leq \sum_{s=1}^t \frac{\mathcal{L}_s}{1 - \beta} \|\mathbf{m}_{s-1}\| \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\| \leq \sum_{s=1}^t \frac{\mathcal{L}_s}{2\eta} \|\mathbf{m}_{s-1}\|^2 + \sum_{s=1}^t \frac{\mathcal{L}_s \eta}{2(1 - \beta)^2} \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2. \quad \square$$

C.6 BOUNDING THE FUNCTION VALUE GAP

Based on the above estimations, we are now ready to provide the bound for the function value gap along the optimization trajectory in the following proposition.

Proposition C.1. *Under the same conditions in Theorem 2, for any given $\delta \in (0, 1)$, the following two inequalities hold with probability at least $1 - \delta$,*

$$\Delta_t^{(x)} \leq \Lambda_x, \quad \mathcal{H}_t \leq \mathcal{H}, \quad \mathcal{L}_t \leq \mathcal{L}, \quad \forall t \in [T + 1], \quad (67)$$

and

$$\Delta_{t+1}^{(y)} \leq \Lambda_y - \frac{\eta}{2(1 - \beta)} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\bar{\mathbf{a}}_s}} \right\|^2, \quad \forall t \in [T], \quad (68)$$

where $\mathcal{H}_t, \mathcal{L}_t$ are as in (45) and $\Lambda_x, \Lambda_y, \mathcal{H}, \mathcal{L}$ are as in Theorem 2.

In what follows, we prove Proposition C.1. First, we verify that (49) holds from the setting of η in (27). Hence, we ensure that (59) and (65) hold.

In the following, we will assume the inequality (63) holds and then deduce the results in (67) and (68). Noting that (63) holds with probability at least $1 - \delta$, we thereby deduce that the desired results hold with probability at least $1 - \delta$.

Plugging the estimations (63) and (64) into (61), we obtain the bound for **A.1**. Then, combining with (65), (60) and (59), and subtracting f^* on both sides of (59),

$$\begin{aligned} \Delta_{t+1}^{(y)} &\leq \Delta_1^{(x)} + \frac{\eta}{1 - \beta} \sum_{s=1}^t \left(\frac{\mathcal{H}_s}{4\mathcal{H}} - \frac{3}{4} \right) \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\bar{\mathbf{a}}_s}} \right\|^2 + \frac{3\mathcal{H}\eta}{1 - \beta} \log \left(\frac{T}{\delta} \right) + \frac{\eta}{1 - \beta} \sum_{s=1}^t \mathcal{H}_s \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 \\ &\quad + \sum_{s=1}^t \frac{\mathcal{L}_s}{2(1 - \beta)} \|\mathbf{m}_{s-1}\|^2 + \sum_{s=1}^t \left(\frac{\mathcal{L}_s \eta^2}{2(1 - \beta)^3} + \frac{\mathcal{L}_s \eta^2}{2(1 - \beta)^2} \right) \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2. \end{aligned} \quad (69)$$

We still rely on an induction argument to deduce the result. First, we define two polynomials $\tilde{\mathcal{L}}_t$ and $\tilde{\mathcal{J}}_t$ with respect to t and present the detailed expressions of Λ_y and Λ_x that are independent from t as follows:

$$\tilde{\mathcal{L}}_t := \Delta_1^{(x)} \cdot t + \frac{C_0 \sqrt{d}}{1 - \beta} \left(\|\bar{\mathbf{g}}_1\| + \frac{C_0 \sqrt{dt}}{1 - \beta} \right) \cdot t^2 + \frac{C_0^2 d}{2(1 - \beta)^2} \cdot t^2, \quad (70)$$

$$\tilde{\mathcal{J}}_t := 1 + \frac{1}{\epsilon^2} \left[2A\tilde{\mathcal{L}}_t + 2(B + 1) \left(\|\bar{\mathbf{g}}_1\| + \frac{C_0 \sqrt{dt}}{1 - \beta} \right)^2 \cdot t + 2Ct \right], \quad (71)$$

$$\begin{aligned} \Lambda_y &:= \Delta_1^{(x)} + \frac{3C_0}{1 - \beta} \log \left(\frac{T}{\delta} \right) + \frac{C_0 d}{1 - \beta} \log \tilde{\mathcal{J}}_T, \\ &\quad + \frac{C_0^2 d}{2(1 - \beta)^3} \log \tilde{\mathcal{J}}_T + \left(\frac{C_0^2 d}{2(1 - \beta)^3} + \frac{C_0^2 d}{2(1 - \beta)^2} \right) \log \tilde{\mathcal{J}}_T, \end{aligned} \quad (72)$$

$$\Lambda_x := (2L_0 + 1)\Lambda_y + 8L_1\Lambda_y^2 + \frac{C_0^2 d}{(1 - \beta)^2} \log \tilde{\mathcal{J}}_T. \quad (73)$$

It's worthy noting that $\tilde{\mathcal{L}}_t$ and $\tilde{\mathcal{J}}_t$ are deterministic polynomials with respect to t and are dependent on hyper-parameters C_0, β, d and problem-parameters A, B, C . We could verify that $\Lambda_y \sim \mathcal{O}(\log(T/\delta))$ and $\Lambda_x \sim \mathcal{O}(\log^2(T/\delta))$.

Induction argument The induction begins by noting that $\Delta_1^{(x)} \leq \Lambda_x$ from (72) and (73). Suppose that for some $t \in [T]$,

$$\Delta_s^{(x)} \leq \Lambda_x, \quad \text{thus, } \mathcal{H}_s \leq \mathcal{H}, \quad \mathcal{L}_s \leq \mathcal{L}, \quad \forall s \in [t], \quad (74)$$

where we rely on $\mathcal{H}_s, \mathcal{L}_s$ in (45), and \mathcal{H}, \mathcal{L} in (28). We thus apply (74) to (69) and get

$$\begin{aligned} \Delta_{t+1}^{(y)} &\leq \Delta_1^{(x)} - \frac{\eta}{2(1-\beta)} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\bar{\mathbf{a}}_s}} \right\|^2 + \frac{3\mathcal{H}\eta}{1-\beta} \log\left(\frac{T}{\delta}\right) + \frac{\eta\mathcal{H}}{1-\beta} \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 \\ &\quad + \frac{\mathcal{L}}{2(1-\beta)} \sum_{s=1}^t \|\mathbf{m}_{s-1}\|^2 + \left(\frac{\mathcal{L}\eta^2}{2(1-\beta)^3} + \frac{\mathcal{L}\eta^2}{2(1-\beta)^2} \right) \sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2. \end{aligned} \quad (75)$$

Then, Lemmas 5.8 should be further revised under the generalized smooth condition as follows.

Lemma C.7. *Suppose that f is (L_0, L_1) -smooth and (49) holds. Then for any $t \geq 1$,*

$$\sum_{s=1}^t \left\| \frac{\mathbf{g}_s}{\mathbf{b}_s} \right\|^2 \leq d \log \mathcal{J}_t, \quad \|\mathbf{m}_t\|^2 \leq \frac{\eta^2 d}{1-\beta} \log \mathcal{J}_t, \quad \sum_{s=1}^t \|\mathbf{m}_s\|^2 \leq \frac{\eta^2 d}{(1-\beta)^2} \log \mathcal{J}_t,$$

where \mathcal{J}_t is a polynomial with respect to t with the detailed expression in (76).

Proof. Recall that the three estimations in (38), (43) and (44) remain unchanged as they are not dependent by smooth-related conditions. However, we shall revise the estimation for the term inside the logarithm operator. We also start from the basic inequality and Assumption (A3),

$$\sum_{j=1}^t g_{j,i}^2 \leq 2 \sum_{j=1}^t (\bar{g}_{j,i}^2 + \xi_{j,i}^2) \leq 2 \sum_{j=1}^t (\|\bar{\mathbf{g}}_j\|^2 + \|\xi_j\|^2) \leq 2 \sum_{j=1}^t (A\Delta_j^{(x)} + (B+1)\|\bar{\mathbf{g}}_j\|^2 + C).$$

Then, applying Lemma C.4 and Lemma C.5,

$$1 + \frac{1}{\epsilon^2} \sum_{j=1}^t g_{j,i}^2 \leq \mathcal{J}_t := 1 + \frac{1}{\epsilon^2} \left[2A\mathcal{I}_t + 2(B+1) \sum_{s=1}^t \left(\|\bar{\mathbf{g}}_s\| + \frac{\eta\sqrt{d}}{1-\beta} \sum_{j=1}^s \mathcal{L}_j \right)^2 + 2Ct \right]. \quad (76)$$

Plugging (76) into (38), (43) and (44), we thereby deduce the final result. \square

Therefore, applying Lemma C.7 to (75), we obtain that

$$\begin{aligned} \Delta_{t+1}^{(y)} &\leq \Delta_1^{(x)} - \frac{\eta}{2(1-\beta)} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\bar{\mathbf{a}}_s}} \right\|^2 + \frac{3\mathcal{H}\eta}{1-\beta} \log\left(\frac{T}{\delta}\right) + \frac{\mathcal{H}\eta d}{1-\beta} \log \mathcal{J}_t \\ &\quad + \frac{\mathcal{L}\eta^2 d}{2(1-\beta)^3} \log \mathcal{J}_t + \left(\frac{\mathcal{L}\eta^2 d}{2(1-\beta)^3} + \frac{\mathcal{L}\eta^2 d}{2(1-\beta)^2} \right) \log \mathcal{J}_t. \end{aligned} \quad (77)$$

Using (27), we have that

$$\mathcal{H}\eta \leq C_0, \quad \mathcal{L}\eta \leq C_0, \quad \mathcal{L}\eta^2 \leq C_0^2. \quad (78)$$

We should also note that \mathcal{I}_t in (57) and \mathcal{J}_t in (76) both include $\mathcal{L}_s, s \leq t$. Then recalling $\tilde{\mathcal{I}}_t$ and $\tilde{\mathcal{J}}_t$ in (70) and (71), and applying $\mathcal{L}_s \leq \mathcal{L}, \forall s \leq t$ in (74) and (78), we have that for any $t \in [T]$,

$$\mathcal{I}_t \leq \tilde{\mathcal{I}}_t \leq \tilde{\mathcal{I}}_T, \quad \mathcal{J}_t \leq \tilde{\mathcal{J}}_t \leq \tilde{\mathcal{J}}_T, \quad \log \mathcal{J}_t \leq \log \tilde{\mathcal{J}}_t \leq \log \tilde{\mathcal{J}}_T. \quad (79)$$

Then, applying (78) and (79) to (77), and recalling Λ_y in (72),

$$\Delta_{t+1}^{(y)} \leq \Lambda_y - \frac{\eta}{2(1-\beta)} \sum_{s=1}^t \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\bar{\mathbf{a}}_s}} \right\|^2 \leq \Lambda_y. \quad (80)$$

Finally, we use the following lemma to control $\Delta_s^{(x)}$ by $\Delta_s^{(y)}$.

Lemma C.8. Suppose that f is (L_0, L_q) -smooth and (49) holds. Let \mathbf{y}_s be defined in (7) and $\beta \in [0, 1)$. Then for any $s \geq 1$,

$$\Delta_s^{(x)} \leq (2L_0 + 1)\Delta_s^{(y)} + 8L_1 \left(\Delta_s^{(y)} \right)^2 + \frac{\mathcal{L}_s + 1}{2(1 - \beta)^2} \|\mathbf{m}_{s-1}\|^2.$$

Proof. Since (49) holds, then using (54),

$$\begin{aligned} f(\mathbf{x}_s) &\leq f(\mathbf{y}_s) + \langle \nabla f(\mathbf{y}_s), \mathbf{x}_s - \mathbf{y}_s \rangle + \frac{\mathcal{L}_s}{2} \|\mathbf{x}_s - \mathbf{y}_s\|^2 \\ &= f(\mathbf{y}_s) - \frac{\beta}{1 - \beta} \langle \nabla f(\mathbf{y}_s), \mathbf{x}_s - \mathbf{x}_{s-1} \rangle + \frac{\mathcal{L}_s \beta^2}{2(1 - \beta)^2} \|\mathbf{x}_s - \mathbf{x}_{s-1}\|^2. \end{aligned}$$

Using Young's inequality and subtracting f^* on both sides,

$$\Delta_s^{(x)} \leq \Delta_s^{(y)} + \frac{1}{2} \|\nabla f(\mathbf{y}_s)\|^2 + \frac{(\mathcal{L}_s + 1)\beta^2}{2(1 - \beta)^2} \|\mathbf{x}_s - \mathbf{x}_{s-1}\|^2 \leq \Delta_s^{(y)} + 8L_1 \left(\Delta_s^{(y)} \right)^2 + 2L_0 \Delta_s^{(y)} + \frac{\mathcal{L}_s + 1}{2(1 - \beta)^2} \|\mathbf{m}_{s-1}\|^2,$$

where we apply Lemma C.1 and $\beta \in [0, 1)$ for the last inequality. \square

Using Lemma C.8, (80), and Lemma C.7, we could bound $\Delta_{t+1}^{(x)}$ as

$$\Delta_{t+1}^{(x)} \leq (2L_0 + 1)\Lambda_y + 8L_1\Lambda_y^2 + \frac{(\mathcal{L}_s + 1)\eta^2 d}{2(1 - \beta)^2} \log \mathcal{J}_t \leq (2L_0 + 1)\Lambda_y + 8L_1\Lambda_y^2 + \frac{2C_0^2 d}{2(1 - \beta)^2} \log \tilde{\mathcal{J}}_T,$$

where the last inequality applies (78) and (79). Recalling Λ_x in (73), we find that

$$\Delta_{t+1}^{(x)} \leq \Lambda_x.$$

Combining with (74), the induction is thus complete. We prove the desired result in (67). Finally, as an intermediate result, (68) is verified by (80). The proof of Proposition C.1 is complete.

C.7 PROOF OF THE MAIN RESULT

Based on Proposition C.1, we could prove the final convergence result as follows.

Proof of Theorem 2. The proof for the final convergence rate follows the similar idea and some same estimations in the proof of Theorem 1. Setting $t = T$ in (68), it holds that with probability at least $1 - \delta$,

$$\frac{\eta}{2(1 - \beta)} \sum_{s=1}^T \frac{\|\bar{\mathbf{g}}_s\|^2}{\|\tilde{\mathbf{a}}_s\|_\infty} \leq \frac{\eta}{2(1 - \beta)} \sum_{s=1}^T \left\| \frac{\bar{\mathbf{g}}_s}{\sqrt{\tilde{\mathbf{a}}_s}} \right\|^2 \leq \Lambda_y. \quad (81)$$

In what follows, we assume that (67) and (68) always hold and deduce the convergence bound under these two inequalities. Note that (67) and (68) hold with probability at least $1 - \delta$ according to Proposition C.1. Thus, the final convergence bound also holds with probability at least $1 - \delta$. Applying $\tilde{\mathbf{a}}_s$ in (46) and (67), and following the similar analysis in (24),

$$\begin{aligned} \|\tilde{\mathbf{a}}_s\|_\infty - \epsilon &\leq \sqrt{2 \sum_{j=1}^{s-1} (A\Delta_j^{(x)} + (B + 1)\|\bar{\mathbf{g}}_j\|^2 + C) + \mathcal{H}_s^2} \leq \sqrt{2(B + 1) \sum_{j=1}^{s-1} \|\bar{\mathbf{g}}_j\|^2 + 2(A\Lambda_x + C)s + \mathcal{H}^2} \\ &\leq \sqrt{2(B + 1) \sum_{s=1}^T \|\bar{\mathbf{g}}_s\|^2 + 2(A\Lambda_x + C)T + \mathcal{H}^2}, \quad \forall s \in [T]. \end{aligned}$$

Then combining with (81), and letting $\tilde{\Lambda}_y = \frac{2\Lambda_y(1-\beta)}{\eta}$,

$$\begin{aligned} \sum_{s=1}^T \|\tilde{\mathbf{g}}_s\|^2 &\leq \tilde{\Lambda}_y \left(\sqrt{2(B+1) \sum_{s=1}^T \|\tilde{\mathbf{g}}_s\|^2 + 2(A\Lambda_x + C)T + \mathcal{H}^2 + \epsilon} \right) \\ &\leq \tilde{\Lambda}_y \left(\sqrt{2(B+1) \sum_{s=1}^T \|\tilde{\mathbf{g}}_s\|^2} + \sqrt{2(A\Lambda_x + C)T} + \mathcal{H} + \epsilon \right) \\ &\leq \sum_{s=1}^T \frac{\|\tilde{\mathbf{g}}_s\|^2}{2} + \tilde{\Lambda}_y^2(B+1) + \tilde{\Lambda}_y \left(\sqrt{2(A\Lambda_x + C)T} + \mathcal{H} + \epsilon \right), \end{aligned}$$

where we apply Young's inequality for the last inequality. Re-arranging the order and dividing T on both sides, we get

$$\frac{1}{T} \sum_{s=1}^T \|\tilde{\mathbf{g}}_s\|^2 \leq 2\tilde{\Lambda}_y \left(\frac{\tilde{\Lambda}_y(B+1) + \mathcal{H} + \epsilon}{T} + \sqrt{\frac{2(A\Lambda_x + C)}{T}} \right).$$

□

D EXPERIMENT

In this section, we present an experiment to verify that AdaGrad can find a stationary point when the noise satisfies Assumption (A3).




Figure 1: Training loss vs. steps and gradient norms vs. steps using SGD and AdaGrad.

Experimental Setup We follow the experimental task outlined in [Khaled and Richtárik, 2023], where the objective is to minimize a regularized logistic regression problem defined as follows:

$$\min_{x \in \mathbb{R}^d} \left\{ \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-a_i^\top x)) + \lambda \sum_{j=1}^d \frac{x_j^2}{1 + x_j^2} \right\}. \quad (82)$$

In [Khaled and Richtárik, 2023], it was verified that using uniform sampling over the a9a dataset and the loss function in (82), the noise conforms to Assumption (A3) with $\hat{A} = 10.09$, $\hat{B} = 0$, $\hat{C} = 0.373$, which closely aligns with the theoretical values where $A = 9$, $B = 0$, $C = 0.994$. We then executed both SGD and AdaGrad to minimize (82) using a batch size of 256. The learning rates were set to 7×10^{-6} for SGD and 5×10^{-4} for AdaGrad.

We utilized the a9a dataset and the PyTorch implementations of SGD and AdaGrad. The experiments were conducted on a single NVIDIA GeForce RTX 4090 GPU.

Results We plotted the training loss and gradient norms against the number of steps in Figure 1, training for 1200 epochs. The results indicate that both SGD and AdaGrad can find a stationary point given a sufficiently large number of steps T , thereby supporting the conclusion in Theorem 1.