
Robustness to Unbounded Smoothness of Generalized SignSGD

Michael Crawshaw*

George Mason University
mcrawsha@gmu.edu

Mingrui Liu*

George Mason University
mingruil@gmu.edu

Francesco Orabona*

Boston University
francesco@orabona.com

Wei Zhang*

IBM T. J. Watson Research Center
weiz@us.ibm.com

Zhenxun Zhuang*

Meta Platforms, Inc.
oldboymls@gmail.com

Abstract

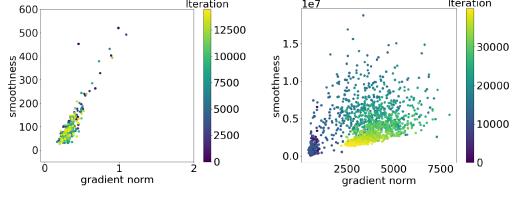
Traditional analyses in non-convex optimization typically rely on the smoothness assumption, namely requiring the gradients to be Lipschitz. However, recent evidence shows that this smoothness condition does not capture the properties of some deep learning objective functions, including the ones involving Recurrent Neural Networks and LSTMs. Instead, they satisfy a much more relaxed condition, with potentially unbounded smoothness. Under this relaxed assumption, it has been theoretically and empirically shown that the gradient-clipped SGD has an advantage over the vanilla one. In this paper, we show that clipping is not indispensable for Adam-type algorithms in tackling such scenarios: we theoretically prove that a generalized SignSGD algorithm can obtain similar convergence rates as SGD with clipping but does not need explicit clipping at all. This family of algorithms on one end recovers SignSGD and on the other end closely resembles the popular Adam algorithm. Our analysis underlines the critical role that momentum plays in analyzing SignSGD-type and Adam-type algorithms: it not only reduces the effects of noise, thus removing the need for large mini-batch in previous analyses of SignSGD-type algorithms, but it also substantially reduces the effects of unbounded smoothness and gradient norms. To the best of our knowledge, this work is the first one showing the benefit of Adam-type algorithms compared with non-adaptive gradient algorithms such as gradient descent in the unbounded smoothness setting. We also compare these algorithms with popular optimizers on a set of deep learning tasks, observing that we can match the performance of Adam while beating others.

1 Introduction

Recent years have witnessed a surge in non-convex machine learning models, with a focus on deep neural networks [27]. DNNs have achieved tremendous progress in a variety of tasks, including computer vision [26, 18, 23], natural language processing [11, 50], and a lot more. Despite their huge empirical success, the theoretical analyses of non-convex optimization [21] prove to be fundamentally more challenging than the established convex optimization theory [4]. Among the numerous literature, many of them assume smoothness of the objective function, namely requiring the gradients to be Lipschitz. Under this scenario, past works have succeeded in proving the convergence rates for a number of algorithms, e.g., Stochastic Gradient Descent [14], AdaGrad [52, 30], and STORM [9, 7].

Nevertheless, it was recently observed that the smoothness assumption does not capture the training of LSTMs [20]: the Hessian can grow with the size of the gradients [56]. Inspired by this, Zhang et

*Authors in alphabetical order.



(a) Wikitext-2 (b) WMT'16 de-en

Figure 1: Local gradient Lipschitz constant vs. Gradient norm on training (a) a 2-layer Transformer Encoder on Wikitext-2 (b) a 6-layer Transformer on WMT'16 Multimodal Machine Translation de-en dataset. The colorbar indicates #Iterations in training. Details in Section 5.2.

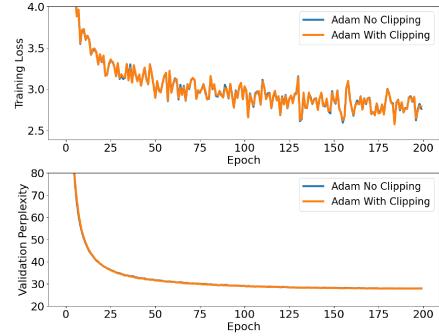


Figure 2: Training GPT-2 on Wikitext-103 using Adam with or without gradient clipping.

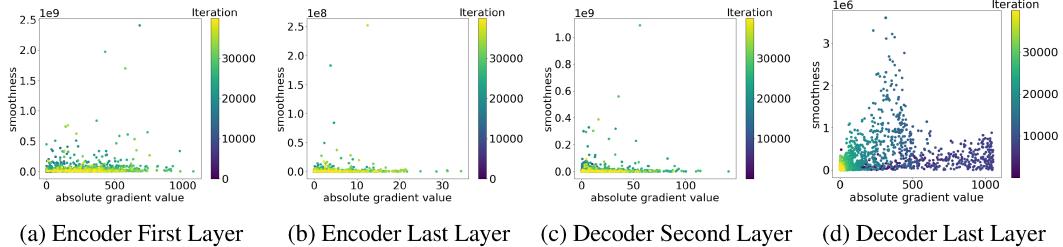


Figure 3: Local gradient Lipschitz constant vs. absolute gradient value on training a Transformer on WMT'16 Multimodal Translation de-en dataset. Each figure represents a randomly picked coordinate in corresponding layers. The colorbar indicates #Iterations during training. Details in Section 5.2.

al. [56] proposed a relaxed smoothness assumption, named (L_0, L_1) smoothness:

$$\|\nabla^2 F(\mathbf{x})\| \leq L_0 + L_1 \|\nabla F(\mathbf{x})\|. \quad (1)$$

They also showed that the well-known gradient clipping technique can ensure Stochastic Gradient Descent's (SGD) convergence in such scenarios. Later, their results were improved to show that SGD with clipping can be made unaffected by the L_1 in (1) and is able to recover the optimal convergence rate of SGD under the original smoothness setting [55, 22].

Nevertheless, the (L_0, L_1) condition has not yet been empirically verified beyond LSTMs. Therefore, our **first contribution** lies in studying the applicability and generalization of the (L_0, L_1) condition. In particular, we have empirically verified that the popular Transformer [50] model also seems to satisfy this assumption, see Figure 1. Yet, we noticed that different coordinates, especially when they are in different layers of the model, exhibit very distinct L_0 and L_1 values as shown in Figure 3. Hence, we propose to refine the (L_0, L_1) assumption in (1) to a coordinate-wise version (Assumption 2) and consider this to better capture the loss surface when training deep neural networks like Transformers.

Given that we assume (a generalization) of the (L_0, L_1) assumption, it would be natural to use some clipping procedure. However, we found out that the use of clipping on Adam [25], while carried out in common practice [e.g., 53], *has no effect on the training and testing performance on optimizing a large transformer model as shown in Figure 2*. In retrospect, this might not be surprising: It is known that Adam has an implicit clipping behavior due to the normalization by the estimated second moment of the gradients. Indeed, Adam can be interpreted as a variant of SignSGD [2].

Inspired by this, our **second contribution** is to propose and analyze a *generalized SignSGD* algorithm under the relaxed smoothness assumption. It is parameterized in such a way that it on one end recovers SignSGD while on the other end closely resembles Adam. Apart from the convergence rates, we also located the critical role the momentum plays in analyzing Adam-type algorithms: it not only reduces the effects of noise but also gives an exponential decaying effect on the unbounded gradient norms and smoothness. This can partly explain the phenomenon that clipping does not help Adam. Moreover, we show a gap between the upper bound of SignSGD and the lower bound of SGD

in the relaxed smoothness setting. This can be considered as a first step in explaining the superior performance of Adam in practical deep learning problems.

The structure of this paper is as follows. Section 2 discusses related works and how our paper builds upon and distinguishes from them. The settings and assumptions are carried out in Section 3. We will introduce formally the generalized SignSGD algorithm and its analysis in Section 4, with a detailed discussion on the bounds and the role of momentum. The experimental results are shown in Section 5, comparing our algorithm with some popular competitors in deep learning tasks. Finally, we draw some conclusions and discuss the limitations of our work in Section 6.

Notations We will use $[d]$ to denote the sequence $[1, 2, \dots, d]$ and use bold letters to represent vectors, e.g., $\mathbf{u} \in \mathbb{R}^d$. The j -th coordinate of a vector \mathbf{u} is u_j . Throughout this paper, we study the Euclidean space \mathbb{R}^d with the inner product $\langle \cdot, \cdot \rangle$. $\mathbb{E}[\mathbf{u}]$ means the expectation with respect to the underlying probability distribution of a random variable \mathbf{u} , and $\mathbb{E}_t[\mathbf{u}]$ is the conditional expectation of \mathbf{u} conditioned on the past of time t . The gradient of F at \mathbf{x} is denoted by $\nabla F(\mathbf{x})$. We use $\mathbb{I}(\cdot)$ to denote the indicator function, $\|\mathbf{u}\|_p$ to denote the p -norm: $\|\mathbf{u}\|_p := (\sum_{j=1}^d |u_j|^p)^{1/p}$ and $\|\mathbf{u}\|_\infty$ the maximum norm: $\|\mathbf{u}\|_\infty := \max\{|u_1|, \dots, |u_d|\}$. We also denote by $\sum_{k=i}^j x_k = 0$ when $i > j$.

2 Related Works

Adaptive Gradient Methods Adaptive gradient methods [34, 12, 25, 19, 44] are popular optimizers for training deep neural networks. The traditional analysis of adaptive gradient methods is providing regret bounds under the online convex optimization framework [12, 25, 44]. Recently, there are some analysis of adaptive gradient methods for nonconvex smooth functions [6, 5, 54, 10, 61]. Zou et al. [60] introduces an intriguing connection between Adam [25] and SignGD [3] when training a two-layer neural network in the deterministic setting, where SignGD is an algorithm following the negative gradient sign direction to perform the update. However, these works cannot be directly extended to nonconvex functions with unbounded smoothness in the stochastic setting. To the best of our knowledge, this work is the first one establishing guarantees for coordinate-wise type optimizers like generalized SignSGD as well as Adam-type updates under a relaxed smoothness condition.

Gradient Clipping The algorithm and analysis of gradient clipping can be traced back to [1, 48, 13] under the assumption that the function is convex and rapidly growing. Hazan et al. [17] considered gradient clipping in quasi-convex optimization. Mai and Johansson [32] showed the stability and convergence of stochastic gradient clipping algorithms for convex problems without the smoothness condition. Gradient clipping is a standard technique in training deep neural networks [39, 40] such as RNNs and LSTMs. The theoretical analysis of gradient clipping for nonconvex models is pioneered by [56], in which the authors analyzed the convergence of gradient clipping under the relaxed smoothness assumption rather than the standard smoothness assumption. Zhang et al. [55] further improved the convergence rate bound under the same assumption as in [56]. Gradient clipping is also used when there is a heavy tail noise in the stochastic gradient to establish high probability convergence rates [8, 15, 57]. Cutkosky and Mehta [7] proved that normalized momentum improves normalized SGD under a second-order smoothness condition. A close algorithm is the one in [22] which employs gradient normalization, momentum, and no gradient clipping to tackle the (L_0, L_1) condition (1) and control noise. Yet, their algorithm normalizes each coordinate with the same scale unlike popular optimizers such as Adam [25]. Moreover, we observe empirically that normalized SGD with momentum performs worse than Adam. Motivated by this, we propose a coordinate-wise optimization algorithm which requires new analysis tools compared with [22].

Employ \mathbf{m}_t^2 to compute \mathbf{v}_t in Adam Designed to combine the advantages of Adagrad [12] and RMSProp [49], the update of Adam [25] employs the ratio between the exponential moving average of the stochastic gradient (\mathbf{m}_t) and the exponential moving average of the squared stochastic gradient (\mathbf{v}_t). Many variants of Adam have been proposed ever since. Among them, one idea is to use \mathbf{m}_t^2 to compute \mathbf{v}_t instead of \mathbf{g}_t^2 . The intuition is that \mathbf{m}_t represents a better update direction than \mathbf{g}_t and can thus better capture the second-moment information. Reddi et al. [43] adopted this change to prove the convergence of Adam in a federated learning setting; yet, they only consider the smooth setting and require a large ϵ to obtain convergence in contrast to the original Adam. Later, Wang et al. [51] explored this idea in more detail, but their analyses are still restricted to the smooth setting. There also exist other variants of the Adam update that attempt to obtain a more stable update changing the order of the normalization and momentum operations [see, e.g., 59].

3 Settings and Preliminaries

In this paper, we focus on the following stochastic optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) := \mathbb{E}_{\xi \sim \mathcal{D}}[f(\mathbf{x}, \xi)],$$

where ξ is a random variable representing a randomly selected data sample or random noise following an unknown distribution \mathcal{D} . We will use the following assumptions.

Assumption 1. $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is differentiable and bounded from below with infimum F^* .

Assumption 2. We say that a differentiable function $F(\mathbf{x})$ is $(\mathbf{L}_0, \mathbf{L}_1)$ -smooth coordinate-wisely, if for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ for which $\|\mathbf{x} - \mathbf{y}\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty}$, we have for any $j \in [d]$ that

$$\left| \frac{\partial F}{\partial x_j}(\mathbf{y}) - \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \leq \left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \right) \|\mathbf{y} - \mathbf{x}\|_2. \quad (2)$$

We will denote $\mathbf{L}_0 := [L_{0,1}, L_{0,2}, \dots, L_{0,d}]^T$ and $\mathbf{L}_1 := [L_{1,1}, L_{1,2}, \dots, L_{1,d}]^T$.

The original (L_0, L_1) smoothness assumption (1) in [56] was proposed as a generalization of the more common smoothness assumption, which says that the gradient should be Lipschitz. Indeed, when $L_{1,j}$ are zero, we recover the smoothness assumption. In contrast, when $L_{1,j}$ are non-zero, the smoothness of the function is potentially *unbounded*. Yet, [56] works with norms and applies to the global scale, while ours is more fine-grained and applies to each coordinate separately. One motivation for this assumption comes from [Remark 2.3, 55] where they noted that (1) can be relaxed to an assumption on gradient differences: there exists $K_0, K_1 > 0$ such that

$$\|\nabla F(\mathbf{x}) - \nabla F(\mathbf{y})\|_2 \leq (K_0 + K_1 \|\nabla F(\mathbf{x})\|_2) \|\mathbf{x} - \mathbf{y}\|_2, \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{y}\|_2 \leq 1/K_1. \quad (3)$$

Indeed, our Assumption 2 implies (3) when $L_{0,j} = L_0$ and $L_{1,j} = L_1$ for all $j \in [d]$, up to constants (See Lemma 3 in the Appendix). Note that the $\frac{1}{\sqrt{d}}$ factor in ours is exactly for easy comparison with (3). The reason we turn to the current coordinate-wise version is that we observed a vast variance across different layers in training Transformer models: (1) is still true globally (Figure 1), but each layer or even each coordinate satisfies has a very different (L_0, L_1) pair (Figure 3). The smoothness assumption has been generalized in orthogonal directions in other work [45, 3, 24].

One merit of Assumption 2 is that it gives us the following descent lemma.

Lemma 1. Let F be $(\mathbf{L}_0, \mathbf{L}_1)$ -smooth coordinate-wisely. Then, for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ for which $\|\mathbf{x} - \mathbf{y}\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty}$, we have

$$F(\mathbf{y}) \leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \sum_{j=1}^d \frac{1}{2} \left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \right) \|\mathbf{y} - \mathbf{x}\|_2 |y_j - x_j|.$$

Our last assumption is common in the literature studying the (L_0, L_1) smooth condition [56, 55].

Assumption 3. For each $j \in [d]$, there exists $\sigma_j > 0$ such that for all $\mathbf{x} \in \mathbb{R}^d$ and $\xi \sim \mathcal{D}$, the noise satisfies $\left| [\nabla f(\mathbf{x}, \xi)]_j - \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \leq \sigma_j$ with probability 1. We will denote $\boldsymbol{\sigma} := [\sigma_1, \sigma_2, \dots, \sigma_d]^T$.

4 A Generalized SignSGD Algorithm

Algorithm 1 Generalized SignSGD (All operations on vectors are element-wise.)

- 1: Inputs: $\mathbf{x}_1, \beta_1, \beta_2, \eta$
- 2: $\mathbf{m}_0 = 0, \mathbf{v}_0 = 0$
- 3: **for** $t = 1, \dots, T$ **do**
- 4: Compute an unbiased estimate $\nabla f(\mathbf{x}_t, \xi_t)$ of $\nabla F(\mathbf{x}_t)$, denoted as \mathbf{g}_t
- 5: $\mathbf{m}_t = \beta_1 \mathbf{m}_{t-1} + (1 - \beta_1) \mathbf{g}_t$
- 6: $\mathbf{v}_t = \beta_2 \mathbf{v}_{t-1} + (1 - \beta_2) \mathbf{m}_t^2$
- 7: $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \frac{\mathbf{m}_t}{\sqrt{\mathbf{v}_t}}$
- 8: **end for**

In this section, we present in Algorithm 1 a generalized SignSGD algorithm. This algorithm encompasses a variety of optimization algorithms.

At first sight, it seems very similar to Adam. Indeed, if we employ \mathbf{g}_t^2 in computing \mathbf{v}_t instead of \mathbf{m}_t^2 , then it is exactly Adam, except for the bias correction terms. We would like to clarify that the idea of this change has been explored before, as detailed in Section 2. In this paper, the motivation for adopting this idea comes from the known effect of momentum on reducing the influence of noises [7]. Indeed, in our analysis the difference between \mathbf{m}_t and $\nabla F(\mathbf{x}_t)$ is much more controllable than between \mathbf{g}_t and $\nabla F(\mathbf{x}_t)$. Thus, we consider employing \mathbf{m}_t in computing \mathbf{v}_t a better choice.

On the other end, the careful reader might observe that Algorithm 1 recovers the SignSGD with Momentum algorithm, also called SIGNUM in [3], when setting $\beta_2 = 0$. Sign-based algorithms are naturally suited to distributed learning [29] and the idea dated back to at least RPROP [46]. The convergence to a stationary point (with ℓ_1 norm) under a coordinate-wise smoothness condition has been established for SignSGD with/without the momentum in [3] though they necessitate large mini-batches to control the variance of the noise. Yet, we are more interested in their property of the update size being bounded without the need for explicit clipping.

Note that both SignSGD and Adam are good candidates for optimization algorithms whose update must be bounded on functions that satisfy the $(\mathbf{L}_0, \mathbf{L}_1)$ condition. Indeed, SignSGD can be seen as an extreme form of gradient clipping. On the other hand, as said in the introduction, Adam does not seem to require gradient clipping at all when used to train the large Transformer model in Figure 2.

Hence, we expect our algorithm, a generalization of SignSGD and a close resemblance to Adam, can enjoy the merits of both and be robust to the unbounded smoothness in the $(\mathbf{L}_0, \mathbf{L}_1)$ scenario. In the next section, we will formalize this claim by presenting the theoretical analysis of Algorithm 1.

4.1 Theoretical Convergence Analysis

Theorem 1. *Under Assumptions 1, 2, and 3, assume $M_j := \sup \left\{ \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| : F(\mathbf{x}) \leq F(\mathbf{x}_1) \right\}$ is finite for each $j \in [d]$, let Δ be any upper bound on $F(\mathbf{x}_1) - F^*$, $\alpha = \min \left(\frac{\sqrt{\|\mathbf{L}_0\|_1} \sqrt{\Delta}}{\|\boldsymbol{\sigma}\|_1 \sqrt{T}}, 1 \right)$, $\beta_1 = 1 - \alpha$, $\frac{\sqrt{\beta_2}}{\beta_1} < 1$, $\rho = 1 - \frac{\sqrt{\beta_2}}{\beta_1}$, $\eta = \frac{\sqrt{\Delta \alpha}}{\sqrt{\|\mathbf{L}_0\|_1} \sqrt{T}}$, for $T \geq \max \left(\frac{100d\Delta \|\mathbf{L}_1\|_\infty^2}{(1-\beta_2)\rho^2 \|\mathbf{L}_0\|_1}, \frac{10000d^2\Delta \|\boldsymbol{\sigma}\|_1^2 \|\mathbf{L}_1\|_\infty^4}{(1-\beta_2)^2 \rho^4 \|\mathbf{L}_0\|_1^3} \right)$, Algorithm 1 guarantees, with probability at least $1 - \delta$, that*

$$\begin{aligned} \min_{t \in [T]} \|\nabla F(\mathbf{x}_t)\|_1 &= \mathcal{O} \left(\frac{\sqrt{\log(dT/\delta)} \|\mathbf{L}_0\|_1^{1/4} \Delta^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2}}{\rho \sqrt{1 - \beta_2} T^{1/4}} + \frac{\log(dT/\delta) \sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\rho \sqrt{T}} \right) \\ &\quad + \mathcal{O} \left(\frac{\|\mathbf{M}\|_1 + \|\boldsymbol{\sigma}\|_1}{\rho} \exp \left(-\frac{\sqrt{1 - \beta_2} \|\mathbf{L}_0\|_1^{3/4}}{\sqrt{d} \|\mathbf{L}_1\|_\infty \|\boldsymbol{\sigma}\|_1^{1/2} \Delta^{1/4}} T^{1/4} \right) + \frac{\|\nabla F(\mathbf{x}_1)\|_1}{T} \right). \end{aligned}$$

Furthermore, for the case when $\beta_2 = 0$, we have the following refined guarantee:

$$\begin{aligned} \min_{t \in [T]} \|\nabla F(\mathbf{x}_t)\|_1 &= \mathcal{O} \left(\frac{\sqrt{\log(dT/\delta)} \|\mathbf{L}_0\|_1^{1/4} \Delta^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2}}{T^{1/4}} + \frac{\log(dT/\delta) \sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\sqrt{T}} \right) \\ &\quad + \mathcal{O} \left(\frac{\|\nabla F(\mathbf{x}_1)\|_1}{\sqrt{T}} \left(\frac{1}{\sqrt{T}} + \frac{\|\boldsymbol{\sigma}\|_1}{\sqrt{\|\mathbf{L}_0\|_1 \Delta}} \right) + \frac{\|\boldsymbol{\sigma}\|_1}{T} \right). \end{aligned}$$

Here, M_j denotes the maximum absolute value of the partial derivative of F for coordinate j among the sub-level set of $F(\mathbf{x}_1)$, namely any point \mathbf{x} with $F(\mathbf{x}) \leq F(\mathbf{x}_1)$. In other words, we assume gradients to be bounded in the sub-level set of $F(\mathbf{x}_1)$; yet, we do not make any restriction on gradients outside of this set. We believe this is not a strong assumption, for example, when the sub-level set of $F(\mathbf{x}_1)$ is bounded, then by the assumed continuity of gradients it trivially holds. Also, we just require an upper bound and it can even be exponentially large as we have an exponentially decaying coefficient to counteract it: notice how the term $\|\mathbf{M}\|_1$ is multiplied by a term that decays exponentially with T . Better still, when $\beta_2 = 0$, we no longer even need this assumption and the algorithm is entirely free of the influence of $\|\mathbf{M}\|_1$. To see why this is good, we show a refined lower bound of Gradient Descent under the relaxed smoothness scenario below which is originally in [56].

Theorem 2. Fix $\epsilon > 0, L_0 > 0, L_1 > 0, M \geq \max(\frac{L_0}{L_1}, \epsilon)$, and $x_0 \in \mathbb{R}$. Pick any constant learning rate η for GD, with the knowledge of the above constants. Then, there exists a 1-d (L_0, L_1) -smooth function, bounded from below by f^* (finite), and such that $\sup\{|f'(x)| : f(x) \leq f(x_0)\} \leq M$ on which the number of iterations T of GD with learning rate η to guarantee $|f'(x_T)| < \epsilon$ is at least

$$\frac{ML_1(f(x_0) - f^* - \frac{15\epsilon^2}{16L_0})}{2\epsilon^2 \left(\ln \frac{ML_1}{L_0} + 1 \right)}.$$

Theorem 2 shows that in the relaxed smoothness setting, GD with any constant step size will suffer from a linear term depending on $L_1 M$. On a side note, it is a fixed version of the lower bound in [56]: we provide in Appendix an explanation of errors in their lower bound and our corrected proof.

Compared with GD, our algorithm only has an exponentially decaying dependence on $L_1 M$. We consider this to be substantial merit of our algorithm. Furthermore, when $\beta_2 = 0$ in which case we recover the SignSGD with Momentum algorithm, we can even show that it completely removes the effects of the unbounded gradient norms. Also notice that in such case we actually no longer need the assumption of $M_j := \sup \left\{ \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| : F(\mathbf{x}) \leq F(\mathbf{x}_1) \right\}$ being finite for each $j \in [d]$ anymore, and the $\|\mathbf{L}_1\|_\infty$ term does not appear in the final bound anymore.

We also would like to point out that this bound closely resembles the one achieved by SGD with gradient clipping algorithm [55] except that we consider the coordinate-wise setting: take the setting of $\beta_2 = 0$ for example, we need at most $\mathcal{O} \left(\Delta \max \left\{ \frac{\|\sigma\|_1^2 \|\mathbf{L}_0\|_1}{\epsilon^4}, \frac{d^2 \|\sigma\|_1^2 \|\mathbf{L}_1\|_\infty^4}{\|\mathbf{L}_0\|_1^3}, \frac{d \|\mathbf{L}_1\|_\infty^2}{\|\mathbf{L}_0\|_1} \right\} \right)$ to get a point \mathbf{x} with $\|\nabla F(\mathbf{x})\|_1 \leq \epsilon$ with high probability.

Remark 1 The almost surely bounded assumption 3 can be relaxed to sub-gaussian noise, using standard extensions of Freedman inequality [e.g., 16].

Remark 2 When $\beta_2 = 0$, we can prove an average-iterate complexity bound (see Proof of Theorem 1 for $\beta_2 = 0$ in Appendix A.3); yet, we use the min form for consistency between the two cases.

Remark 3 Our bound is incomparable with the one in [55, Theorem 3.2]. Yet, as we said, if $L_{0,j} = L_0$ and $L_{1,j} = L_1$ for all $j \in [d]$, then the function satisfies (3). In this case, assuming the noise vector and the gradient vector to be dense to be able to compare the ℓ_1 -norm and the ℓ_2 -norm, we recover the same bound of [55, Theorem 3.2] in terms of dependencies on L_1, L_0 , and T . Instead, in the more general case when $L_{0,j}$ and $L_{1,j}$ are not uniform vectors, our bound allows a finer control of the unbounded smoothness.

Remark 4 Careful readers might be concerned on the relations between $\alpha, \beta_1, \beta_2, \rho$, and T when $\alpha \neq 1$. We would like to note that, when β_2 is fixed, α is inversely proportional to \sqrt{T} . In turn, the definition of ρ means that as T grows, ρ grows and approaches $1 - \sqrt{\beta_2}$. Thus, the two conditions for T decreases when T grows. This means that there must exists a threshold of T above which the two conditions on T always hold. In summary, Theorem 1 conveys the same message as [55] that as long as the expected ϵ is sufficiently small, the complexity no longer has a dependency on \mathbf{L}_1 .

The proof of the theorem is highly technical and it uses recent advancements in the analysis of momentum methods [7], key techniques to deal with the (L_0, L_1) assumption [55], as well as a novel and essential inductive argument to control the norm of past gradients. We want to stress that the difficulty mainly comes from analyzing Adam-type updates when $\beta_2 > 0$, while for the other case of $\beta_2 = 0$ the proof is significantly simpler. The full proof is in the Appendix, but here we present a proof sketch that underlines the main steps. First, we list some key lemmas we used but move their proofs to the appendix due to space constraints.

Lemma 5. With notations in Algorithm 1, for $\tau \leq \bar{\tau} = \frac{\sqrt{1-\beta_2}}{\eta\sqrt{d}\|\mathbf{L}_1\|_\infty}$, we have $\|\mathbf{x}_{t-\tau} - \mathbf{x}_t\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty}$.

Lemma 5 limits our focus to the most recent $\bar{\tau}$ steps on which Assumption 2 and Lemma 1 can apply.

Lemma 7. Assume Assumption 3. With the notation of Algorithm 1, let $j \in [d]$ and $\beta_1 \leq 1$. Then, with probability at least $1 - 3\delta$, for any $t_0 \in [t]$, we have

$$\left| \sum_{\tau=1}^{t_0} \beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right| \leq 3\sigma_j \max(1, \log(1/\delta)) + \frac{3}{\sqrt{1-\beta_1^2}} \sqrt{\sigma_j^2 \max(1, \log(1/\delta))} \triangleq E_j.$$

Lemma 7 is the major tool we use to handle the noise we incur during drawing stochastic gradients. It is derived based on Lemma 12 in [8].

Lemma 10. *With the notation of Algorithm 1 and under the assumptions of Theorem 1, if $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right| \leq M_j$ holds for all $\tau \leq t$ and $j \in [d]$, and $D > 0$, then, with probability at least $1 - 3t\delta$ we have that,*

$$\text{either } \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| < \frac{5B_j}{D} \text{ or } \frac{|m_{t,j}|}{\sqrt{v_{t,j}}} \geq \frac{\rho D}{5\sqrt{1 - \beta_2}},$$

where $B_j \triangleq \frac{\eta L_{0,j}}{\sqrt{1 - \beta_2}(1 - \beta_1)} + \beta_1^\tau (M_j + \sigma_j) + (1 - \beta_1)E_j$ and $D \triangleq 1 - \frac{2\eta\sqrt{d}\|\mathbf{L}_1\|_\infty}{\sqrt{1 - \beta_2}(1 - \beta_1)}$.

Lemma 10 is similar to Lemma A.2 in [60] which considered the deterministic and smooth setting; in contrast, our proof is much more challenging in that we need to tackle both the noise and the unbounded smoothness. With this lemma, we know that either the true gradient is small or that the update of our Algorithm 1 can be lower bounded.

Lemma 12. *Under Assumptions 1, 2, and 3, using the hyperparameters in Theorem 1, denoting $\alpha = 1 - \beta_1$ and $\epsilon_t = \mathbf{m}_t - \nabla F(\mathbf{x}_t)$, for all t and $j \in [d]$ we have, with probability at least $1 - 3\delta$,*

$$|\epsilon_{t+1,j}| \leq (1 - \alpha)^t \left(\alpha\sigma_j + (1 - \alpha) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| \right) + \frac{\eta L_{0,j}}{\alpha} + \alpha E_j + (1 - \alpha)\eta\sqrt{d}L_{1,j} \sum_{\tau=0}^{t-1} (1 - \alpha)^\tau \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t-\tau}) \right|.$$

Lemma 12 shows how the use of momentum can help control the noise by choosing β_1 wisely. It is adapted from the proof of Theorem 2 in [8] but with the added difficulty of unbounded smoothness.

Proof sketch of Theorem 1. Observing the formula of setting β_1 , we can see that when $\|\boldsymbol{\sigma}\|_1 \leq \sqrt{\|\mathbf{L}_0\|_1\Delta}/\sqrt{T}$, $\beta_1 = 0$. As $\beta_2 < \beta_1$, Algorithm 1 reduces to SignSGD. In this case, the key component is Lemma 12 using which we are able to show that $\sum_{t=1}^T \left| m_{t,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|$ can be controlled as $C_1 \sum_{t=1}^T \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + C_2$. The summation of true gradients over time can then be offsetted by choosing η and β_1 wisely when we invoke the descent lemma 1. The rest is standard.

Now for the other case in which $\|\boldsymbol{\sigma}\|_1 > \sqrt{\|\mathbf{L}_0\|_1\Delta}/\sqrt{T}$, we take a different route.

First, notice that Assumption 2 and the Descent Lemma 1 only hold when two points are not too far away. Thus, we need to restrict our attention to the recent updates (Lemma 5), beyond which we would have no control. This means we want the influence of those updates too long ago to not have a big effect on the current one. To make this happen, one natural idea is to use a bounded gradient assumption, then with the use of exponential averaging, their effect would be quickly reduced. Yet, assuming directly that all gradients are bounded would trivialize the $(\mathbf{L}_0, \mathbf{L}_1)$ assumption. Thus, we pose a much weaker condition, assuming that $M_j := \sup \left\{ \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| : F(\mathbf{x}) \leq F(\mathbf{x}_1) \right\}$ being finite for each $j \in [d]$. Then, we prove that M_j will provide an upper bound to all the true gradients the algorithm see. We prove it using induction, analyzing separately the case that either the true gradient is already very small and we have reached the proximity of a stationary point, or the objective function is monotonically non-increasing and the gradient remains bounded.

Having controlled the past gradients, we prove in Lemma 10 that the update of Algorithm 1 is either very small that we can pass or having a constant lower bound that we can use in the Descent Lemma 1.

Also, considering that this is the stochastic setting, noise typically slows down convergence or can even cause the algorithm to diverge if the hyperparameters are not chosen wisely. To handle this, we invoke Freedman's inequality to show that the addition of adjacent stochastic noise almost cancels out each other and the absolute value of the sum remains controlled (Lemma 7).

Yet, we still need another block to handle the difference between the true gradient and the momentum as we are updating in the direction of the momentum instead of the true gradient. Turns our that we can prove that $\text{sign}(m_{t,j}) = \text{sign} \left(\frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right)$ when $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|$ is not too small. As before, in the case $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|$ is small, we have converged on that coordinate. Combining all these blocks together, we are able to arrive at the final results. \square

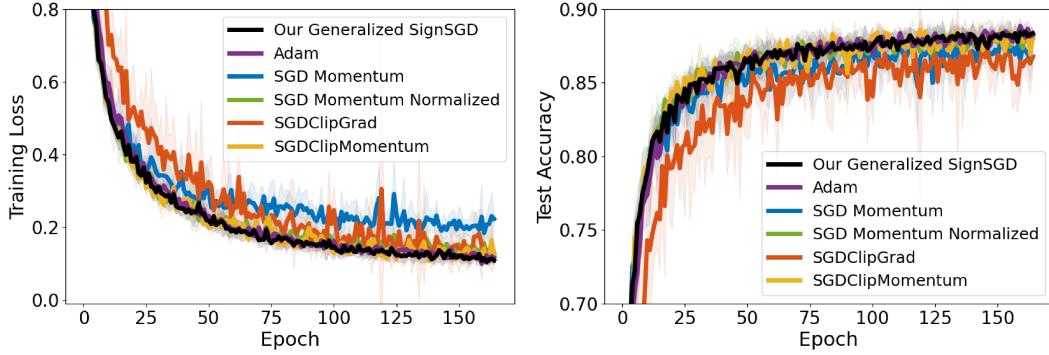


Figure 4: Training a 20-layer Resnet on CIFAR10. The shading of each curve represents the 95% confidence interval computed across 5 independent runs from different random seeds.

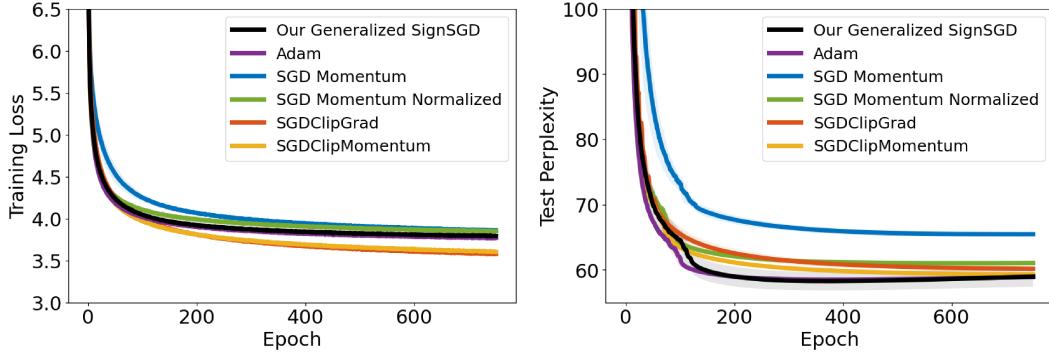


Figure 5: Training an AWD-LSTM to do language modeling (word level) on Penn Treebank. The shading of each curve represents the 95% confidence interval over 5 independent runs.

5 Experiments

We conducted our experiments using PyTorch [41] on Nvidia V100 GPUs. Codes can be found at <https://github.com/zhenxun-zhuang/Generalized-SignSGD>.

5.1 Comparison with Other Optimizers

To validate the efficacy of our Algorithm 1, we compare it with Adam [25], SGD [47], SGD Momentum Normalized [22], SGDClipGrad, and SGDClipMomentum. The latter two are from Algorithm 1 in [55] where SGDClipGrad corresponds to the case when $\nu = 0$ and SGDClipMomentum corresponds to when $\nu = 1$.

Training Unless otherwise specified, we use grid-search to fine-tune the initial learning rate for all optimizers, as well as the clipping threshold for SGDClipGrad and SGDClipMomentum, and β_2 for Adam and our algorithm, to select the one giving the best validation performance on a separated validation set. We then employ the best performing hyperparameters to train the model over all training data and report the testing performance. The testing is repeated with random seeds 5 times to eliminate the influence of stochasticity. For more details, please refer to Section A.4.

Resnet for Image Classification on CIFAR-10 We employ the 20-layer Residual Network model [18] to do image classification on the CIFAR-10 dataset. Images are normalized per channel using the means and standard deviations computed from all training images. We adopt the data augmentation technique following [28] (for training only): 4 pixels are padded on each side of an image and a 32×32 crop is randomly sampled from the padded image or its horizontal flip. The mini-batch size is 128 and we train all algorithms for 164 epochs. We do not employ any learning rate decay schedule in order to focus on the comparison of the optimizers themselves. We fixed the weight decay value to be

Table 1: Average final training loss and test accuracy achieved by each method when optimizing respective models on each dataset. The \pm shows 95% confidence intervals of the mean loss/accuracy/perplexity value over 5 runs starting from different random seeds.

Methods	CIFAR10		Penn Treebank	
	Training loss	Test accuracy	Training loss	Test perplexity
SGD Momentum	0.2226 \pm 0.0169	0.8674 \pm 0.0048	3.8587 \pm 0.0058	65.4622 \pm 0.3842
SGD Momentum Normalized	0.1262 \pm 0.0170	0.8795 \pm 0.0086	3.8487 \pm 0.0073	61.0558 \pm 0.3224
SGDClipGrad	0.1288 \pm 0.0403	0.8677 \pm 0.0106	3.5774 \pm 0.0081	60.1604 \pm 0.2797
SGDClipMomentum	0.1220 \pm 0.0162	0.8809 \pm 0.0022	3.6038 \pm 0.0102	59.3052 \pm 0.2798
Adam	0.1161 \pm 0.0111	0.8823 \pm 0.0041	3.7692 \pm 0.0062	58.9005 \pm 0.3058
Our Algorithm 1	0.1086 \pm 0.0129	0.8835 \pm 0.0032	3.7928 \pm 0.0425	58.9661 \pm 1.5218

0.0001 and the momentum parameter (β_1) to be 0.9. Figure 4 and Table 1 report the training and testing performance for each algorithm, showing that ours is among the best.

LSTM for Language Modeling on Penn Treebank We adopt a 3-layer AWD-LSTM [35] to do language modeling on the Penn Treebank (PTB) dataset [33](word level). The mini-batch size is 40 and we trained each algorithm for 750 epochs. Apart from the hyperparameters we stated above, we further fine-tuned the weight decay value for all algorithms noticing its significant influence on the performance. We choose the set of hyperparameters that give the smallest final validation perplexity. We report the results in Figure 5 and Table 1. It can be seen that we can match the performance of Adam while beating the others.

5.2 Transformers Observe (L_0, L_1) -smoothness

For Figure 1 which verifies the original form (1) of the (L_0, L_1) condition using the norm, we followed the method in Section H.3 of [56]. Specifically, given \mathbf{x}_t and \mathbf{x}_{t+1} , denote $\mathbf{d} := \mathbf{x}_{t+1} - \mathbf{x}_t$. We estimate the smoothness at \mathbf{x}_t by

$$\hat{L}_t = \max_{\gamma \in \{\delta_1, \delta_2, \dots, \delta_N\}} \frac{\|\nabla F(\mathbf{x}_t + \gamma\mathbf{d}) - \nabla F(\mathbf{x}_t)\|_2}{\|\gamma\mathbf{d}\|_2},$$

where $\{\delta_1, \delta_2, \dots, \delta_N\}$ denotes the sample locations and we use $\{\frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}\}$.

For Figure 3 verifying the coordinate-wise version (2) of the (L_0, L_1) condition, note that the equation is symmetric in that if we just swap \mathbf{x} and \mathbf{y} it shall still holds. Thus, during plotting, we compare $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t+1}) - \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| / |x_{t+1,j} - x_{t,j}|$ vs. $\min \left(\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|, \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t+1}) \right| \right)$.

Figure 1(a) is on training a 2-layer Transformer Encoder to do language modeling on the Wikitext-2 dataset. The implementation, settings, and parameter choices follow this.¹ We only plot the first 5 training epochs. Figure 1(b) and 3 are on training a 6-layer Transformer [50] to do machine translation on the WMT’16 Multimodal Machine Translation Task German-English dataset. The implementation of the transformer is forked from here² and we also follow their default settings. The mini-batch size is 256 and we trained for 400 epochs using Adam and report the whole training trajectory.

5.3 Clipping does not Affect Adam’s Performance

We compare clipping and non-clipping for Adam optimizer on the Wikitext-103 (103 million tokens, 180MB) [36] language modeling task, with a 16-layer GPT-2 transformer model [42]. This GPT-2 model has an input length of 256 tokens, 410-dimension word embedding, 16 Attention layers with 10 Attention heads and 2100 hidden dimensions. Model size is 201.58 MB. The vocabulary size is 28996. We use the hyper-parameter settings prescribed in [53]: batch size 256, warm up learning rate from 0 to 2.5×10^{-4} in the first 64000 samples (i.e., 250 iterations) and then cosine-anneal learning rate to zero, on top of an Adam optimizer. It takes about 40 hours to train 200 epochs on 8 V100 GPUs. We use clipping threshold max_norm 0.25 for the entire model as prescribed in the literature [53]. We also count that with this clipping scheme, clipping occurs in every single batch. As we can see from Figure 2, neither training loss (2.79 vs 2.76) nor perplexity score (27.92 vs 27.97) differs much in the clipping and non-clipping case, which is consistent with our theory that Adam naturally achieves gradients clipping effect.

¹https://pytorch.org/tutorials/beginner/transformer_tutorial.html

²<https://github.com/jadore801120/attention-is-all-you-need-pytorch>

6 Conclusion and Limitations

Smoothness has been a widely adopted condition for proving convergence rates of algorithms in the non-convex optimization scenario. Yet, it has been found that this assumption does not capture losses when employing some deep learning models including RNNs and LSTMs. In light of this, a relaxed smoothness assumption was proposed that aligns well with the practice. We observed that the loss surface of training using Transformers also exhibits this relaxed smoothness. Under this assumption, SGD with clipped gradient has been proven to work well. However, we found that clipping is not necessary for achieving convergence in such a setting. Indeed, we showed that a generalized SignSGD algorithm does not require explicit clipping but can almost guarantee the same bound as SGD with clipping. In the analyses, we identified the key effect of using momentum in analyzing Adam-type algorithms, that it reduces both the noise and the unbounded gradient norms. Finally, we conducted a variety of deep learning tasks showing that our algorithm can match Adam’s performance while exceeding others.

Limitations The current work is in no way a perfect one and there are many directions worth exploring beyond it. First of all, though our algorithm could be seen as a close resemblance to the original Adam algorithm, they are still not equal. Considering the huge popularity of Adam and its established effectiveness in practice, it is worth studying whether Adam in its original form can converge in the relaxed smooth setting. Second, while our Theorem 1 are upper bounds and cannot be directly compared between the two cases of β_2 , it does suggest that $\beta_2 = 0$ minimizes the worst-case convergence rate. However, it still does not fully explain the phenomenon that a choice of β_2 close to 1 yields better performance in using our Algorithm 1 as well as Adam in practice. Third, despite there are lower bounds showing that, for example, GD with a constant step size can be arbitrarily worse than GD with clipping, it would be more meaningful to study whether the relaxed smooth condition is inherently more difficult, possibly by establishing a lower bound for all first-order optimization algorithms. Fourth, we did show that Transformers observe the relaxed smoothness condition, but we consider it more beneficial to research in-depth what properties or structures make a model satisfy such conditions. Finally, when conducting our experiments, we observed that the weight decay value plays a prominent role in each optimizer’s performance, and that the best weight decay value varies for different optimizers. Thus, one potential direction would be to explore different ways of incorporating the regularization in a way to preserve the scale-freeness [37, 38] of Algorithm 1, just as AdamW [31] does [58].

Acknowledgements

Michael Crawshaw is supported by the Institute for Digital Innovation fellowship from George Mason University. Michael Crawshaw and Mingrui Liu are both supported by a grant from George Mason University. Francesco Orabona is supported by the National Science Foundation under the grants no. 1908111 “AF: Small: Collaborative Research: New Representations for Learning Algorithms and Secure Computation”, no. 2022446 “Foundations of Data Science Institute”, and no. 2046096 “CAREER: Parameter-free Optimization Algorithms for Machine Learning”. The majority of work of Zhenxun Zhuang was done when he was a Ph.D. student at Boston University.

References

- [1] Ya I. Alber, Alfredo N. Iusem, and Mikhail V. Solodov. On the projected subgradient method for nonsmooth convex optimization in a Hilbert space. *Mathematical Programming*, 81(1):23–35, 1998.
- [2] Lukas Balles and Philipp Hennig. Dissecting Adam: The sign, magnitude and variance of stochastic gradients. In *International Conference on Machine Learning*, pages 404–413. PMLR, 2018.
- [3] Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. signSGD: Compressed optimisation for non-convex problems. In *International Conference on Machine Learning*, pages 560–569. PMLR, 2018.
- [4] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

- [5] Xiangyi Chen, Sijia Liu, Ruoyu Sun, and Mingyi Hong. On the convergence of a class of Adam-type algorithms for non-convex optimization. In *International Conference on Learning Representations*, 2019.
- [6] Zaiyi Chen, Zhuoning Yuan, Jinfeng Yi, Bowen Zhou, Enhong Chen, and Tianbao Yang. Universal stagewise learning for non-convex problems with convergence on averaged solutions. In *International Conference on Learning Representations*, 2019.
- [7] Ashok Cutkosky and Harsh Mehta. Momentum improves normalized SGD. In *International Conference on Machine Learning*, pages 2260–2268. PMLR, 2020.
- [8] Ashok Cutkosky and Harsh Mehta. High-probability bounds for non-convex stochastic optimization with heavy tails. *Advances in Neural Information Processing Systems*, 34, 2021.
- [9] Ashok Cutkosky and Francesco Orabona. Momentum-based variance reduction in non-convex SGD. In *Advances in Neural Information Processing Systems*, pages 15236–15245, 2019.
- [10] Alexandre Défossez, Léon Bottou, Francis Bach, and Nicolas Usunier. A simple convergence proof of Adam and Adagrad. *arXiv preprint arXiv:2003.02395*, 2020.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [12] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- [13] Yuri Ermoliev. Stochastic quasigradient methods. In *Numerical techniques for stochastic optimization*, number 10 in Springer Series in Computational Mathematics, pages 141–185. Springer, 1988.
- [14] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [15] Eduard Gorbunov, Marina Danilova, and Alexander Gasnikov. Stochastic optimization with heavy-tailed noise via accelerated gradient clipping. *arXiv preprint arXiv:2005.10785*, 2020.
- [16] Nicholas JA Harvey, Christopher Liaw, Yaniv Plan, and Sikander Randhawa. Tight analyses for non-smooth stochastic gradient descent. In *Conference on Learning Theory*, pages 1579–1613. PMLR, 2019.
- [17] Elad Hazan, Kfir Y Levy, and Shai Shalev-Shwartz. Beyond convexity: Stochastic quasi-convex optimization. In *Advances in Neural Information Processing Systems*, 2015.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Geoffrey Hinton, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, and Brian Kingsbury. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal processing magazine*, 29, 2012.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [21] Prateek Jain and Purushottam Kar. Non-convex optimization for machine learning. *Found. Trends Mach. Learn.*, 10:142–336, 2017.
- [22] Jikai Jin, Bohang Zhang, Haiyang Wang, and Liwei Wang. Non-convex distributionally robust optimization: Non-asymptotic analysis. In *Advances in Neural Information Processing Systems*, volume 34, pages 2771–2782. Curran Associates, Inc., 2021.

[23] Michael Kamp, Linara Adilova, Joachim Sicking, Fabian Hüger, Peter Schlicht, Tim Wirtz, and Stefan Wrobel. Efficient decentralized deep learning by dynamic model averaging. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 393–409. Springer, 2018.

[24] Ahmed Khaled and Peter Richtárik. Better theory for SGD in the nonconvex world. *arXiv preprint arXiv:2002.03329*, 2020.

[25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

[26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[27] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[28] Chen-Yu Lee, Saining Xie, Patrick Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-supervised nets. In *Artificial intelligence and statistics*, pages 562–570. PMLR, 2015.

[29] Mu Li, David G Andersen, Jun Woo Park, Alexander J Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J Shekita, and Bor-Yiing Su. Scaling distributed machine learning with the parameter server. In *OSDI*, volume 14, pages 583–598, 2014.

[30] Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 983–992. PMLR, 2019.

[31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

[32] Vien V Mai and Mikael Johansson. Stability and convergence of stochastic gradient clipping: Beyond lipschitz continuity and smoothness. In *International Conference on Machine Learning*, pages 7325–7335. PMLR, 2021.

[33] Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. Building a large annotated corpus of English: The Penn Treebank. *Comput. Linguist.*, 19(2):313–330, June 1993.

[34] H Brendan McMahan and Matthew Streeter. Adaptive bound optimization for online convex optimization. *arXiv preprint arXiv:1002.4908*, 2010.

[35] Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and optimizing LSTM language models. In *International Conference on Learning Representations*, 2018.

[36] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *ICLR*, 2017.

[37] Francesco Orabona and Dávid Pál. Scale-free algorithms for online linear optimization. In *International Conference on Algorithmic Learning Theory*, pages 287–301. Springer, 2015.

[38] Francesco Orabona and Dávid Pál. Scale-free online learning. *Theoretical Computer Science*, 716:50–69, 2018. Special Issue on ALT 2015.

[39] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. Understanding the exploding gradient problem. *corr abs/1211.5063* (2012). *arXiv preprint arXiv:1211.5063*, 2012.

[40] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pages 1310–1318. PMLR, 2013.

[41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.

[42] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

[43] Sashank J. Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Hugh Brendan McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2021.

[44] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of Adam and beyond. In *International Conference on Learning Representations*, 2018.

[45] Peter Richtárik and Martin Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, 144:1–38, 2014.

[46] Martin Riedmiller and Heinrich Braun. A direct adaptive method for faster backpropagation learning: The RPROP algorithm. In *IEEE international conference on neural networks*, pages 586–591. IEEE, 1993.

[47] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

[48] Naum Zuselevich Shor. *Minimization methods for non-differentiable functions*, volume 3. Springer Science & Business Media, 2012.

[49] Tijmen Tieleman, Geoffrey Hinton, et al. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012.

[50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[51] Yizhou Wang, Yue Kang, Can Qin, Huan Wang, Yilun Xu, Yulun Zhang, and Yun Raymond Fu. Rethinking Adam: A twofold exponential moving average approach. *arXiv preprint arXiv:2106.11514*, 2021.

[52] Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: Sharp convergence over nonconvex landscapes. In *International Conference on Machine Learning*, pages 6677–6686. PMLR, 2019.

[53] Thomas Wolf. *Transfer Learning in Natural Language Processing*, 2019. Available at https://github.com/huggingface/naacl_transfer_learning_tutorial.

[54] Manzil Zaheer, Sashank Reddi, Devendra Sachan, Satyen Kale, and Sanjiv Kumar. Adaptive methods for nonconvex optimization. *Advances in neural information processing systems*, 31, 2018.

[55] Bohang Zhang, Jikai Jin, Cong Fang, and Liwei Wang. Improved analysis of clipping algorithms for non-convex optimization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 15511–15521. Curran Associates, Inc., 2020.

[56] 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*, 2020.

[57] Jingzhao Zhang, Sai Praneeth Karimireddy, Andreas Veit, Seungyeon Kim, Sashank Reddi, Sanjiv Kumar, and Suvrit Sra. Why are adaptive methods good for attention models? *Advances in Neural Information Processing Systems*, 33:15383–15393, 2020.

- [58] Zhenxun Zhuang, Mingrui Liu, Ashok Cutkosky, and Francesco Orabona. Understanding AdamW through proximal methods and scale-freeness. *Transactions on Machine Learning Research*, 2022.
- [59] Liu Ziyin, Zhikang T Wang, and Masahito Ueda. LaProp: Separating momentum and adaptivity in Adam. *arXiv preprint arXiv:2002.04839*, 2020.
- [60] Difan Zou, Yuan Cao, Yuanzhi Li, and Quanquan Gu. Understanding the generalization of Adam in learning neural networks with proper regularization. *arXiv preprint arXiv:2108.11371*, 2021.
- [61] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 11127–11135, 2019.

Checklist

- 1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]** See Section 6
 - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]** This paper is mainly a theoretical one analyzing machine learning algorithms, we do not foresee any potential negative societal impacts of our work.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
- 2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** See Section 3.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See the Appendix.
- 3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** In the supplemental material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See Section 5.1 and A.4.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** See Section 5
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]**
 - (b) Did you mention the license of the assets? **[Yes]** Included in the code folder of the supplemental materials.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[No]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[No]** Data use follows the permission of corresponding licenses.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[No]**
- 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

A Appendix

A.1 Properties of the coordinate-wise (L_0, L_1) Assumption 2

Below, we prove the descent lemma for coordinate-wisely (L_0, L_1) -smooth functions satisfying Assumption 2.

Lemma 1. *Let F be (L_0, L_1) -smooth coordinate-wisely. For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ for which $\|\mathbf{x} - \mathbf{y}\|_2 \leq \frac{1}{\|L_1\|_\infty}$, we have*

$$F(\mathbf{y}) \leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \sum_{j=1}^d \frac{\left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \right) \|\mathbf{y} - \mathbf{x}\|_2}{2} |y_j - x_j|.$$

Proof of Lemma 1.

$$\begin{aligned} F(\mathbf{y}) &= F(\mathbf{x}) + \int_0^1 \langle \nabla F(\mathbf{x} + u(\mathbf{y} - \mathbf{x})), \mathbf{y} - \mathbf{x} \rangle du \\ &= F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \int_0^1 \langle \nabla F(\mathbf{x} + u(\mathbf{y} - \mathbf{x})) - \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle du \\ &\leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \left| \int_0^1 \langle \nabla F(\mathbf{x} + u(\mathbf{y} - \mathbf{x})) - \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle du \right| \\ &\leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \int_0^1 | \langle \nabla F(\mathbf{x} + u(\mathbf{y} - \mathbf{x})) - \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle | du \\ &\leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \int_0^1 \sum_{j=1}^d \left| \left[\frac{\partial F}{\partial x_j}(\mathbf{x} + u(\mathbf{y} - \mathbf{x})) - \frac{\partial F}{\partial x_j}(\mathbf{x}) \right] (y_j - x_j) \right| du \\ &\leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \int_0^1 \sum_{j=1}^d \left| \frac{\partial F}{\partial x_j}(\mathbf{x} + u(\mathbf{y} - \mathbf{x})) - \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| |y_j - x_j| du \\ &= F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \sum_{j=1}^d \int_0^1 \left| \frac{\partial F}{\partial x_j}(\mathbf{x} + u(\mathbf{y} - \mathbf{x})) - \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| |y_j - x_j| du \\ &\leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \sum_{j=1}^d \int_0^1 u \left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \right) \|\mathbf{y} - \mathbf{x}\|_2 |y_j - x_j| du \\ &= F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \sum_{j=1}^d \frac{\left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \right) \|\mathbf{y} - \mathbf{x}\|_2}{2} |y_j - x_j|, \end{aligned}$$

where the second inequality uses the fact that $\left| \int_a^b F(x) dx \right| \leq \int_a^b |F(x)| dx$ and the final one is due to Assumption 2. \square

The following Lemma shows that our coordinate-wise (L_0, L_1) smooth assumption 2 is equivalent to the original (L_0, L_1) smooth assumption (1) at least in 1-d case.

Lemma 2. *Let $F : \mathbb{R} \rightarrow \mathbb{R}$ be a twice continuously differentiable function. Then if (1) there exists some $K_0, K_1 \geq 0$ such that it holds for any $x, y \in \mathbb{R}$ with $|y - x| \leq \frac{1}{K_1}$ that $|F'(y) - F'(x)| \leq (K_0 + K_1|F'(x)|)|y - x|$, then (2) there exists some $L_0, L_1 \geq 0$ such that it holds for any $x \in \mathbb{R}^d$ that $|F''(x)| \leq L_0 + L_1|F'(x)|$, and vice versa.*

Proof of Lemma 2. (1) \Rightarrow (2) By definition, for any $x \in \mathbb{R}$, we know that

$$F''(x) = \lim_{h \rightarrow 0} \frac{F'(x + h) - F'(x)}{h} \leq \lim_{h \rightarrow 0} \frac{|F'(x + h) - F'(x)|}{|h|}$$

$$\leq \lim_{h \rightarrow 0} \frac{(K_0 + K_1 |F'(\mathbf{x})|) |h|}{|h|} = K_0 + K_1 |F'(\mathbf{x})|.$$

(2) \Rightarrow (1) This is a special case for 1-d and $c = 1$ of Corollary A.4 in [55]. \square

Lemma 3. *When $L_{0,j} = L_0$ and $L_{1,j} = L_1$ for all $j \in [d]$, Assumption 2 implies (3) (up to constants).*

Proof. Suppose all $L_{0,j}, L_{1,j}$ are the same across j , then we have

$$\begin{aligned} \|\nabla F(\mathbf{y}) - \nabla F(\mathbf{x})\|_2 &= \sqrt{\sum_{j=1}^d \left| \frac{\partial F}{\partial x_j}(\mathbf{y}) - \frac{\partial F}{\partial x_j}(\mathbf{x}) \right|^2} \\ &\leq \sqrt{\sum_{j=1}^d \left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| \right)^2 \times \|\mathbf{y} - \mathbf{x}\|_2^2} \\ &\leq \sqrt{\sum_{j=1}^d \left(\frac{2L_{0,j}^2}{d} + 2L_{1,j}^2 \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right|^2 \right) \times \|\mathbf{y} - \mathbf{x}\|_2^2} \\ &\leq \sqrt{2L_0} \|\mathbf{y} - \mathbf{x}\|_2 + \sqrt{2L_1} \|\mathbf{y} - \mathbf{x}\|_2 \sqrt{\sum_{j=1}^d \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right|^2} \\ &= \left(\sqrt{2L_0} + \sqrt{2L_1} \|\nabla F(\mathbf{x})\|_2 \right) \times \|\mathbf{y} - \mathbf{x}\|_2. \end{aligned} \quad \square$$

A.2 Proof of Lower Bound

Theorem 2. *Fix $\epsilon > 0, L_0 > 0, L_1 > 0, M \geq \max(\frac{L_0}{L_1}, \epsilon)$, and $x_0 \in \mathbb{R}$. Pick any constant learning rate η for GD, with the knowledge of the above constants. Then, there exists a 1-d (L_0, L_1) -smooth function, bounded from below by f^* (finite), and such that $\sup\{|f'(x)| : f(x) \leq f(x_0)\} \leq M$ on which the number of iterations T of GD with learning rate η to guarantee $|f'(x_T)| < \epsilon$ is at least*

$$\frac{ML_1(f(x_0) - f^* - \frac{15\epsilon^2}{16L_0})}{2\epsilon^2 \left(\ln \frac{ML_1}{L_0} + 1 \right)}.$$

Proof of Theorem 2. By Lemma 2, we know that, in 1-d case, our coordinate-wise $(\mathbf{L}_0, \mathbf{L}_1)$ assumption 2 is equivalent as the original one (1). Thus, without loss of generality, we use the original condition (1) in the proof. We will construct two different (L_0, L_1) -smooth functions based on the value of η .

Case $\eta > \frac{2}{ML_1} \left(\ln \frac{ML_1}{L_0} + 1 \right)$. In this case, we can construct a function on which GD does not converge, hence the lower bound is trivially true. Consider the function

$$f(x) = \begin{cases} L_0 \frac{e^{-L_1 x - 1}}{L_1^2} & x < -\frac{1}{L_1} \\ L_0 \frac{x^2}{2} + \frac{L_0}{2L_1^2} & x \in [-\frac{1}{L_1}, \frac{1}{L_1}] \\ L_0 \frac{e^{L_1 x - 1}}{L_1^2} & x > \frac{1}{L_1} \end{cases}$$

Note that f is (L_0, L_1) -smooth. Without loss of generality, we can assume $x_0 = \frac{1}{L_1} \left(\ln \frac{ML_1}{L_0} + 1 \right)$, in fact if this is not the case we can translate the function f accordingly. This setting of x_0 guarantees that the bound on the gradient is correct. Moreover, with this choice, we claim the function will diverge. To see this, we use mathematical induction to show that $|x_{t+1}| > |x_t|$ and $\text{sign}(x_{t+1}) \neq \text{sign}(x_t)$ for any $t \geq 0$. First, for the case when $t = 0$, we have

$$x_1 = x_0 - \eta f'(x_0) = x_0 - \frac{\eta L_0}{L_1} e^{L_1 x_0 - 1} = x_0 - \eta M < x_0 - 2x_0 = -x_0.$$

Then suppose the condition holds up until t and we prove for $t + 1$. From the formula of f , we have that $\text{sign}(f'(x)) = \text{sign}(x)$ and that f is monotonically increasing with $|x|$. Thus, from the update of gradient descent which moves along the negative direction of the gradient, if we can show that $|x_{t+1}| > |x_t|$, then $\text{sign}(x_{t+1}) \neq \text{sign}(x_t)$. This yields

$$|x_{t+1}| = |x_t - \eta f'(x_t)| > |x_t| \Leftarrow \eta |f'(x_t)| > 2|x_t| \Leftarrow \eta L_0 > \frac{2|x_t|L_1}{\exp(L_1|x_t| - 1)}.$$

Now, note that $\psi(x) = \frac{2|x|L_1}{\exp(L_1|x| - 1)}$ is decreasing for $x > \frac{1}{L_1}$ and increasing for $x < -\frac{1}{L_1}$. Hence, we have that

$$\eta L_0 > \frac{2|x_0|L_1}{\exp(L_1|x_0| - 1)} > \frac{2|x_t|L_1}{\exp(L_1|x_t| - 1)},$$

where the first inequality is true by the choice of $x_0 > \frac{1}{L_1}$ and the condition on η and the second one is true by the induction hypothesis.

Case $\eta \leq \frac{2}{ML_1} \left(\ln \frac{ML_1}{L_0} + 1 \right)$.

Now, consider

$$f(x) = \begin{cases} -\epsilon x, & x < -\frac{3\epsilon}{2L_0} \\ \frac{L_0}{2}x^2 - \frac{L_0^3x^4}{27\epsilon^2} + \frac{9\epsilon^2}{16L_0}, & x \in [-\frac{3\epsilon}{2L_0}, \frac{3\epsilon}{2L_0}] \\ \epsilon x, & x > \frac{3\epsilon}{2L_0} \end{cases}$$

We have that f is $(L_0, 0)$ -smooth, hence also (L_0, L_1) -smooth. Note that the presence of the fourth power makes this function twice differentiable. Moreover, the maximum gradient in this case is $\epsilon \leq M$.

As before, without loss of generality, let the initial point $x_0 = \frac{3\epsilon}{2L_0} + \Delta$, where $\Delta > 0$. We have that $f(x_0) - f^* = \epsilon \left(\Delta + \frac{3\epsilon}{2L_0} \right) - \frac{9\epsilon^2}{16L_0}$, hence $\Delta = \frac{1}{\epsilon} (f(x_0) - f^*) - \frac{15\epsilon}{16L_0}$. Now, while we stay on the last branch of the function, we have

$$x_{t+1} = x_t - \eta \epsilon \geq x_t - \epsilon \frac{2}{ML_1} \left(\ln \frac{ML_1}{L_0} + 1 \right).$$

Hence, we have that, for

$$t \leq \frac{ML_1 \Delta}{2\epsilon \left(\ln \frac{ML_1}{L_0} + 1 \right)} = \frac{ML_1 \left(f(x_0) - f^* - \frac{15\epsilon^2}{16L_0} \right)}{2\epsilon^2 \left(\ln \frac{ML_1}{L_0} + 1 \right)},$$

we guarantee $|f'(x_t)| = \epsilon$. \square

Errors in the lower bound in [56] As we said in the main text, unfortunately, the lower bound theorem in [56] is wrong, both statement and proof. First of all, they have a logarithm of a quantity with units, M , which is an undefined mathematical operation. A closer look at the proof reveals that, differently from the statement of their theorem, they construct a function with $L_0 = L_1$, which explains why these terms are missing in the logarithm. Moreover, it is also unclear if the second constructed function satisfies the assumptions of the theorem. We correct all these issues by properly scaling the constructed functions so that they always satisfy the (L_0, L_1) condition and all the units are coherent. This result in the correct term inside the logarithm and the right conditions on L_0, L_1, M , and ϵ .

A.3 Proof of Theorem 1

We first write down some notations here that we will use heavily later for easier reference:

$$\begin{aligned} \bar{\tau} &= \frac{\sqrt{1 - \beta_2}}{\eta \sqrt{d} \|\mathbf{L}_1\|_\infty}, & \alpha &= 1 - \beta_1, & \rho &= 1 - \beta_2^{1/2} \beta_1^{-1}, \\ \epsilon_t &= \mathbf{m}_t - \nabla F(\mathbf{x}_t), & \tilde{\epsilon}_t &= \mathbf{g}_t - \nabla F(\mathbf{x}_t), \end{aligned}$$

$$\begin{aligned}
E_j &= 6\sigma_j \max(1, \log(1/\delta)) + \frac{6}{\sqrt{1-\beta_1^2}} \sqrt{\sigma_j^2 \max(1, \log(1/\delta))}, \\
B_j &= \frac{\eta L_{0,j}}{\sqrt{1-\beta_2}(1-\beta_1)} + \beta_1^{\bar{\tau}}(M_j + \sigma_j) + (1-\beta_1)E_j, \\
C_j &= 1 + \frac{\eta\sqrt{d}L_{1,j}}{(1-\beta_1)\sqrt{1-\beta_2}}, \quad D = 1 - \frac{2\eta\sqrt{d}\|\mathbf{L}_1\|_\infty}{\sqrt{1-\beta_2}(1-\beta_1)}, \\
A &= \frac{\rho}{10\sqrt{1-\beta_2}}.
\end{aligned}$$

Also, we would need the following formula many times:

$$\beta_1^{\bar{\tau}} = (1-\alpha)^{\frac{1}{\alpha}} \frac{\alpha\sqrt{1-\beta_2}}{\eta\sqrt{d}\|\mathbf{L}_1\|_\infty} \leq e^{-\frac{\alpha\sqrt{1-\beta_2}}{\eta\sqrt{d}\|\mathbf{L}_1\|_\infty}}, \quad (4)$$

where in the first inequality we used the fact that $(1-x)^{1/x} \leq \frac{1}{e}$ for $0 < x < 1$.

Lemma 4. *With the notations in Algorithm 1, for each coordinate $j \in [d]$ we have*

$$m_{t,j} = (1-\beta_1) \sum_{\tau=1}^t \beta_1^{t-\tau} g_{\tau,j}, \quad v_{t,j} = (1-\beta_2) \sum_{\tau=1}^t \beta_2^{t-\tau} m_{\tau,j}^2, \quad \frac{|m_{t,j}|}{\sqrt{v_{t,j}}} \leq \frac{1}{\sqrt{1-\beta_2}}.$$

Proof of Lemma 4. For all $t \geq 1$, we have

$$\begin{aligned}
m_{t,j} &= \beta_1 m_{t-1,j} + (1-\beta_1) g_{t,j} \\
&= \beta_1 [\beta_1 m_{t-2,j} + (1-\beta_1) g_{t-1,j}] + (1-\beta_1) g_{t,j} \\
&= \dots = (1-\beta_1) \sum_{\tau=1}^t \beta_1^{t-\tau} g_{\tau,j}.
\end{aligned}$$

Similarly for $v_{t,j}$. Next,

$$\frac{|m_{t,j}|}{\sqrt{v_{t,j}}} = \frac{|m_{t,j}|}{\sqrt{(1-\beta_2) \sum_{\tau=1}^t \beta_2^{t-\tau} m_{\tau,j}^2}} \leq \frac{1}{\sqrt{1-\beta_2}}. \quad \square$$

The following lemma shows when we can apply Assumption 2 and Lemma 1.

Lemma 5. *With notations in Algorithm 1, for $\tau \leq \bar{\tau} = \frac{\sqrt{1-\beta_2}}{\eta\sqrt{d}\|\mathbf{L}_1\|_\infty}$, we have $\|\mathbf{x}_{t-\tau} - \mathbf{x}_t\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty}$.*

Proof of Lemma 5. Using Lemma 4 we have

$$|x_{t-\tau,j} - x_{t,j}| \leq \sum_{i=1}^{\tau} |x_{t-i,j} - x_{t-i+1,j}| \leq \frac{\eta\tau}{\sqrt{1-\beta_2}} \leq \frac{1}{\sqrt{d}\|\mathbf{L}_1\|_\infty} \Rightarrow \|\mathbf{x} - \mathbf{y}\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty}. \quad \square$$

The following two lemmas are the major tools we use to analyze the effects of noises.

Lemma 6 (Lemma 12, [8]). *Suppose X_1, \dots, X_T is a martingale difference sequence in a Hilbert space such that $\|X_t\| \leq R$ almost surely for some constant R . Further, assume $\mathbb{E}_t[\|X_t\|^2] \leq \sigma_t^2$ with probability 1 for some constants σ_t , where $\mathbb{E}_t[\cdot] \triangleq \mathbb{E}[\cdot | \xi_1, \xi_2, \dots, \xi_{t-1}]$ denotes the expectation conditioned on all past randomnesses. Then, with probability at least $1 - 3\delta$, for all $k \leq T$ we have*

$$\left\| \sum_{t=1}^k X_t \right\| \leq 3R \max(1, \log(1/\delta)) + 3 \sqrt{\sum_{t=1}^k \sigma_t^2 \max(1, \log(1/\delta))}.$$

Lemma 7. *Assume Assumption 3. With the notation of Algorithm 1, let $j \in [d]$ and $\beta_1 < 1$. Then, with probability at least $1 - 3\delta$, for any $t_0 \in [t]$, we have*

$$\left| \sum_{\tau=1}^{t_0} \beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right| \leq 3\sigma_j \max(1, \log(1/\delta)) + \frac{3}{\sqrt{1-\beta_1^2}} \sqrt{\sigma_j^2 \max(1, \log(1/\delta))}.$$

Proof of Lemma 7. Recall Assumption 3 and notice that $\beta_1^{t-\tau} \leq 1$ for all $\tau \in [1, t]$, we know that $\left| \beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right| \leq \sigma_j$ almost surely. It also means $\mathbb{E}_\tau \left[\left(\beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right)^2 \right] \leq \beta_1^{2(t-\tau)} \sigma_j^2$. Now, note that in Algorithm 1 \mathbf{g}_τ is an unbiased estimate of $\nabla F(\mathbf{x}_\tau)$ namely $\mathbb{E}_\tau \left[\beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right] = 0$. Thus, $\left\{ \beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right\}_{1, \dots, t}$ is a martingale difference sequence. Then, using Lemma 6, with probability at least $1 - 3\delta$, we have for all $t_0 \in [t]$ that

$$\begin{aligned} \left| \sum_{\tau=1}^{t_0} \beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right| &\leq 3\sigma_j \max(1, \log(1/\delta)) + 3\sqrt{\sum_{\tau=1}^{t_0} \beta_1^{2(t-\tau)} \sigma_j^2 \max(1, \log(1/\delta))} \\ &\leq 3\sigma_j \max(1, \log(1/\delta)) + \frac{3}{\sqrt{1-\beta_1^2}} \sqrt{\sigma_j^2 \max(1, \log(1/\delta))}. \quad \square \end{aligned}$$

The following lemma upper bounds the differences between recent true gradients and the current one.

Lemma 8. *With the notation of Algorithm 1 and under the assumptions in Theorem 1, for any $j \in [d]$ and any t_0 with $t - t_0 \leq \bar{\tau} = \frac{\sqrt{1-\beta_2}}{\eta\sqrt{d}\|\mathbf{L}_1\|_\infty}$, we have*

$$\sum_{\tau=t_0}^t \beta_1^{t-\tau} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right| \leq \left(L_{0,j} + L_{1,j}\sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \frac{\eta}{(1-\beta_1)^2\sqrt{1-\beta_2}}.$$

Proof of Lemma 8.

$$\begin{aligned} &\sum_{\tau=t_0}^t \beta_1^{t-\tau} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right| \\ &\leq \sum_{\tau=t_0}^t \beta_1^{t-\tau} \left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \|\mathbf{x}_t - \mathbf{x}_\tau\|_2 \\ &\leq \sum_{\tau=t_0}^t \beta_1^{t-\tau} \left(L_{0,j} + L_{1,j}\sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \frac{\eta(t-\tau)}{\sqrt{1-\beta_2}} \\ &= \left(L_{0,j} + L_{1,j}\sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \frac{\eta}{\sqrt{1-\beta_2}} \sum_{\tau=t_0}^t (t-\tau) \beta_1^{t-\tau} \\ &\leq \left(L_{0,j} + L_{1,j}\sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \frac{\eta}{(1-\beta_1)^2\sqrt{1-\beta_2}}, \end{aligned}$$

where the first inequality is due to Assumption 2 and Lemma 5, the second inequality uses Lemma 4, and the final inequality uses the fact that $\sum_{k=1}^N ka^k \leq \frac{1}{(1-a)^2}$ for any $0 < a < 1$. \square

The following lemma upper bounds a past momentum with the current one.

Lemma 9. *With the notation of Algorithm 1 and under the assumptions of Theorem 1, for any $\tau \leq \bar{\tau} = \frac{\sqrt{1-\beta_2}}{\eta\sqrt{d}\|\mathbf{L}_1\|_\infty}$, with probability at least $1 - 3\delta$, it holds that*

$$|m_{t-\tau,j}| \leq \beta_1^{-\tau} \left(|m_{t,j}| + \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + \left(L_{0,j} + L_{1,j}\sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \frac{\eta}{(1-\beta_1)\sqrt{1-\beta_2}} + (1-\beta_1)E_j \right).$$

Proof of Lemma 9. Denoting by $\tilde{\mathbf{e}}_t = \mathbf{g}_t - \nabla F(\mathbf{x}_t)$ and using Lemma 4, we have

$$|m_{t-\tau,j} - \beta_1^{-\tau} m_{t,j}| = \left| (1-\beta_1) \sum_{i=1}^t \beta_1^{t-\tau-i} g_{i,j} - (1-\beta_1) \sum_{i=1}^{t-\tau} \beta_1^{t-\tau-i} g_{i,j} \right|$$

$$\begin{aligned}
&= (1 - \beta_1) \left| \sum_{i=t-\tau+1}^t \beta_1^{t-\tau-i} g_{i,j} \right| \\
&\leq (1 - \beta_1) \left| \sum_{i=t-\tau+1}^t \beta_1^{t-\tau-i} \frac{\partial F}{\partial x_j}(\mathbf{x}_i) \right| + (1 - \beta_1) \left| \sum_{i=t-\tau+1}^t \beta_1^{t-\tau-i} \tilde{\epsilon}_{i,j} \right| \quad (5)
\end{aligned}$$

We now upper bound the first term of (5) using Lemma 8 by using the fact that $\tau \leq \frac{\sqrt{1-\beta_2}}{\eta \|\mathbf{L}_1\|_1}$:

$$\begin{aligned}
\left| \sum_{i=t-\tau+1}^t \beta_1^{t-\tau-i} \frac{\partial F}{\partial x_j}(\mathbf{x}_i) \right| &\leq \left| \sum_{i=t-\tau+1}^t \beta_1^{t-\tau-i} \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + \sum_{i=t-\tau+1}^t \beta_1^{t-\tau-i} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) - \frac{\partial F}{\partial x_j}(\mathbf{x}_i) \right| \\
&\leq \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \frac{\beta_1^{-\tau}}{1 - \beta_1} + \left(L_{0,j} + L_{1,j} \sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \frac{\eta \beta_1^{-\tau}}{(1 - \beta_1)^2 \sqrt{1 - \beta_2}}.
\end{aligned}$$

Finally, the second term of (5) can be bounded using Lemma 7 by noticing that

$$\left| \sum_{i=t-\tau+1}^t \beta_1^{t-\tau-i} \tilde{\epsilon}_{i,j} \right| \leq \beta_1^{-\tau} \left(\left| \sum_{i=1}^t \beta_1^{t-i} \tilde{\epsilon}_{i,j} \right| + \left| \sum_{i=1}^{t-\tau} \beta_1^{t-i} \tilde{\epsilon}_{i,j} \right| \right). \quad \square$$

The following Lemma is adapted from [60]. Yet, they only considered Adam under the L -smooth setting and when there is no noise. The existence of noise and the relaxed smoothness assumption makes the proofs substantially more challenging. With this lemma, we know that either the true gradient is small or that the update of our Algorithm 1 can be lower bounded.

Lemma 10. *With the notation of Algorithm 1 and under the assumptions of Theorem 1, if $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right| \leq M_j$ holds for all $\tau \leq t$ and $j \in [d]$, and $D > 0$, then, with probability at least $1 - 3t\delta$ we have that,*

$$\text{either } \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| < \frac{5B_j}{D} \text{ or } \frac{|m_{t,j}|}{\sqrt{v_{t,j}}} \geq \frac{\rho D}{5\sqrt{1 - \beta_2}}.$$

Proof. Given that $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right| \leq M_j$ for any $\tau \leq t$ and $j \in [d]$, using Lemma 4 and Assumption 3, it is immediate to show that $|m_{t,j}| \leq M_j + \sigma_j$. Then, denote $\hat{\tau} = \lfloor \bar{\tau} \rfloor = \left\lfloor \frac{\sqrt{1-\beta_2}}{\eta \sqrt{d} \|\mathbf{L}_1\|_\infty} \right\rfloor$ namely the largest integer that is no greater than $\bar{\tau}$, from Lemma 4, we have

$$\begin{aligned}
\frac{|m_{t,j}|}{\sqrt{v_{t,j}}} &= \frac{|m_{t,j}|}{\sqrt{(1 - \beta_2) \sum_{\tau=0}^{t-1} \beta_2^\tau m_{t-\tau,j}^2}} \\
&= \frac{|m_{t,j}|}{\sqrt{1 - \beta_2} \sqrt{\sum_{\tau=\hat{\tau}+1}^{t-1} \beta_2^\tau m_{t-\tau,j}^2 + \sum_{\tau=0}^{\hat{\tau}} \beta_2^\tau m_{t-\tau,j}^2}} \\
&\geq \frac{|m_{t,j}|}{\sqrt{1 - \beta_2} \sqrt{(M_j + \sigma_j)^2 \frac{\beta_2^{\hat{\tau}+1}}{1 - \beta_2} + \sum_{\tau=0}^{\hat{\tau}} \beta_2^\tau m_{t-\tau,j}^2}} \\
&\geq \frac{|m_{t,j}|}{(M_j + \sigma_j) \beta_2^{\bar{\tau}/2} + \sqrt{1 - \beta_2} \sum_{\tau=0}^{\hat{\tau}} \beta_2^{\tau/2} |m_{t-\tau,j}|} \\
&> \frac{|m_{t,j}|}{(M_j + \sigma_j) \beta_1^{\bar{\tau}} + \sqrt{1 - \beta_2} \sum_{\tau=0}^{\hat{\tau}} \beta_2^{\tau/2} |m_{t-\tau,j}|},
\end{aligned}$$

where the final inequality uses the assumption that $\sqrt{\beta_2} < \beta_1$. Using Lemma 9 and the definition of $\rho = 1 - \beta_2^{1/2} \beta_1^{-1} \in (0, 1]$, with probability at least $1 - 3t\delta$, as we need to invoke Lemma 7 for at most t times, we have

$$\frac{\sqrt{v_{t,j}}}{\sqrt{1 - \beta_2}} \leq (|m_{t,j}| + (1 - \beta_1) E_j) \sum_{\tau=0}^{\hat{\tau}} \beta_2^{\tau/2} \beta_1^{-\tau} + \frac{\beta_1^{\bar{\tau}} (M_j + \sigma_j)}{\sqrt{1 - \beta_2}}$$

$$\begin{aligned}
& + \left(\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + \left(L_{0,j} + L_{1,j} \sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \frac{\eta}{(1 - \beta_1) \sqrt{1 - \beta_2}} \right) \sum_{\tau=0}^{\hat{\tau}} \beta_2^{\tau/2} \beta_1^{-\tau} \\
& \leq \left(|m_{t,j}| + C_j \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + B_j \right) \frac{1}{\rho},
\end{aligned}$$

where in the last inequality we used the fact that $\frac{1}{\rho} \geq \frac{1}{\sqrt{1 - \beta_2}}$.

Thus, we consider following two cases depending on the relative size of $|m_{t,j}|$ vs. $C_j \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + B_j$.

Case 1: $|m_{t,j}| > C_j \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + B_j$, then

$$\frac{|m_{t,j}|}{\sqrt{v_{t,j}}} > \frac{\rho}{2\sqrt{1 - \beta_2}}. \quad (6)$$

Case 2: $|m_{t,j}| \leq C_j \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + B_j$, then we have

$$\sqrt{v_{t,j}} \leq \frac{2\sqrt{1 - \beta_2}}{\rho} \left(C_j \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + B_j \right).$$

Also, for $|m_{t,j}|$ we have from Lemma 4 that

$$\begin{aligned}
|m_{t,j}| &= (1 - \beta_1) \left| \sum_{\tau=1}^t \beta_1^{t-\tau} g_{\tau,j} \right| \geq (1 - \beta_1) \left| \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} g_{\tau,j} \right| - (1 - \beta_1) \left| \sum_{\tau=1}^{t-\hat{\tau}-1} \beta_1^{t-\tau} g_{\tau,j} \right| \\
&\geq \underbrace{(1 - \beta_1) \left| \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right|}_{R_1} - \underbrace{(1 - \beta_1) \left| \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} \left(\frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) - g_{\tau,j} \right) \right|}_{R_2} - \underbrace{(1 - \beta_1) \left| \sum_{\tau=1}^{t-\hat{\tau}-1} \beta_1^{t-\tau} g_{\tau,j} \right|}_{R_3}.
\end{aligned}$$

The first term can be bounded below by using Lemma 8 and that $\hat{\tau} + 1 \geq \bar{\tau}$:

$$\begin{aligned}
R_1 &\geq (1 - \beta_1) \left| \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right| - (1 - \beta_1) \left| \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} \left(\frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) - \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right) \right| \\
&\geq \left(1 - \beta_1^{\bar{\tau}} - \frac{\sqrt{d} L_{1,j} \eta}{(1 - \beta_1) \sqrt{1 - \beta_2}} \right) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| - \frac{L_{0,j} \eta}{(1 - \beta_1) \sqrt{1 - \beta_2}}.
\end{aligned}$$

The second term can be bounded using Lemma 7. Thus,

$$\begin{aligned}
|m_{t,j}| &\geq \left(1 - \beta_1^{\bar{\tau}} - \frac{\sqrt{d} L_{1,j} \eta}{(1 - \beta_1) \sqrt{1 - \beta_2}} \right) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| - \frac{L_{0,j} \eta}{(1 - \beta_1) \sqrt{1 - \beta_2}} - \beta_1^{\bar{\tau}} (M_j + \sigma_j) - (1 - \beta_1) E_j \\
&\geq D \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| - B_j,
\end{aligned}$$

where we used (4) and that $e^{-x} \leq \frac{1}{x}$ for $x > 0$.

Therefore, with probability at least $1 - 3t\delta$ we have

$$\frac{|m_{t,j}|}{\sqrt{v_{t,j}}} \geq \frac{\rho \left(D \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| - B_j \right)}{2\sqrt{1 - \beta_2} \left(C_j \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + B_j \right)}. \quad (7)$$

Given that $D > 0$, depending on the relative size of $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|$ vs. B_j , we have following two cases.

Case 2.1: $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| < \frac{5B_j}{D}$.

Case 2.2: $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \geq \frac{5B_j}{D}$, using the fact that the r.h.s of (7) is decreasing in B_j , we have

$$\frac{|m_{t,j}|}{\sqrt{v_{t,j}}} \geq \frac{\frac{4D}{5}\rho \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|}{2\sqrt{1-\beta_2} \left(C_j + \frac{D}{5} \right) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|} \geq \frac{2\rho D}{5\sqrt{1-\beta_2}(C_j + D)} \geq \frac{\rho D}{5\sqrt{1-\beta_2}},$$

where in the last inequality we used the fact that $C_j + D \leq 2$. Note that $D < 1$ so the above lower bound is smaller than (6). \square

The following two lemmas are for the special case of $\beta_2 = 0$.

Lemma 11. *With choices of parameters in Theorem 1, when $\beta_2 = 0$, we have $\|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2 = \eta\sqrt{d} \leq \frac{1}{\|\mathbf{L}_1\|_\infty}$.*

Proof of Lemma 11. Using the fact that $\alpha \leq 1$ and the condition on T , we have

$$\eta = \frac{\sqrt{\Delta\alpha}}{\sqrt{\|\mathbf{L}_0\|_1}\sqrt{T}} \leq \frac{\sqrt{\Delta}}{\sqrt{\|\mathbf{L}_0\|_1}\sqrt{T}} \leq \frac{\sqrt{\Delta}}{\sqrt{\|\mathbf{L}_0\|_1}} \frac{\sqrt{\|\mathbf{L}_0\|_1}}{10\sqrt{d}\sqrt{\Delta}\|\mathbf{L}_1\|_\infty} \leq \frac{1}{\sqrt{d}\|\mathbf{L}_1\|_\infty}. \quad \square$$

The following lemma is adapted from the proof of Theorem 2 in [8].

Lemma 12. *Under Assumptions 1, 2, and 3, using the settings of the hyperparameters in Theorem 1, denoting $\alpha = 1 - \beta_1$ and $\epsilon_t = \mathbf{m}_t - \nabla F(\mathbf{x}_t)$, for all $t \geq 1$ and $j \in [d]$ we have, with probability at least $1 - 3\delta$,*

$$|\epsilon_{t+1,j}| \leq (1-\alpha)^t \left(\alpha\sigma_j + (1-\alpha) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| \right) + \frac{\eta L_{0,j}}{\alpha} + \alpha E_j + (1-\alpha)\eta\sqrt{d}L_{1,j} \sum_{\tau=0}^{t-1} (1-\alpha)^\tau \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t-\tau}) \right|.$$

Proof of Lemma 12. Denote $\tilde{\epsilon}_t = \mathbf{g}_t - \nabla F(\mathbf{x}_t)$ and $S_j(\mathbf{a}, \mathbf{b}) = \frac{\partial F}{\partial x_j}(\mathbf{a}) - \frac{\partial F}{\partial x_j}(\mathbf{b})$. Then, from Assumption 2 and Lemma 11, for all $t \geq 1$ and all $j \in [d]$ we have

$$\epsilon_1 = \alpha\tilde{\epsilon}_1 - (1-\alpha)\nabla F(\mathbf{x}_1), \quad (8)$$

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty} \Rightarrow |S_j(\mathbf{x}_{t+1}, \mathbf{x}_t)| \leq \left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2. \quad (9)$$

We can derive the following recursive formulation for any $t \geq 1$:

$$\begin{aligned} m_{t+1,j} &= (1-\alpha)m_{t,j} + \alpha g_{t+1,j} \\ &= (1-\alpha) \frac{\partial F}{\partial x_j}(\mathbf{x}_t) + (1-\alpha)\epsilon_{t,j} + \alpha \frac{\partial F}{\partial x_j}(\mathbf{x}_{t+1}) + \alpha\tilde{\epsilon}_{t+1,j} \\ &= \frac{\partial F}{\partial x_j}(\mathbf{x}_{t+1}) + (1-\alpha)S_j(\mathbf{x}_t, \mathbf{x}_{t+1}) + (1-\alpha)\epsilon_{t,j} + \alpha\tilde{\epsilon}_{t+1,j}, \end{aligned}$$

which implies

$$\epsilon_{t+1,j} = (1-\alpha)\epsilon_{t,j} + (1-\alpha)S_j(\mathbf{x}_t, \mathbf{x}_{t+1}) + \alpha\tilde{\epsilon}_{t+1,j}. \quad (10)$$

Unravel (10) from 1 to t gives us

$$\epsilon_{t+1,j} = (1-\alpha)^t \epsilon_{1,j} + (1-\alpha) \sum_{\tau=0}^{t-1} (1-\alpha)^\tau S_j(\mathbf{x}_{t-\tau}, \mathbf{x}_{t+1-\tau}) + \alpha \sum_{\tau=0}^{t-1} (1-\alpha)^\tau \tilde{\epsilon}_{t+1-\tau,j}.$$

Take the absolute value of both sides, to obtain

$$|\epsilon_{t+1,j}| \leq (1-\alpha)^t |\epsilon_{1,j}| + (1-\alpha) \sum_{\tau=0}^{t-1} (1-\alpha)^\tau |S_j(\mathbf{x}_{t-\tau}, \mathbf{x}_{t+1-\tau})| + \alpha \left| \sum_{\tau=0}^{t-1} (1-\alpha)^\tau \tilde{\epsilon}_{t+1-\tau,j} \right|$$

$$\begin{aligned}
&\leq (1-\alpha)^t |\epsilon_{1,j}| + (1-\alpha) \sum_{\tau=0}^{t-1} (1-\alpha)^\tau \left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t-\tau}) \right| \right) \|\mathbf{x}_{t+1-\tau} - \mathbf{x}_{t-\tau}\|_2 + \alpha E_j \\
&\leq (1-\alpha)^t |\epsilon_{1,j}| + (1-\alpha) \eta L_{0,j} \sum_{\tau=0}^{t-1} (1-\alpha)^\tau + (1-\alpha) \eta \sqrt{d} L_{1,j} \sum_{\tau=0}^{t-1} (1-\alpha)^\tau \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t-\tau}) \right| + \alpha E_j \\
&\leq (1-\alpha)^t \left(\alpha \sigma_j + (1-\alpha) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| \right) + \frac{\eta L_{0,j}}{\alpha} + (1-\alpha) \eta \sqrt{d} L_{1,j} \sum_{\tau=0}^{t-1} (1-\alpha)^\tau \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t-\tau}) \right| + \alpha E_j ,
\end{aligned}$$

where the second inequality uses (9) and Lemma 7, the fourth and fifth inequalities use (9), and the final one is due to (8). \square

Theorem 1. *Under Assumptions 1, 2, and 3, assume $M_j := \sup \left\{ \left| \frac{\partial F}{\partial x_j}(\mathbf{x}) \right| : F(\mathbf{x}) \leq F(\mathbf{x}_1) \right\}$ is finite for each $j \in [d]$, let Δ be any upper bound on $F(\mathbf{x}_1) - F^*$, $\alpha = \min \left(\frac{\sqrt{\|\mathbf{L}_0\|_1} \sqrt{\Delta}}{\|\boldsymbol{\sigma}\|_1 \sqrt{T}}, 1 \right)$, $\beta_1 = 1 - \alpha$, $\frac{\sqrt{\beta_2}}{\beta_1} < 1$, $\rho = 1 - \frac{\sqrt{\beta_2}}{\beta_1}$, $\eta = \frac{\sqrt{\Delta \alpha}}{\sqrt{\|\mathbf{L}_0\|_1} \sqrt{T}}$, for $T \geq \max \left(\frac{100d\Delta \|\mathbf{L}_1\|_\infty^2}{(1-\beta_2)\rho^2 \|\mathbf{L}_0\|_1}, \frac{10000d^2 \Delta \|\boldsymbol{\sigma}\|_1^2 \|\mathbf{L}_1\|_\infty^4}{(1-\beta_2)^2 \rho^4 \|\mathbf{L}_0\|_1^3} \right)$, Algorithm 1 guarantees, with probability at least $1 - \delta$, that*

$$\begin{aligned}
\min_{t \in [T]} \|\nabla F(\mathbf{x}_t)\|_1 &= \mathcal{O} \left(\frac{\sqrt{\log(dT/\delta)} \|\mathbf{L}_0\|_1^{1/4} \Delta^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2}}{\rho \sqrt{1 - \beta_2} T^{1/4}} + \frac{\log(dT/\delta) \sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\rho \sqrt{T}} \right) \\
&\quad + \mathcal{O} \left(\frac{\|\mathbf{M}\|_1 + \|\boldsymbol{\sigma}\|_1}{\rho} \exp \left(-\frac{\sqrt{1 - \beta_2} \|\mathbf{L}_0\|_1^{3/4}}{\sqrt{d} \|\mathbf{L}_1\|_\infty \|\boldsymbol{\sigma}\|_1^{1/2} \Delta^{1/4}} T^{1/4} \right) + \frac{\|\nabla F(\mathbf{x}_1)\|_1}{T} \right) .
\end{aligned}$$

Furthermore, for the case when $\beta_2 = 0$, we have the following refined guarantee:

$$\begin{aligned}
\min_{t \in [T]} \|\nabla F(\mathbf{x}_t)\|_1 &= \mathcal{O} \left(\frac{\sqrt{\log(dT/\delta)} \|\mathbf{L}_0\|_1^{1/4} \Delta^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2}}{T^{1/4}} + \frac{\log(dT/\delta) \sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\sqrt{T}} \right) \\
&\quad + \mathcal{O} \left(\frac{\|\nabla F(\mathbf{x}_1)\|_1}{\sqrt{T}} \left(\frac{1}{\sqrt{T}} + \frac{\|\boldsymbol{\sigma}\|_1}{\sqrt{\|\mathbf{L}_0\|_1 \Delta}} \right) + \frac{\|\boldsymbol{\sigma}\|_1}{T} \right) .
\end{aligned}$$

Proof of Theorem 1 for $\beta_2 = 0$. From Lemma 11 we know that $\|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty}$ for all $t \in [T]$. Thus, we can apply Lemma 1 to have

$$\begin{aligned}
&F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \\
&\leq \langle \nabla F(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \sum_{j=1}^d \frac{\left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2}{2} |x_{t+1,j} - x_{t,j}| \\
&= \langle \nabla F(\mathbf{x}_t), -\eta \text{sign}(\mathbf{m}_t) \rangle + \sum_{j=1}^d \frac{L_{0,j} + L_{1,j} \sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|}{2} \eta^2 \\
&= -\eta \|\nabla F(\mathbf{x}_t)\|_1 + \eta \langle \nabla F(\mathbf{x}_t), \text{sign}(\nabla F(\mathbf{x}_t)) - \text{sign}(\mathbf{m}_t) \rangle + \sum_{j=1}^d \frac{L_{0,j} + L_{1,j} \sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|}{2} \eta^2 \\
&= -\eta \|\nabla F(\mathbf{x}_t)\|_1 + 2\eta \sum_{j=1}^d \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \mathbb{I} \left[\text{sign} \left(\frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right) \neq \text{sign}(m_{t,j}) \right] + \sum_{j=1}^d \frac{L_{0,j} + L_{1,j} \sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|}{2} \eta^2 ,
\end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function and the first inequality uses Lemma 1,

Now, note that

$$\mathbb{I}\left[\text{sign}\left(\frac{\partial F}{\partial x_j}(\mathbf{x}_t)\right) \neq \text{sign}(m_{t,j})\right] \leq \mathbb{I}\left[\left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t) - m_{t,j}\right| \geq \left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t)\right|\right] \leq \frac{\left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t) - m_{t,j}\right|}{\left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t)\right|}.$$

Thus, denoting by $\boldsymbol{\epsilon}_t = \mathbf{m}_t - \nabla F(\mathbf{x}_t)$ gives

$$\begin{aligned} F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) &\leq -\eta \|\nabla F(\mathbf{x}_t)\|_1 + 2\eta \|\boldsymbol{\epsilon}_t\|_1 + \sum_{j=1}^d \frac{L_{0,j} + L_{1,j}\sqrt{d} \left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t)\right|}{2} \eta^2 \\ &= -\eta \|\nabla F(\mathbf{x}_t)\|_1 + 2\eta \|\boldsymbol{\epsilon}_t\|_1 + \frac{\|\mathbf{L}_0\|_1 \eta^2}{2} + \frac{\eta^2 \sqrt{d}}{2} \sum_{j=1}^d L_{1,j} \left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t)\right|. \end{aligned}$$

Sum both sides over $t = 1, \dots, T$, to have

$$F^* - F(\mathbf{x}_1) \leq -\eta \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 + 2\eta \sum_{t=1}^T \|\boldsymbol{\epsilon}_t\|_1 + \frac{\|\mathbf{L}_0\|_1 \eta^2 T}{2} + \frac{\eta^2 \sqrt{d}}{2} \sum_{t=1}^T \sum_{j=1}^d L_{1,j} \left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t)\right|. \quad (11)$$

Use Lemma 12 to bound each coordinate of $\sum_{t=1}^T \|\boldsymbol{\epsilon}_t\|_1$:

$$\begin{aligned} &\sum_{t=0}^{T-1} |\epsilon_{t+1,j}| \\ &\leq \sum_{t=0}^{T-1} \left[(1-\alpha)^t \left(\alpha \sigma_j + (1-\alpha) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| \right) + \frac{\eta L_{0,j}}{\alpha} + \alpha E_j \right] \\ &\quad + (1-\alpha) \eta \sqrt{d} L_{1,j} \sum_{t=0}^{T-1} \sum_{\tau=0}^{t-1} (1-\alpha)^\tau \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{t-\tau}) \right| \\ &= \sigma_j + \frac{1}{\alpha} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| + \frac{\eta L_{0,j} T}{\alpha} + \alpha E_j T + (1-\alpha) \eta \sqrt{d} L_{1,j} \sum_{t=1}^{T-1} \sum_{\tau'=1}^t (1-\alpha)^{t-\tau'} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{\tau'}) \right| \\ &= \sigma_j + \frac{1}{\alpha} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| + \frac{\eta L_{0,j} T}{\alpha} + \alpha E_j T + (1-\alpha) \eta \sqrt{d} L_{1,j} \sum_{\tau'=1}^{T-1} \left(\sum_{t=\tau'}^{T-1} (1-\alpha)^t \right) (1-\alpha)^{-\tau'} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_{\tau'}) \right| \\ &\leq \sigma_j + \frac{1}{\alpha} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| + \frac{\eta L_{0,j} T}{\alpha} + \alpha E_j T + \frac{(1-\alpha) \eta \sqrt{d} L_{1,j}}{\alpha} \sum_{t=1}^{T-1} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|. \end{aligned}$$

The above one holds with probability at least $1 - 3T\delta$ as we invoked Lemma 12 for T times which in turns means invoking Lemma 7 for T times. Yet, the above inequality would need to sum from $j = 1$ to d , meaning in total we would invoke Lemma 7 for dT times. Thus, following results hold with probability at least $1 - 3dT\delta$.

Now, put the above inequality back into (11) to have

$$\begin{aligned} &F^* - F(\mathbf{x}_1) \\ &\leq -\eta \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 + \frac{\|\mathbf{L}_0\|_1 \eta^2 T}{2} + \frac{\eta^2 \sqrt{d}}{2} \sum_{t=1}^T \sum_{j=1}^d L_{1,j} \left|\frac{\partial F}{\partial x_j}(\mathbf{x}_t)\right| \end{aligned}$$

$$\begin{aligned}
& + 2\eta \sum_{j=1}^d \left(\sigma_j + \frac{1}{\alpha} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_1) \right| + \frac{\eta L_{0,j} T}{\alpha} + \alpha E_j T + \frac{(1-\alpha)\eta\sqrt{d}L_{1,j}}{\alpha} \sum_{t=1}^T \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \\
& = -\eta \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 + \frac{\|\mathbf{L}_0\|_1 \eta^2 T}{2} + \frac{\eta^2 \sqrt{d}}{2} \sum_{t=1}^T \sum_{j=1}^d L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + 2\eta \|\boldsymbol{\sigma}\|_1 \\
& \quad + \frac{2\eta}{\alpha} \|\nabla F(\mathbf{x}_1)\|_1 + \frac{2\eta^2 \|\mathbf{L}_0\|_1 T}{\alpha} + 2\eta\alpha T \sum_{j=1}^d E_j + \frac{2\eta^2 \sqrt{d}}{\alpha} \sum_{t=1}^T \sum_{j=1}^d L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \\
& = -\eta \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 + 2\eta \|\boldsymbol{\sigma}\|_1 + \frac{2\eta}{\alpha} \|\nabla F(\mathbf{x}_1)\|_1 + \left(\frac{1}{2} + \frac{2}{\alpha} \right) \|\mathbf{L}_0\|_1 \eta^2 T + 2\eta\alpha T \sum_{j=1}^d E_j \\
& \quad + \left(\frac{1}{2} + \frac{2}{\alpha} \right) \eta^2 \sqrt{d} \sum_{t=1}^T \sum_{j=1}^d L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|.
\end{aligned}$$

Now, using the definitions of η and α , and the conditions on T , we have

$$\begin{aligned}
\left(\frac{1}{2} + \frac{2}{\alpha} \right) \eta^2 \sqrt{d} L_{1,j} & \leq \frac{\eta}{2} \left(1 + \frac{4}{\alpha} \right) \frac{\sqrt{d} \sqrt{\Delta \alpha}}{\sqrt{T}} \frac{\|\mathbf{L}_1\|_\infty}{\sqrt{\|\mathbf{L}_0\|_1}} \leq \frac{\eta}{2} \frac{5\sqrt{d} \sqrt{\Delta}}{\sqrt{\alpha T}} \frac{\|\mathbf{L}_1\|_\infty}{\sqrt{\|\mathbf{L}_0\|_1}} \\
& = \frac{\eta}{2} \frac{5\sqrt{d} \sqrt{\Delta}}{\sqrt{T}} \frac{\|\mathbf{L}_1\|_\infty}{\sqrt{\|\mathbf{L}_0\|_1}} \cdot \max \left(\frac{\sqrt{\|\boldsymbol{\sigma}\|_1} T^{1/4}}{\|\mathbf{L}_0\|_1^{1/4} \Delta^{1/4}}, 1 \right) \\
& = \frac{\eta}{2} \max \left(\frac{5\sqrt{d} \sqrt{\|\boldsymbol{\sigma}\|_1} \|\mathbf{L}_1\|_\infty \Delta^{1/4}}{\|\mathbf{L}_0\|_1^{3/4} T^{1/4}}, \frac{5\sqrt{d} \sqrt{\Delta}}{\sqrt{T}} \frac{\|\mathbf{L}_1\|_\infty}{\sqrt{\|\mathbf{L}_0\|_1}} \right) \leq \frac{\eta}{2}.
\end{aligned}$$

Thus, we have

$$F^* - F(\mathbf{x}_1) \leq -\frac{\eta}{2} \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 + 2\eta \|\boldsymbol{\sigma}\|_1 + \frac{2\eta}{\alpha} \|\nabla F(\mathbf{x}_1)\|_1 + \left(\frac{1}{2} + \frac{2}{\alpha} \right) \|\mathbf{L}_0\|_1 \eta^2 T + 2\eta\alpha T \sum_{j=1}^d E_j.$$

Divide both sides by T and rearrange terms to give

$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 & \leq \frac{2}{\eta T} [F(\mathbf{x}_1) - F^*] + \frac{4}{T} \|\boldsymbol{\sigma}\|_1 + \frac{4}{\alpha T} \|\nabla F(\mathbf{x}_1)\|_1 + \frac{5}{\alpha} \|\mathbf{L}_0\|_1 \eta \\
& \quad + 24 \|\boldsymbol{\sigma}\|_1 (\alpha \max(1, \log(1/\delta)) + \sqrt{\alpha} \sqrt{\max(1, \log(1/\delta))}).
\end{aligned}$$

Now, we need to consider the following two cases:

$$\begin{aligned}
1. \quad & \|\boldsymbol{\sigma}\|_1 < \frac{\sqrt{\|\mathbf{L}_0\|_1} \sqrt{\Delta}}{\sqrt{T}}: \text{ then } \alpha = 1 \text{ and } \eta = \frac{\sqrt{\Delta}}{\sqrt{\|\mathbf{L}_0\|_1} \sqrt{T}} \\
& \frac{1}{T} \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 \leq \frac{2\sqrt{\|\mathbf{L}_0\|_1}}{\sqrt{\Delta} \sqrt{T}} [F(\mathbf{x}_1) - F^*] + \frac{5\|\mathbf{L}_0\|_1 \sqrt{\Delta}}{\sqrt{\|\mathbf{L}_0\|_1} \sqrt{T}} + \frac{4\sqrt{\|\mathbf{L}_0\|_1} \sqrt{\Delta}}{T^{3/2}} + \frac{4}{T} \|\nabla F(\mathbf{x}_1)\|_1 \\
& \quad + \frac{24\sqrt{\|\mathbf{L}_0\|_1} \sqrt{\Delta} (\max(1, \log(1/\delta)) + \sqrt{\max(1, \log(1/\delta))})}{\sqrt{T}} \\
& \leq \frac{59 \max(1, \log(1/\delta)) \sqrt{\|\mathbf{L}_0\|_1} \Delta}{\sqrt{T}} + \frac{4}{T} \|\nabla F(\mathbf{x}_1)\|_1. \tag{12}
\end{aligned}$$

$$2. \quad \|\boldsymbol{\sigma}\|_1 \geq \frac{\sqrt{\|\mathbf{L}_0\|_1} \Delta}{\sqrt{T}}: \text{ then } \alpha = \frac{\sqrt{\|\mathbf{L}_0\|_1} \Delta}{\|\boldsymbol{\sigma}\|_1 \sqrt{T}} \leq 1 \text{ and } \eta = \frac{\Delta^{3/4}}{\|\mathbf{L}_0\|_1^{1/4} \sqrt{\|\boldsymbol{\sigma}\|_1} T^{3/4}}$$

$$\frac{1}{T} \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1$$

$$\begin{aligned}
&\leq \frac{2\|\mathbf{L}_0\|_1^{1/4}\sqrt{\|\boldsymbol{\sigma}\|_1}}{\Delta^{3/4}T^{1/4}}[F(\mathbf{x}_1) - F^*] + \frac{4\|\boldsymbol{\sigma}\|_1}{T} + \frac{4\|\boldsymbol{\sigma}\|_1}{\sqrt{\|\mathbf{L}_0\|_1\Delta T}}\|\nabla F(\mathbf{x}_1)\|_1 + \frac{5\|\mathbf{L}_0\|_1^{1/4}\sqrt{\|\boldsymbol{\sigma}\|_1}\Delta^{1/4}}{T^{1/4}} \\
&\quad + \frac{24\max(1, \log(1/\delta))\sqrt{\|\mathbf{L}_0\|_1\Delta}}{\sqrt{T}} + \frac{24\sqrt{\max(1, \log(1/\delta))}\|\mathbf{L}_0\|_1^{1/4}\sqrt{\|\boldsymbol{\sigma}\|_1}\Delta^{1/4}}{T^{1/4}} \\
&\leq \frac{31\sqrt{\max(1, \log(1/\delta))}\|\mathbf{L}_0\|_1^{1/4}\Delta^{1/4}\sqrt{\|\boldsymbol{\sigma}\|_1}}{T^{1/4}} + \frac{24\max(1, \log(1/\delta))\sqrt{\|\mathbf{L}_0\|_1\Delta}}{\sqrt{T}} \\
&\quad + \frac{4\|\boldsymbol{\sigma}\|_1\|\nabla F(\mathbf{x}_1)\|_1}{\sqrt{\|\mathbf{L}_0\|_1\Delta T}} + \frac{4\|\boldsymbol{\sigma}\|_1}{T}.
\end{aligned}$$

Put $\delta' = \frac{\delta}{3dT}$ concludes the proof. \square

Proof of Theorem 1 for general β_2 . Note that when $\|\boldsymbol{\sigma}\|_1 \leq \frac{\sqrt{\|\mathbf{L}_0\|_1\Delta}}{\sqrt{T}}$, $\alpha = 1$, $\beta_2 < \beta_1^2 = 0$. Then our Generalized SignSGD algorithm 1 reduces to the SignSGD algorithm, and thus has the same guarantee of (12). Therefore, we only consider the other case when $\|\boldsymbol{\sigma}\|_1 \geq \frac{\sqrt{\|\mathbf{L}_0\|_1\Delta}}{\sqrt{T}}$.

Note that the only randomness comes from evaluating stochastic gradients. In the following proof, we will need to invoke Lemma 7 for T times for each coordinate $j \in [d]$. Therefore, the following results hold with probability at least $1 - 3dT\delta$. For simplicity, we use the term "with high probability" later in the proof to denote this.

We derive the following quantity which will be used multiple times later:

$$\frac{\eta}{\alpha} = \frac{1}{\sqrt{\|\mathbf{L}_0\|_1}} \cdot \frac{\sqrt{\Delta}}{\sqrt{T}} \cdot \frac{1}{\sqrt{\alpha}} = \frac{1}{\sqrt{\|\mathbf{L}_0\|_1}} \cdot \frac{\sqrt{\Delta}}{\sqrt{T}} \cdot \frac{\|\boldsymbol{\sigma}\|_1^{1/2}T^{1/4}}{\|\mathbf{L}_0\|_1^{1/4}\Delta^{1/4}} = \frac{\|\boldsymbol{\sigma}\|_1^{1/2}\Delta^{1/4}}{\|\mathbf{L}_0\|_1^{3/4}T^{1/4}}. \quad (13)$$

First, from Lemma 10 we have $D = 1 - \frac{2\sqrt{d}\|\mathbf{L}_1\|_\infty\eta}{(1-\beta_1)\sqrt{1-\beta_2}}$. Then, from the choice of the hyperparameters, we have

$$\frac{\sqrt{d}\|\mathbf{L}_1\|_\infty\eta}{(1-\beta_1)\sqrt{1-\beta_2}} = \frac{\|\mathbf{L}_1\|_\infty}{\sqrt{1-\beta_2}} \cdot \frac{\sqrt{d}\|\boldsymbol{\sigma}\|_1^{1/2}\Delta^{1/4}}{\|\mathbf{L}_0\|_1^{3/4}T^{1/4}} \leq \frac{\rho}{10} \leq \frac{1}{10}.$$

Thus, we have $D \geq 1 - \frac{1}{5} \geq \frac{1}{2}$ and, as $\sqrt{\beta_2}/\beta_1 < 1$,

$$\frac{\rho D}{5\sqrt{1-\beta_2}} \geq \frac{\rho}{10\sqrt{1-\beta_2}} = A. \quad (14)$$

Also, for those coordinates with small gradients $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| < \frac{5B_j}{D} \leq 10B_j$, we have

$$\begin{aligned}
\frac{\partial F}{\partial x_j}(\mathbf{x}_t) \cdot (x_{t+1,j} - x_{t,j}) &= -\eta \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \cdot \frac{m_{t,j}}{\sqrt{v_{t,j}}} \\
&= -A\eta \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + \eta \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \cdot \left(A - \text{sign} \left(\frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right) \frac{m_{t,j}}{\sqrt{v_{t,j}}} \right) \\
&\leq -A\eta \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + 10B_j\eta \left(\frac{1}{\sqrt{1-\beta_2}} + A \right).
\end{aligned} \quad (15)$$

We are now ready to prove the theorem. We will need to use Lemma 10, hence we first need to show that all past true gradients are bounded by M_j , namely that, for any t , $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right| \leq M_j$ holds for all $\tau \leq t$ and all $j \in [d]$. From the definition of M_j stated in the theorem, in order to guarantee this, we only need to prove that $F(\mathbf{x}_\tau) \leq F(\mathbf{x}_1)$ for all $\tau \leq t$. We will prove this by induction.

For $t = 1$ the condition is trivially true.

We then assume that the condition holds for t and prove based on this that it still holds for $t + 1$.

For those coordinates with $\left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \geq \frac{5B_j}{D}$, denote $\hat{\tau} = \lfloor \bar{\tau} \rfloor = \left\lfloor \frac{\sqrt{1-\beta_2}}{\eta\sqrt{d}\|\mathbf{L}_1\|_\infty} \right\rfloor$, with high probability, we have

$$\begin{aligned}
\left| m_{t,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| &= \left| (1 - \beta_1) \sum_{\tau=1}^t \beta_1^{t-\tau} g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \\
&\leq \left| (1 - \beta_1) \sum_{\tau=1}^{t-\hat{\tau}-1} \beta_1^{t-\tau} g_{\tau,j} \right| + \left| (1 - \beta_1) \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \\
&\leq \beta_1^{\bar{\tau}} (M_j + \sigma_j) + \beta_1^{\bar{\tau}} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \\
&\quad + (1 - \beta_1) \left| \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} \left(\frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) - \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right) \right| \\
&\quad + (1 - \beta_1) \left| \sum_{\tau=t-\hat{\tau}}^t \beta_1^{t-\tau} \left(g_{\tau,j} - \frac{\partial F}{\partial x_j}(\mathbf{x}_\tau) \right) \right| \\
&\leq (1 - D) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + B_j \\
&\leq \left(1 - \frac{4D}{5} \right) \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \leq \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|, \tag{16}
\end{aligned}$$

where the first equality comes from Lemma 4; for the second inequality, the third term can be bounded using Lemma 8, and the final term can be bounded using Lemma 7; for the third inequality, we used (4) and that $e^{-x} \leq \frac{1}{x}$ for $x > 0$. This inequality implies that $\text{sign}(m_{t,j}) = \text{sign} \left(\frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right)$ with high probability.

Denote $\mathcal{U}_t = \left\{ j \in [d] : \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \geq \frac{5B_j}{D} \right\}$. From the choices of hyperparameters we can show that $1 \leq \bar{\tau}$ which means $\|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2 \leq \frac{1}{\|\mathbf{L}_1\|_\infty}$ (Lemma 5). Thus, using Lemma 1, with high probability we have

$$\begin{aligned}
&F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \\
&\leq \langle \nabla F(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \sum_{j=1}^d \frac{\left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2}{2} |x_{t+1,j} - x_{t,j}| \\
&= \sum_{j=1}^d \left(-\frac{\partial F}{\partial x_j}(\mathbf{x}_t) \cdot \eta \text{sign}(m_{t,j}) \frac{|m_{t,j}|}{\sqrt{v_{t,j}}} + \frac{\left(\frac{L_{0,j}}{\sqrt{d}} + L_{1,j} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \right) \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2}{2} |x_{t+1,j} - x_{t,j}| \right) \\
&\leq \sum_{j \in \mathcal{U}_t} -\eta \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| \cdot \frac{|m_{t,j}|}{\sqrt{v_{t,j}}} + \sum_{j \notin \mathcal{U}_t} \left(-A\eta \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right| + 10\eta B_j \left(\frac{1}{\sqrt{1-\beta_2}} + A \right) \right) + \sum_{j=1}^d \frac{L_{0,j} + L_{1,j}\sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|}{2(1-\beta_2)} \eta^2 \\
&\leq -A\eta \|\nabla F(\mathbf{x}_t)\|_1 + \sum_{j=1}^d \frac{L_{0,j} + L_{1,j}\sqrt{d} \left| \frac{\partial F}{\partial x_j}(\mathbf{x}_t) \right|}{2(1-\beta_2)} \eta^2 + 10\eta \left(\frac{1}{\sqrt{1-\beta_2}} + A \right) \sum_{j=1}^d B_j, \tag{17}
\end{aligned}$$

where the second inequality uses (15), (16), and Lemma 4, and the third inequality uses Lemma 10 and (14).

Now, noticing the conditions on $\eta, \alpha, \beta_2 < \beta_1^2 < \beta_1$, and T , use (13) to have

$$\frac{\eta^2 \sqrt{d} L_{1,j}}{2(1-\beta_2)} \leq \frac{\eta \sqrt{d} \|\mathbf{L}_1\|_\infty}{2} \cdot \frac{\eta}{\alpha} = \frac{\eta}{2} \frac{\sqrt{d} \|\mathbf{L}_1\|_\infty \|\boldsymbol{\sigma}\|_1^{1/2} \Delta^{1/4}}{\|\mathbf{L}_0\|_1^{3/4} T^{1/4}} \leq \frac{\eta}{2} \frac{\rho \sqrt{1-\beta_2}}{10} \leq \frac{\eta}{2} A.$$

Thus, (17) becomes

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \leq -\frac{A\eta}{2} \|\nabla F(\mathbf{x}_t)\|_1 + \frac{\eta^2 \|\mathbf{L}_0\|_1}{2(1-\beta_2)} + 10\eta \left(\frac{1}{\sqrt{1-\beta_2}} + A \right) \sum_{j=1}^d B_j. \quad (18)$$

Therefore, either

$$\|\nabla F(\mathbf{x}_t)\|_1 \leq \frac{\eta \|\mathbf{L}_0\|_1}{A(1-\beta_2)} + 20 \left(\frac{1}{A\sqrt{1-\beta_2}} + 1 \right) \sum_{j=1}^d B_j, \quad (19)$$

or $F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \leq 0$.

This concludes the mathematical induction up until (19) is met for the first time which we denote as T_0 . In the following, we will explain that if (19) holds then the algorithm has found an approximate stationary point.

Now, suppose $T \leq T_0$, then (18) holds all the time and we sum both sides of it from 1 to T to have, with high probability,

$$F^* - F(\mathbf{x}_1) \leq -\frac{A\eta}{2} \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 + \frac{\|\mathbf{L}_0\|_1 \eta^2 T}{2(1-\beta_2)} + 10\eta T \left(\frac{1}{\sqrt{1-\beta_2}} + A \right) \sum_{j=1}^d B_j.$$

Rearrange terms to obtain

$$\min_{t \in [T]} \|\nabla F(\mathbf{x}_t)\|_1 \leq \frac{1}{T} \sum_{t=1}^T \|\nabla F(\mathbf{x}_t)\|_1 \leq \frac{2}{A\eta T} [F(\mathbf{x}_1) - F^*] + \frac{\eta \|\mathbf{L}_0\|_1}{A(1-\beta_2)} + 20 \left(\frac{1}{A\sqrt{1-\beta_2}} + 1 \right) \sum_{j=1}^d B_j. \quad (20)$$

Note that RHS of (19) is less than RHS of (20). Thus, for the other case of $T > T_0$, (20) still holds.

Recalling that $\rho = 1 - \frac{\sqrt{\beta_2}}{\beta_1}$, $A = \frac{\rho}{10\sqrt{1-\beta_2}}$, $\beta_2 < \beta_1^2 < \beta_1$, and $B_j \triangleq \frac{\eta L_{0,j}}{(1-\beta_1)\sqrt{1-\beta_2}} + \beta_1^{\bar{\tau}} (M_j + \sigma_j) + 6(1-\beta_1)\sigma_j \max(1, \log(1/\delta)) + \frac{6(1-\beta_1)}{\sqrt{1-\beta_1^2}} \sqrt{\sigma_j^2 \max(1, \log(1/\delta))}$, we have

$$\begin{aligned} \min_{t \in [T]} \|\nabla F(\mathbf{x}_t)\|_1 &\leq \frac{20}{\rho\eta T} [F(\mathbf{x}_1) - F^*] + \frac{10\eta \|\mathbf{L}_0\|_1}{\rho\sqrt{1-\beta_1}} + 20 \left(\frac{10}{\rho} + 1 \right) \left(\frac{\|\mathbf{L}_0\|_1 \eta}{(1-\beta_1)\sqrt{1-\beta_2}} + \beta_1^{\bar{\tau}} (\|\mathbf{M}\|_1 + \|\boldsymbol{\sigma}\|_1) \right) \\ &\quad + 120 \|\boldsymbol{\sigma}\|_1 (1-\beta_1) \left(\frac{10}{\rho} + 1 \right) \left(\max(1, \log(1/\delta)) + \frac{1}{\sqrt{1-\beta_1^2}} \sqrt{\max(1, \log(1/\delta))} \right). \end{aligned}$$

When $\|\boldsymbol{\sigma}\|_1 \geq \frac{\sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\sqrt{T}}$, then $\alpha = \frac{\sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\|\boldsymbol{\sigma}\|_1 \sqrt{T}}$ and $\eta = \frac{\Delta^{3/4}}{\|\mathbf{L}_0\|_1^{1/4} \sqrt{\|\boldsymbol{\sigma}\|_1 T^{3/4}}}$. Hence, we obtain

$$\begin{aligned} &\min_{t \in [T]} \|\nabla F(\mathbf{x}_t)\|_1 \\ &\leq \frac{20\Delta^{1/4} \|\mathbf{L}_0\|_1^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2} T^{3/4}}{\rho T} + \frac{10\sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\rho\sqrt{T}} + 20 \left(\frac{10}{\rho} + 1 \right) \left(\frac{\Delta^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2} \|\mathbf{L}_0\|_1}{\sqrt{1-\beta_2} \|\mathbf{L}_0\|_1^{3/4} T^{1/4}} + \beta_1^{\bar{\tau}} (\|\mathbf{M}\|_1 + \|\boldsymbol{\sigma}\|_1) \right) \\ &\quad + 120 \left(\frac{10}{\rho} + 1 \right) \left(\frac{\sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\sqrt{T}} \max(1, \log(1/\delta)) + \frac{\Delta^{1/4} \|\mathbf{L}_0\|_1^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2}}{T^{1/4}} \sqrt{\max(1, \log(1/\delta))} \right) \\ &\leq \left(\frac{20}{\rho T^{1/4}} + \left(\frac{10}{\rho} + 1 \right) \frac{20 + 120\sqrt{\max(1, \log(1/\delta))}}{\sqrt{1-\beta_2} T^{1/4}} \right) \Delta^{1/4} \|\mathbf{L}_0\|_1^{1/4} \|\boldsymbol{\sigma}\|_1^{1/2} \\ &\quad + \left(\frac{10}{\rho} + 120 \left(\frac{10}{\rho} + 1 \right) \max(1, \log(1/\delta)) \right) \frac{\sqrt{\|\mathbf{L}_0\|_1 \Delta}}{\sqrt{T}} + 20 \left(\frac{10}{\rho} + 1 \right) \beta_1^{\bar{\tau}} (\|\mathbf{M}\|_1 + \|\boldsymbol{\sigma}\|_1) \end{aligned}$$

Table 2: Hyperparameter grid search ranges and choices yielding the highest validation accuracy for each optimizer for training a 20-layer Resnet to do image classification on CIFAR-10. ("lr" denotes the initial learning rate, "clip" denotes the clipping parameter γ in Algorithm 1 of [55], and " β_2 " is defined in [25] for Adam and in Algorithm 1 for ours.)

Optimizer	Grid Search Range	Best Choice
SGD Momentum	lr {1e-5, 0.0001, 0.001, 0.01, 0.05, 0.07, 0.1, 0.2, 0.3, 1, 10}	lr=0.07
SGD Momentum Normalized	lr {0.0001, 0.001, 0.01, 0.05, 0.07, 0.09, 0.1, 0.2, 0.3, 0.5 1, 10}	lr=0.1
SGDClipGrad	lr {0.001, 0.01, 0.05, 0.1, 0.5, 1, 10} clip {0.1, 1, 10}	lr=0.5 clip=1
SGDClipMomentum	lr {0.001, 0.01, 0.1, 1, 5, 10, 20, 50} clip {0.01, 0.1, 1, 10}	lr=10 clip=0.1
Adam	lr {1e-5, 0.0001, 0.0007, 0.0009, 0.001, 0.002, 0.003, 0.01, 0.1} β_2 {0.4, 0.8, 0.999}	lr=0.0009 β_2 =0.999
Our Algorithm 1	lr {5e-5, 8e-5, 0.0001, 0.0002, 0.0005, 0.001, 0.01} β_2 {0.4, 0.8, 0.999}	lr=0.0002 β_2 = 0.999

$$\begin{aligned} &\leq \frac{1560\Delta^{1/4}\|\mathbf{L}_0\|_1^{1/4}\|\boldsymbol{\sigma}\|_1^{1/2}\sqrt{\max(1, \log(1/\delta))}}{\rho\sqrt{1-\beta_2}T^{1/4}} + \frac{1330\max(1, \log(1/\delta))\sqrt{\|\mathbf{L}_0\|_1\Delta}}{\rho\sqrt{T}} \\ &\quad + \frac{220}{\rho}(\|\mathbf{M}\|_1 + \|\boldsymbol{\sigma}\|_1)\exp\left(-\frac{\sqrt{1-\beta_2}\|\mathbf{L}_0\|_1^{3/4}}{\sqrt{d}\|\mathbf{L}_1\|_\infty\|\boldsymbol{\sigma}\|_1^{1/2}\Delta^{1/4}}T^{1/4}\right). \end{aligned}$$

Finally, taking $\delta' = \frac{\delta}{3dT}$, we obtain the stated result. \square

A.4 More Experiment Details

In this section, we provide more details for the experiments we show in Section 5.

Hyperparameter Tuning During the validation stage, we used grid-search to fine-tune respective hyperparameters and choose the ones that yield the best validation results. We tuned the hyperparameters using the following two-stage grid searching strategy: First, search over a coarse grid, and select the one yielding the best validation result. Next, continue searching in a fine grid centering at the best-performing hyperparameters found in the coarse stage, and in turn, take the best one as the final choice. Also, whenever the best-performing hyperparameters lie in the boundary of the searching grid, we always extend the grid to make the final best-performing hyperparameters fall into the interior of the grid, if possible.

Resnet on CIFAR-10 We randomly selected 10% images from the training dataset for validation. Yet, during testing, we trained on the whole training dataset. The detailed search ranges and the hyperparameter choices yielding the highest validation accuracy for each optimizer are listed in Table 2.

AWD-LSTM on Penn Treebank We used the original train-validation-test split that comes with the dataset. The momentum parameter (β_1) is fixed to be 0.9 except for SGDClipGrad which does not use momentum. The detailed search ranges and the hyperparameter choices yielding the lowest validation perplexity for each optimizer are listed in Table 3.

A.5 Training a Transformer Model on WMT’16 German-English Translation Task

We noted that Transformers [50] are gaining huge popularities recently and reported in Figure 1 and 3 that Transformers observe the relaxed smoothness conditions. Thus, to further showcase the effectivity of our algorithm 1 compared with other optimizers listed in 5.1, we train a 6-layer Transformer model to do machine translation on the WMT’16 Multimodal Machine Translation Task German-English dataset. The implementation of the transformer is forked from here³ and we inherited the default model structure. We also adopted the warm-up steps of 128000 and the learning rate decay strategy as recommended by the GitHub repo. The mini-batch size is 256 and we trained for 400 epochs.

³<https://github.com/jadore801120/attention-is-all-you-need-pytorch>

Table 3: Hyperparameter grid search ranges and choices yielding the lowest validation perplexity for each optimizer for training an AWD-LSTM to do language modeling on Penn Treebank. ("wd" denotes the weight decay value, "lr" denotes the initial learning rate, "clip" denotes the clipping parameter γ in Algorithm 1 of [55], and " β_2 " is defined in [25] for Adam and in Algorithm 1 for ours.)

Optimizer	Grid Search Range	Best Choice
SGD Momentum	wd {1e-7, 1.2e-6, 5e-6, 1e-5, 1e-4, 1e-3} lr {0.001, 0.01, 0.1, 0.5, 0.8, 1, 2, 4, 5}	wd=1e-5 lr=1
SGD Momentum Normalized	wd {1e-7, 1.2e-6, 5e-6, 1e-5, 5e-5, 1e-4} lr {0.01, 0.05, 0.1, 0.5, 0.8, 1, 2, 4, 5, 10}	wd=5e-6 lr=2
SGDClipGrad	wd {1e-7, 1.2e-6, 5e-6, 1e-5} lr {0.1, 0.5, 1, 5, 10, 20, 30, 40, 50, 60, 70} clip {1, 2.5, 7.5, 10, 15, 20}	wd=1.2e-6 lr=50 clip=10
SGDClipMomentum	wd {1e-7, 1.2e-6, 5e-6, 1e-5} lr {5, 10, 20, 30, 50, 100} clip {1, 2.5, 7.5}	wd=1.2e-6 lr=20 clip=2.5
Adam	wd {1e-7, 1.2e-6, 5e-6, 1e-5} lr {0.0001, 0.001, 0.002, 0.003, 0.01, 0.1} β_2 {0.4, 0.8, 0.999}	wd=5e-6 lr=0.002 β_2 =0.999
Our Algorithm 1	wd {1e-7, 1.2e-6, 5e-6, 1e-5} lr {0.0001, 0.001, 0.002, 0.003, 0.01, 0.1} β_2 {0.4, 0.8, 0.999}	wd=1.2e-6 lr=0.001 β_2 =0.999

Table 4: Average final training loss, test perplexity, and test accuracy achieved by each method when optimizing the Transformer model on the WMT'16 Multimodal Machine Translation Task German-English dataset. The \pm shows 95% confidence intervals of the mean value over 5 runs starting from different random seeds.

Methods	Training loss	Test perplexity	Test accuracy
SGD Momentum	2.8045 \pm 0.0209	10.4319 \pm 0.1973	63.9108 \pm 0.5797
SGD Momentum Normalized	2.9268 \pm 0.0512	10.5793 \pm 0.6383	63.0774 \pm 0.8231
SGDClipGrad	3.0214 \pm 0.0508	10.5974 \pm 0.3527	62.2534 \pm 0.3145
SGDClipMomentum	2.8128 \pm 0.0295	10.4677 \pm 0.3619	63.6562 \pm 0.5370
Adam	1.4303 \pm 0.0009	8.9088 \pm 0.1294	68.9828 \pm 0.2786
Our Algorithm 1	1.6263 \pm 0.0024	7.2731 \pm 0.0870	68.5790 \pm 0.4693

For the hyperparameter tuning, the momentum parameter (β_1) is fixed to be 0.9 except for SGDClipGrad which does not use momentum. We then used grid-search to fine-tune the initial learning rate and the weight decay value for all optimizers, as well as the clipping threshold for SGDClipGrad and SGDClipMomentum, and β_2 for Adam and our algorithm. We select the combination of hyperparameters that gives the lowest validation perplexity. The detailed search ranges and the hyperparameter choices yielding the lowest validation perplexity for each optimizer are listed in Table 3.

We then employ the best-performing hyperparameters to report the testing performance. The testing is repeated with random seeds 5 times to eliminate the influence of stochasticity. The results are reported in Figure 6 and Table 4. Here the accuracy means the proportion of correct correspondences of words, namely the same word at the same location, between the machine-translated output and the target. It can be seen that we can still match the performance of Adam while beating the others. Also, note that the curves of SGD Momentum and the ones of SGDClipMomentum overlap as they utilize the same weight decay values and initial learning rates, and turns out clipping is seldomly performed when employing SGDClipMomentum.

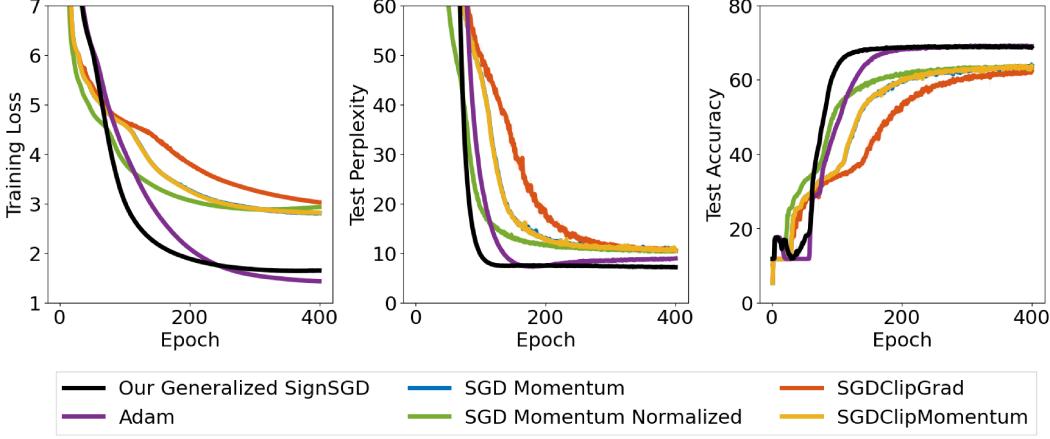


Figure 6: Training a 6-layer Transformer model to do machine translation on the WMT’16 Multimodal Machine Translation Task German-English dataset. The shading of each curve represents the 95% confidence interval computed across five independent runs from different random seeds.

Table 5: Hyperparameter grid search ranges and choices yielding the lowest validation perplexity for each optimizer for training a 6-layer Transformer to do machine translation on the WMT’16 Multimodal Machine Translation Task German-English dataset. (“wd” denotes the weight decay value, “lr” denotes the initial learning rate, “clip” denotes the clipping parameter γ in Algorithm 1 of [55], and “ β_2 ” is defined in [25] for Adam and in Algorithm 1 for ours.)

Optimizer	Grid Search Range	Best Choice
SGD Momentum	wd {1e-6, 0.001, 0.01, 0.1, 1, 10, 100, 1000} lr {0.1, 1, 10, 100, 1000, 10000}	wd=1 lr=1
SGD Momentum Normalized	wd {1e-6, 0.001, 0.01, 0.1, 1, 10, 100, 1000} lr {0.01, 0.1, 1, 10, 100, 1000, 10000, 100000}	wd=1 lr=10000
SGDClipGrad	wd {1e-6, 0.01, 0.1, 1, 10, 100, 1000} lr {0.01, 0.1, 1, 10, 100, 1000, 10000} clip {0.01, 0.1, 1, 10}	wd=1 lr=10 clip=1
SGDClipMomentum	wd {1e-6, 0.01, 0.1, 1, 10} lr {0.01, 0.1, 1, 10, 100, 1000, 10000} clip {0.01, 0.1, 1, 10}	wd=1 lr=1 clip=1
Adam	wd {1e-6, 0.001, 0.01, 0.1, 1, 10, 100} lr {0.1, 1, 10, 100, 1000, 10000} β_2 {0.9, 0.98, 0.999}	wd=0.1 lr=10 $\beta_2=0.98$
Our Algorithm 1	wd {1e-6, 0.001, 0.01, 0.1, 1, 10, 100} lr {0.1, 1, 10, 100, 1000, 10000} β_2 {0.9, 0.98, 0.999}	wd=0.1 lr=10 $\beta_2=0.98$