

A Methodology Establishing Linear Convergence of Adaptive Gradient Methods Under PL Inequality

Kushal Chakrabarti^{a,*} and Mayank Baranwal^{a,b}

^aTata Consultancy Services Research, Mumbai, India

^bDepartment of Systems & Control Engineering, Indian Institute of Technology, India

ORCID (Kushal Chakrabarti): <https://orcid.org/0000-0002-6747-8709>, ORCID (Mayank Baranwal):

<https://orcid.org/0000-0001-9354-2826>

Abstract. Adaptive gradient-descent optimizers are the standard choice for training neural network models. Despite their faster convergence than gradient-descent and remarkable performance in practice, the adaptive optimizers are not as well understood as vanilla gradient-descent. A reason is that the dynamic update of the learning rate that helps in faster convergence of these methods also makes their analysis intricate. Particularly, the simple gradient-descent method converges at a *linear* rate for a class of optimization problems, whereas the practically faster adaptive gradient methods lack such a theoretical guarantee. The Polyak-Łojasiewicz (PL) inequality is the weakest known class, for which *linear* convergence of gradient-descent and its momentum variants has been proved. Therefore, in this paper, we prove that AdaGrad and Adam, two well-known adaptive gradient methods, converge *linearly* when the cost function is smooth and satisfies the PL inequality. Our theoretical framework follows a simple and unified approach, applicable to both batch and stochastic gradients, which can potentially be utilized in analyzing *linear* convergence of other variants of Adam.

1 Introduction

In this paper, we consider the problem of minimizing a possibly non-convex *objective function* $f : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\min_{x \in \mathbb{R}^d} f(x). \quad (1)$$

Among other applications, non-convex optimization appears in training neural network models. It is a standard practice to use adaptive gradient optimizers for training such models. Compared to gradient-descent, these methods have been observed to converge faster and do not require line search to determine the learning rates.

AdaGrad [7] is possibly the earliest adaptive gradient optimizer. To address the gradient accumulation problem in AdaGrad, the Adam algorithm [11] was proposed. Adam and its variants have been widely used to train deep neural networks in the past decade. Despite the success of these adaptive gradient optimizers, we lack an understanding of why these methods work so well in practice. The theory of convergence of the adaptive optimizers has not been completely developed. Since these methods work better than simple gradient-descent or its momentum variants, it is natural to expect a similar or better convergence guarantee of adaptive optimizers compared

to gradient-descent. One such theoretical difference between these two classes of methods is *linear* convergence. Specifically, gradient-descent and its accelerated variants, such as the Nesterov accelerated gradient or Heavy-Ball method, are known to exhibit *linear* convergence for a class of cost functions [16]. On the other hand, most of the adaptive gradient optimizers lack such a guarantee.

Linear convergence guarantees of gradient-descent, its accelerated variants, coordinate descent, and AdaGrad-Norm [22] (among the adaptive gradient methods) have been proved for the class of smooth and possibly non-convex cost functions that satisfy the Polyak-Łojasiewicz (PL) inequality [10]. The PL inequality is the weakest condition among others, such as strong convexity, essential strong convexity, weak strong convexity, and restricted secant inequality, that leads to *linear* convergence of gradient-descent and its accelerated variants to the solution of (1) [10]. The objective functions in standard machine learning problems like linear regression and logistic regression satisfy the PL inequality. For solving over-parameterized non-linear equations, [13] establishes a relation between PL inequality and the condition number of the tangent kernel, and argued that sufficiently wide neural networks generally satisfy the PL inequality. Motivated by *linear* convergence guarantees of gradient-descent and its aforementioned variants under the PL condition and the applicability of PL inequality on a set of machine learning problems, we investigate *linear* convergence of AdaGrad and Adam under the PL inequality.

To present our results, we define the following notations and make a set of assumptions as stated below. Let the gradient of f evaluated at $x \in \mathbb{R}^d$ be denoted by $\nabla f(x) \in \mathbb{R}^d$, and its i -th element be denoted by $\nabla_i f(x)$ for each dimension $i \in \{1, \dots, d\}$.

Assumption 1. *The minimum f_* of f exists and is finite, i.e., $|\min_{x \in \mathbb{R}^d} f(x)| < \infty$.*

Assumption 2. *f is twice differentiable over its domain \mathbb{R}^d and is L -smooth, i.e., $\exists L > 0$ such that $\|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\|$ for all $x, y \in \mathbb{R}^d$.*

Assumption 3. *f satisfies the Polyak-Łojasiewicz (PL) inequality, i.e., $\exists \mu > 0$ such that $\frac{1}{2} \|\nabla f(x)\|^2 \geq \mu(f(x) - f_*)$ for all $x \in \mathbb{R}^d$.*

Assumption 3 has been justified in the preceding paragraph. Assumption 1-2 are standard in the literature of gradient-based opti-

* Corresponding Author. Email: chakrabarti.k@tcs.com.

mization. An implication of Assumption 2 is that [2]

$$f(y) - f(x) \leq (y - x)^\top \nabla f(x) + \frac{L}{2} \|x - y\|^2, \forall x, y \in \mathbb{R}^d. \quad (2)$$

Next, we review the AdaGrad [7] algorithm and its existing convergence guarantees. AdaGrad and Adam are iterative methods. At each iteration $k = 0, 1, \dots$, AdaGrad maintains an estimate x_k of the solution of (1) and an auxiliary state y_k that determines the learning rate. x_k and y_k are updated as

$$y_{k+1,i} = y_{k,i} + |\nabla_i f(x_k)|^2, \quad (3a)$$

$$x_{k+1,i} = x_{k,i} - h \frac{\nabla_i f(x_k)}{\sqrt{y_{k+1,i}} + \epsilon}. \quad (3b)$$

Here, $h > 0$ is a stepsize, and $\epsilon > 0$ is a small-valued parameter to avoid division by zero. The original AdaGrad paper [7] proves $\mathcal{O}(\sqrt{K})$ regret bound for convex f , where K is the number of iterations, assuming the stochastic gradients and the estimates of minima are uniformly bounded. Under the same boundedness assumptions, for strongly convex f , AdaGrad's regret bound analysis [6] imply a $\mathcal{O}(\log(K)/K)$ convergence rate [4]. Following [7], there are several works on the convergence of AdaGrad. The notable among them include [12, 5, 20, 18]. Among them, [18] most recently proved $\mathcal{O}(1/K)$ convergence rate of AdaGrad for non-convex f , without the bounded gradient assumption. To the best of our knowledge, *linear* convergence of AdaGrad has not been proved in the literature. On the other hand, the gradient-descent algorithm is known to have *linear* convergence when f is *smooth* and satisfies the PL inequality [10]. So, we ask the question of whether the adaptive gradient-descent methods, such as AdaGrad and Adam, can have guaranteed *linear* convergence for the same class of optimization problems.

Among the variants of AdaGrad, [22] proves *linear* convergence of the AdaGrad-Norm (norm version of AdaGrad) in the stochastic settings under strong convexity of f and in the deterministic setting under PL inequality. The AdaGrad-Norm algorithm, described as

$$y_{k+1} = y_k + \|\nabla f(x_k)\|^2, \\ x_{k+1} = x_k - h \frac{\nabla f(x_k)}{\sqrt{y_{k+1}} + \epsilon},$$

differs from the coordinate-wise update in AdaGrad (3a)-(3b), which is more commonly used in practice [18]. So, the analysis in [22] does not trivially extend to AdaGrad. Furthermore, in the deterministic settings, if y_0 is initialized as a small value, AdaGrad-Norm is shown to have a *sublinear* rate until a finite number of iterations where y_k crosses a certain threshold [22]. In the same settings, we prove that AdaGrad has a *linear* rate for each iteration $k \geq 0$. [15] proved a $\mathcal{O}(\log K)$ regret bound of the SC-AdaGrad algorithm, with $y_{k+1,i}$ replacing $\sqrt{y_{k+1,i}}$ in the denominator of AdaGrad (3b) and a coordinate-wise ϵ_k , when f is strongly convex and the gradients and the estimates of minima are uniformly bounded. The SAdaGrad algorithm, a double-loop algorithm with AdaGrad in its inner loop, has a convergence rate of $\mathcal{O}(1/K)$ for weakly strongly convex f [4]. For two-layer networks with a positive definite kernel matrix, [21] proves *linear* convergence of the AdaLoss algorithm if the width of the hidden layer is sufficiently large. Compared to the aforementioned works on the analysis of AdaGrad and its variants, we prove *linear* convergence of the original AdaGrad (3a)-(3b) when f is L -smooth and satisfies the PL inequality.

Next, we review the Adam algorithm [11] and its existing convergence results. At each iteration $k = 0, 1, \dots$, Adam maintains an

estimate x_k of the solution of (1) and two moment estimates μ_k and ν_k that determine the learning rate. These estimates are updated as

$$\mu_{k+1,i} = \beta_{1k} \mu_{k,i} + (1 - \beta_{1k}) \nabla_i f(x_k), \quad (4a)$$

$$\nu_{k+1,i} = \beta_{2k} \nu_{k,i} + (1 - \beta_{2k}) |\nabla_i f(x_k)|^2, \quad (4b)$$

$$x_{k+1,i} = x_{k,i} - h \frac{\mu_{k+1,i}}{\sqrt{\nu_{k+1,i}} + \epsilon}. \quad (4c)$$

Here, $\beta_{1k}, \beta_{2k} \in [0, 1)$ are two algorithm parameters. For non-convex optimization in stochastic settings, [23, 8] have proved $\mathcal{O}(1/\sqrt{K})$ convergence rate of Adam and a family of its variants when the stochastic gradients are uniformly bounded. [9] proved $\mathcal{O}(1/K)$ convergence rate of double-loop algorithms with Adam-style update in its inner loop in the stochastic settings under PL inequality and bounded stochastic gradients. Without the bounded gradients assumption, [24] proved $\mathcal{O}(\log(K)/\sqrt{K})$ convergence rate of Adam for non-convex optimization in stochastic settings. [19] proved a $\mathcal{O}(\log K)$ regret bound of the SAdam algorithm, with $\nu_{k+1,i}$ replacing $\sqrt{\nu_{k+1,i}}$ in the denominator of Adam (4c) and a vanishing ϵ_k , when f is strongly convex and the gradients and the estimates of minima are uniformly bounded. [1] proved exponential convergence of a continuous-time version of Adam under the PL inequality. However, [1] assumes that the continuous-time version of the momentum parameters β_1 and β_2 converge to one as the stepsize h converge to zero. Also, the exponential convergence guarantee in [1] does not extend to discrete-time, as discussed by the authors. From the above literature review, we note that a *linear* convergence of Adam in discrete-time has not been proved. We prove *linear* convergence of the discrete-time Adam algorithm (4a)-(4c) in the deterministic settings when f is L -smooth and satisfies the PL inequality.

We aim to present a unified proof sketch for AdaGrad and Adam in this paper, which can potentially be utilized for variants of Adam in the future. Thus, our convergence analyses of AdaGrad and Adam both follow a similar approach. It is well known that the difficulty in obtaining a convergence rate, or even showing asymptotic convergence, of the adaptive gradient methods stems from both numerator and denominator of $x_{k+1} - x_k$ being dependent on the gradient (and its history) and, thus, the first order component in (2) does not admit a straightforward descent direction, unlike the vanilla gradient-descent. The existing analyses of adaptive methods usually tackle this challenge by adding and subtracting a surrogate term, leading to an additional complicated term in the error.

Our approach involves splitting the denominator in $x_{k+1,i} - x_{k,i}$ into two cases, depending on whether the denominator is always less than one or crosses one after a finite iterations. In the former case, the analysis from (2) becomes equivalent to that of the vanilla gradient-descent or its momentum-variant, without any adaptive stepsize. In the latter case, the challenge is that the denominator of $x_{k+1,i} - x_{k,i}$ is not apriori bounded and, hence, the argument used in gradient-descent does not trivially apply (ref. (6) later). To address this challenge in this case, our analysis involves two broader steps. In the first step, the condition on the denominator allows us to prove decrement in the cost function over the iterations, simply for AdaGrad and after a more involved argument for Adam due to its momentum in the numerator of $x_{k+1,i} - x_{k,i}$. In the next step, we use the telescopic sum, followed by a simple argument to prove the boundedness of the denominator. Once we have a bounded denominator, the analysis again becomes similar to the former case. Finally, in the general case, where the numerator of $x_{k+1,i} - x_{k,i}$ crosses one only along a subset of the dimensions $\{1, \dots, d\}$ and stays less than one along the rest of the dimensions, the linear rate is proved by the weakest

(maximum) of the linear rates obtained in the previous two cases.

To not make our exposition too burdensome, we analyze AdaGrad and Adam first in the deterministic setting, followed by analysis of AdaGrad with stochastic gradients. Due to space limit, we do not present the analysis of Adam with stochastic gradients, which can be done by following the technique presented in this paper.

The key contributions of our paper are as follows.

- We prove *linear* convergence of the discrete-time AdaGrad (3a)-(3b) and the Adam algorithm (4a)-(4c) in the deterministic setting when f is L -smooth and satisfies the PL inequality. To the best of our knowledge, the *linear* convergence guarantees of these two methods do not exist in the literature. We do not require convexity or boundedness of the gradients. Our analysis of *linear* convergence is existential in the sense that we prove *linear* convergence of the optimizers, but do not characterize the coefficient of convergence explicitly in terms of the problem and algorithm parameters.
- Our analyses of AdaGrad and Adam follow a similar approach, with the difference lying in the additional arguments for Adam due to its momentum term of the gradients. We split the convergence analysis into two broad cases depending on the magnitude of the denominator in $x_{k+1,i} - x_{k,i}$ and leverage the condition for each such case to present a simple convergence analysis. Thus, our proofs of AdaGrad and Adam are simpler than the existing works. Moreover, the presented proof methodology, as outlined earlier, provides us with a unified recipe for analyzing other adaptive gradient methods. The novelty of our technique lies in addressing the denominator of $x_{k+1,i} - x_{k,i}$, which is not bounded a priori, in a simpler way, by proving its boundedness.
- Considering the practical scenario, we prove *linear* convergence in *expectation* of AdaGrad with stochastic gradients to a neighborhood of the minima. The analysis with stochastic gradients follow the same outline described above in the deterministic setting.

Although we prove existence of a *linear* convergence rate for AdaGrad and Adam, a limitation of our analysis is that it does not characterize the exact coefficient by explicitly relating it with the algorithm parameters and the problem constants. However, at least for AdaGrad, such explicit relation can be obtained by a closer look into our analysis, which is difficult for Adam due to the additional momentum term. Due to limited space, we do not present the analysis of Adam with stochastic gradients, as it can be done by following the steps in the proofs of Theorem 2 and Theorem 3.

2 Linear convergence of AdaGrad

Theorem 1. Consider the AdaGrad algorithm in (3a)-(3b) with initialization $y_0 = 0_d$ and $x_0 \in \mathbb{R}^d$ and the parameter $\epsilon \in (0, 1)$. If Assumptions 1-3 hold, then there exists $\bar{h} > 0$ such that if the step size $0 < h < \bar{h}$, then $(f(x_{k+1}) - f_*) \leq \rho(f(x_k) - f_*)$ for all $k \geq 0$, where $\rho \in (0, 1)$.

Proof. Consider an arbitrary iteration $k \geq 0$. Under Assumption 2, from (2) we have

$$f(x_{k+1}) - f(x_k) \leq (x_{k+1} - x_k)^\top \nabla f(x_k) + \frac{L}{2} \|x_{k+1} - x_k\|^2.$$

Upon substituting above from (3b),

$$\begin{aligned} & f(x_{k+1}) - f(x_k) \\ & \leq \sum_{i=1}^d \left(-h \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k+1,i}} + \epsilon} + \frac{Lh^2}{2} \frac{|\nabla_i f(x_k)|^2}{|\sqrt{y_{k+1,i}} + \epsilon|^2} \right). \end{aligned} \quad (5)$$

Case-I: First, we consider the case when $y_{k+1,i} > (1 - \epsilon)^2$ for all $i \in \{1, \dots, d\}$ for all $k \geq T$, where $T < \infty$. Recall that $\{y_{k,i}\}$ is a non-decreasing sequence for all i . Consider any iteration $k \geq T$. Then, $|\sqrt{y_{k+1,i}} + \epsilon|^2 > \sqrt{y_{k+1,i}} + \epsilon$. From (5), then we have

$$f(x_{k+1}) - f(x_k) \leq -h \left(1 - \frac{Lh}{2}\right) \sum_{i=1}^d \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k+1,i}} + \epsilon}. \quad (6)$$

If $0 < h < \frac{2}{L}$, from above we have $f(x_{k+1}) \leq f(x_k)$. Since $\{f(x_k) : k \geq T\}$ is a decreasing sequence and, under Assumption 1, bounded from below by f_* , the sequence converges with $\lim_{k \rightarrow \infty} f(x_k) < \infty$. Next, upon summation from $t = T$ to $t = k$ on both sides of (6), due to telescopic cancellation on the L.H.S.,

$$f(x_{k+1}) - f(x_T) \leq -h \left(1 - \frac{Lh}{2}\right) \sum_{t=T}^k \sum_{i=1}^d \frac{|\nabla_i f(x_t)|^2}{\sqrt{y_{t+1,i}} + \epsilon},$$

which means

$$\begin{aligned} 0 & \leq \lim_{k \rightarrow \infty} h \left(1 - \frac{Lh}{2}\right) \sum_{t=T}^k \sum_{i=1}^d \frac{|\nabla_i f(x_t)|^2}{\sqrt{y_{t+1,i}} + \epsilon} \\ & \leq \lim_{k \rightarrow \infty} (f(x_T) - f(x_{k+1})) \leq f(x_T) - f_*. \end{aligned}$$

So, $\lim_{k \rightarrow \infty} \sum_{t=T}^k \frac{|\nabla_i f(x_t)|^2}{\sqrt{y_{t+1,i}} + \epsilon}$ is bounded, which implies that

$$\lim_{k \rightarrow \infty} \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k+1,i}} + \epsilon} = 0, \quad \forall i. \quad (7)$$

Thus, either $\lim_{k \rightarrow \infty} |\nabla_i f(x_k)| = 0$ or $\lim_{k \rightarrow \infty} \sqrt{y_{k+1,i}} = \infty$. Consider the case $\lim_{k \rightarrow \infty} \sqrt{y_{k+1,i}} = \infty$. From (7), we have $|\nabla_i f(x_k)|^2 = o(\sqrt{y_{k+1,i}} + \epsilon)$, which implies $\frac{|\nabla_i f(x_k)|}{\sqrt{y_{k+1,i}} + \epsilon} = o((\sqrt{y_{k+1,i}} + \epsilon)^{-0.5})$. So, if $\lim_{k \rightarrow \infty} \sqrt{y_{k+1,i}} = \infty$, we have $\lim_{k \rightarrow \infty} \frac{|\nabla_i f(x_k)|}{\sqrt{y_{k+1,i}} + \epsilon} = 0$. Since both $\lim_{k \rightarrow \infty} \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k+1,i}} + \epsilon} = 0$ and $\lim_{k \rightarrow \infty} \frac{|\nabla_i f(x_k)|}{\sqrt{y_{k+1,i}} + \epsilon} = 0$, it is possible only if $\lim_{k \rightarrow \infty} |\nabla_i f(x_k)| = 0$. So, we have proved that $\lim_{k \rightarrow \infty} |\nabla_i f(x_k)| = 0$ is true. Under Assumption 3, then we have $\lim_{k \rightarrow \infty} f(x_k) = f_*$. The above argument further shows that $\lim_{k \rightarrow \infty} \sqrt{y_{k+1,i}} = \infty$ is possible only in the trivial case where the function is at the minimum point. In the non-trivial case, therefore, $\sqrt{y_{k+1,i}} + \epsilon$ is bounded. Then, for $0 < h < \frac{2}{L}$, $\exists M \in (0, \infty)$ such that $\sqrt{y_{k+1,i}} + \epsilon \leq M \forall i$. From (6),

$$f(x_{k+1}) - f(x_k) \leq -\frac{h}{M} \left(1 - \frac{Lh}{2}\right) \|\nabla f(x_k)\|^2.$$

Under Assumption 3, from above we have

$$f(x_{k+1}) - f(x_k) \leq -\frac{h}{M} \left(1 - \frac{Lh}{2}\right) 2l(f(x_k) - f_*).$$

Upon defining $c_1 = h \left(1 - \frac{Lh}{2}\right) \frac{2l}{M}$, we rewrite the above as

$$f(x_{k+1}) - f(x_k) \leq -c_1(f(x_k) - f_*),$$

which means

$$f(x_{k+1}) - f_* \leq (1 - c_1)(f(x_k) - f_*), \quad \forall k \geq T. \quad (8)$$

Moreover, $0 < h < \frac{2}{L}$ implies that $c_1 > 0$ and $0 < h < \frac{M}{2l}$ implies that $c_1 < 1$. So, $(1 - c_1) \in (0, 1)$ for $0 < h < \min\{\frac{2}{L}, \frac{M}{2l}\}$.

Case-II: Next, we consider the case when $y_{k+1,i} \leq (1 - \epsilon)^2$ for all $i \in \{1, \dots, d\}$ for all $k \geq 0$. Then, $\sqrt{y_{k+1,i}} + \epsilon \leq 1$. Also, $y_0 = 0_d$ and (3a) implies that $y_{k,i} \geq 0 \forall k, i$. From (5), then we have

$$f(x_{k+1}) - f(x_k) \leq -h \left(1 - \frac{Lh}{2\epsilon^2}\right) \|\nabla f(x_k)\|^2.$$

Under Assumption 3, from above we have

$$f(x_{k+1}) - f(x_k) \leq -h \left(1 - \frac{Lh}{2\epsilon^2}\right) 2l(f(x_k) - f_*).$$

Upon defining $c_2 = h \left(1 - \frac{Lh}{2\epsilon^2}\right) 2l$, we obtain that

$$f(x_{k+1}) - f_* \leq (1 - c_2)(f(x_k) - f_*), \forall k \geq 0. \quad (9)$$

Moreover, $(1 - c_2) \in (0, 1)$ for $0 < h < \min\left\{\frac{2\epsilon^2}{L}, \frac{1}{2l}\right\}$.

Case-III: Finally, we consider the case when, for each $i \in \{1, \dots, d\}$, $\exists T_i \in [0, \infty)$ such that $y_{k+1,i} \leq (1 - \epsilon)^2$ for $k < T_i$ and $y_{k+1,i} > (1 - \epsilon)^2$ for $k \geq T_i$. Then, $\exists T = \max\{T_1, \dots, T_d\}$ such that $y_{k+1,i} > (1 - \epsilon)^2$ for all $i \in \{1, \dots, d\}$ for all $k \geq T$. For $k \geq T$, the analysis in Case-I directly applies. For $k < T$, from the analysis in Case-I and Case-II, it follows that

$$f(x_{k+1}) - f_* \leq \max\{(1 - c_1), (1 - c_2)\}(f(x_k) - f_*), \quad (10)$$

where $\max\{(1 - c_1), (1 - c_2)\}$ for $0 < h < \min\left\{\frac{2}{L}, \frac{M}{2l}, \frac{2\epsilon^2}{L}, \frac{1}{2l}\right\}$. Since $\epsilon < 1$, $\frac{2\epsilon^2}{L} < \frac{2}{L}$.

We conclude that, for $0 < h < \min\left\{\frac{2\epsilon^2}{L}, \frac{M}{2l}, \frac{1}{2l}\right\}$,

$$(f(x_{k+1}) - f_*) \leq \max\{(1 - c_1), (1 - c_2)\}(f(x_k) - f_*), \forall k \geq 0,$$

where $\max\{(1 - c_1), (1 - c_2)\} \in (0, 1)$. The proof is complete with $\rho = \max\{(1 - c_1), (1 - c_2)\}$ and $\bar{h} = \min\left\{\frac{2\epsilon^2}{L}, \frac{M}{2l}, \frac{1}{2l}\right\}$. \square

Since $\rho \in (0, 1)$, Theorem 1 implicates that the sub-optimality gap $(f(x_k) - f_*)$ of the AdaGrad algorithm (3a)-(3b) linearly converges to zero at the worst-case rate ρ , for small enough stepsize h .

3 Linear convergence of Adam

Theorem 2. Consider the Adam algorithm in (4a)-(4c) with initialization $\mu_0 = \nu_0 = 0_d$ and $x_0 \in \mathbb{R}^d$ and the parameters $\beta_{1k}, \beta_2 \in [0, 1)$, $\epsilon \in (0, 1)$. If Assumptions 1-3 hold, then $\exists \beta_k \in (0, 1)$ such that for $\beta_{1k} \in [0, \beta_k)$ there exists $\bar{h} > 0$ such that if the stepsize $0 < h < \bar{h}$, then $(f(x_{k+1}) - f_*) \leq \rho(f(x_k) - f_*)$ for all $k \geq 0$ where $\rho \in (0, 1)$.

Proof. Under Assumption 2, from (2) we have

$$f(x_{k+1}) - f(x_k) \leq (x_{k+1} - x_k)^\top \nabla f(x_k) + \frac{L}{2} \|x_{k+1} - x_k\|^2.$$

Upon substituting above from (4c),

$$\begin{aligned} & f(x_{k+1}) - f(x_k) \\ & \leq \sum_{i=1}^d \left(-h \frac{\mu_{k+1,i} \nabla_i f(x_k)}{\sqrt{\nu_{k+1,i}} + \epsilon} + \frac{Lh^2}{2} \frac{|\mu_{k+1,i}|^2}{|\sqrt{\nu_{k+1,i}} + \epsilon|^2} \right). \quad (11) \end{aligned}$$

Case-I: First, we consider the case $\nu_{k+1,i} > (1 - \epsilon)^2$ for all $i \in \{1, \dots, d\}$ for all $k \geq T$, where $T < \infty$. Then, $|\sqrt{\nu_{k+1,i}} + \epsilon|^2 > \sqrt{\nu_{k+1,i}} + \epsilon$. From (11), then we have

$$\begin{aligned} & f(x_{k+1}) - f(x_k) \\ & \leq -h \sum_{i=1}^d \left(\frac{\mu_{k+1,i} \nabla_i f(x_k)}{\sqrt{\nu_{k+1,i}} + \epsilon} - \frac{Lh}{2} \frac{|\mu_{k+1,i}|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \right). \quad (12) \end{aligned}$$

Upon substituting $\mu_{k+1,i}$ above from (4a) and rearranging the terms in the numerator,

$$\begin{aligned} & f(x_{k+1}) - f(x_k) \\ & \leq -h \sum_{i=1}^d \frac{(1 - \beta_{1k})(1 - h\frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2 - h\frac{L}{2} \beta_{1k}^2 |\mu_{k,i}|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ & \quad - h \sum_{i=1}^d \frac{\beta_{1k}(1 - hL(1 - \beta_{1k})) \mu_{k,i} \nabla_i f(x_k)}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ & = -h \sum_{i=1}^d \frac{(1 - \theta_k)(1 - \beta_{1k})(1 - h\frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ & \quad - h \sum_{i=1}^d \frac{\theta_k(1 - \beta_{1k})(1 - h\frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ & \quad - h \sum_{i=1}^d \frac{-h\frac{L}{2} \beta_{1k}^2 |\mu_{k,i}|^2 + \beta_{1k}(1 - hL(1 - \beta_{1k})) \mu_{k,i} \nabla_i f(x_k)}{\sqrt{\nu_{k+1,i}} + \epsilon}, \end{aligned}$$

for any $\theta_k \in (0, 1)$. We define two set of indices $I = \{i | i \in \{1, \dots, d\}, \mu_{k,i} \neq 0\}$ and its complement I' . Then, we rewrite the above inequality as

$$\begin{aligned} & f(x_{k+1}) - f(x_k) \\ & \leq -h \sum_{i \in I} \frac{(1 - \theta_k)(1 - \beta_{1k})(1 - h\frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ & \quad - h \sum_{i \in I} \frac{\theta_k(1 - \beta_{1k})(1 - h\frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ & \quad - h \sum_{i \in I} \frac{-h\frac{L}{2} \beta_{1k}^2 |\mu_{k,i}|^2 + \beta_{1k}(1 - hL(1 - \beta_{1k})) \mu_{k,i} \nabla_i f(x_k)}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ & \quad - h \sum_{i \in I'} \frac{(1 - \beta_{1k})(1 - h\frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon}. \quad (13) \end{aligned}$$

For $\beta_{1k} < 1$ and $h < \frac{2}{L(1 - \beta_{1k})}$, the last term on the R.H.S. in (13) is negative. So, for $i \in I$, we want to show that

$$\begin{aligned} & \theta_k(1 - \beta_{1k})(1 - h\frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2 \\ & \geq h\frac{L}{2} \beta_{1k}^2 |\mu_{k,i}|^2 + \beta_{1k}(1 - hL(1 - \beta_{1k})) |\mu_{k,i} \nabla_i f(x_k)|, \quad (14) \end{aligned}$$

for some $\theta_k \in (0, 1)$ and some $h > 0$, which would imply that $f(x_{k+1}) - f(x_k) \leq 0$, from (13). Consider any $i \in I$. We note that (14) is equivalent to

$$\begin{aligned} & \theta_k(1 - \beta_{1k}) |\nabla_i f(x_k)|^2 - \beta_{1k} |\mu_{k,i} \nabla_i f(x_k)| \geq h\frac{L}{2} \beta_{1k}^2 |\mu_{k,i}|^2 \\ & \quad + h\frac{L}{2} (1 - \beta_{1k}) (\theta_k(1 - \beta_{1k}) |\nabla_i f(x_k)|^2 - 2\beta_{1k} |\mu_{k,i} \nabla_i f(x_k)|), \end{aligned}$$

which is implied if

$$\begin{aligned} & \theta_k(1 - \beta_{1k}) |\nabla_i f(x_k)|^2 - \beta_{1k} |\mu_{k,i} \nabla_i f(x_k)| \geq h\frac{L}{2} \beta_{1k}^2 |\mu_{k,i}|^2 \\ & \quad + h\frac{L}{2} (1 - \beta_{1k}) (\theta_k(1 - \beta_{1k}) |\nabla_i f(x_k)|^2 - \beta_{1k} |\mu_{k,i} \nabla_i f(x_k)|). \quad (15) \end{aligned}$$

For the case $|\mu_{k,i}| = \infty$, we choose $\beta_{1k} = 0$ and $h < \frac{2}{L}$ so that (15) holds. Otherwise, for the case $|\mu_{k,i}| < \infty$, if the L.H.S. in (15) is positive, then (15) holds for

$$h \leq \frac{\rho_k}{q_k}, \quad (16)$$

where we denote

$$\begin{aligned} p_k &= \theta_k(1 - \beta_{1k}) |\nabla_i f(x_k)|^2 - \beta_{1k} |\mu_{k,i} \nabla_i f(x_k)|, \\ q_k &= \frac{L}{2}(1 - \beta_{1k}) (\theta_k(1 - \beta_{1k}) |\nabla_i f(x_k)|^2 - \beta_{1k} |\mu_{k,i} \nabla_i f(x_k)|) \\ &\quad + \frac{L}{2} \beta_{1k}^2 |\mu_{k,i}|^2. \end{aligned}$$

Now, the L.H.S. in (15) being positive is equivalent to

$$\theta_k > \frac{\beta_{1k}}{1 - \beta_{1k}} \frac{|\mu_{k,i}|}{|\nabla_i f(x_k)|}. \quad (17)$$

Since $|\mu_{k,i}| < \infty$, we choose $\beta_{1k} < \frac{1}{1 + \frac{|\mu_{k,i}|}{|\nabla_i f(x_k)|}}$ so that

$$\begin{aligned} \beta_{1k} &\in (0, 1) \text{ and } \frac{\beta_{1k}}{1 - \beta_{1k}} \frac{|\mu_{k,i}|}{|\nabla_i f(x_k)|} < 1. \text{ Such a choice of } \\ \beta_{1k} &\text{ allows us to choose } \theta_k \in \left(\frac{\beta_{1k}}{1 - \beta_{1k}} \frac{|\mu_{k,i}|}{|\nabla_i f(x_k)|}, 1 \right). \text{ So, for } \\ \beta_{1k} &< \frac{1}{1 + \frac{|\mu_{k,i}|}{|\nabla_i f(x_k)|}} \text{ and } \theta_k \in \left(\frac{\beta_{1k}}{1 - \beta_{1k}} \frac{|\mu_{k,i}|}{|\nabla_i f(x_k)|}, 1 \right), \text{ we have} \end{aligned}$$

that the L.H.S. in (15) is positive. Moreover, since $|\mu_{k,i}| < \infty$, both numerator and denominator in the R.H.S. of (16) are $\mathcal{O}(\theta_k(1 - \beta_{1k}) |\nabla_i f(x_k)|^2 - \beta_{1k} |\mu_{k,i} \nabla_i f(x_k)|)$. So, the R.H.S. of (16) is positive. So, we can choose $h > 0$ satisfying (16) so that (15) holds. We conclude that, $\exists \bar{\beta}_k, \theta_k \in (0, 1)$ such that for $\beta_{1k} \in [0, \bar{\beta}_k]$ and $\theta_k \in (\theta_k, 1)$ there exists $h_1 \in (0, \infty)$ such that for $h < h_1$, (15) holds, which implies that (14) holds for $i \in I$.

Upon substituting from (14) in (13), we have

$$\begin{aligned} &f(x_{k+1}) - f(x_k) \\ &\leq -h \sum_{i \in I} \frac{(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ &\quad - h \sum_{i \in I'} \frac{(1 - \beta_{1k})(1 - h \frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ &\leq -h \sum_{i \in I} \frac{(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ &\quad - h \sum_{i \in I'} \frac{(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2}(1 - \beta_{1k})) |\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ &= -h(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2}(1 - \beta_{1k})) \sum_{i=1}^d \frac{|\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} \\ &\leq 0. \end{aligned} \quad (18)$$

Under Assumption 1, we follow the argument in the proof of Theorem 1 after (6), and obtain

$$\lim_{k \rightarrow \infty} \frac{|\nabla_i f(x_k)|^2}{\sqrt{\nu_{k+1,i}} + \epsilon} = 0, \forall i. \quad (19)$$

Under Assumption 3, following the argument in the proof of Theorem 1 after (7), we have $\lim_{k \rightarrow \infty} \nabla f(x_k) = 0_d$, and $\lim_{k \rightarrow \infty} f(x_k) = f_*$, and $\exists M \in (0, \infty)$ such that $\sqrt{\nu_{k+1,i}} + \epsilon \leq M \forall i$. From (18), then we obtain

$$\begin{aligned} &f(x_{k+1}) - f(x_k) \\ &\leq -\frac{h}{M}(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2}(1 - \beta_{1k})) \|\nabla f(x_k)\|^2. \end{aligned}$$

Under Assumption 3 and defining $c_{1k} = h(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2}(1 - \beta_{1k})) \frac{2l}{M}$, from above we get

$$f(x_{k+1}) - f_* \leq (1 - c_{1k})(f(x_k) - f_*), \forall k \geq T. \quad (20)$$

Moreover, $0 < h < \frac{2}{L(1 - \beta_{1k})}$ implies that $c_{1k} > 0$ and $0 < h < \frac{M}{2l}$ implies that $c_1 < 1$. So, $(1 - c_1) \in (0, 1)$ for $0 < h < \min\{\frac{2}{L(1 - \beta_{1k})}, \frac{M}{2l}, h_1\}$.

Case-II: Next, we consider the case when $\nu_{k+1,i} \leq (1 - \epsilon)^2$ for all $i \in \{1, \dots, d\}$ for all $k \geq 0$. Then, $\sqrt{\nu_{k+1,i}} + \epsilon \leq 1$. Also, $\nu_0 = 0_d$ and (4b) implies that $\nu_{k,i} \geq 0 \forall k, i$. From (11), then we have

$$f(x_{k+1}) - f(x_k) \leq \sum_{i=1}^d \left(-h \mu_{k+1,i} \nabla_i f(x_k) + \frac{Lh^2}{2\epsilon^2} |\mu_{k+1,i}|^2 \right).$$

We note that the above inequality is similar to (12) where $\sqrt{\nu_{k+1,i}} + \epsilon$ in (12) is replaced by 1 and L is replaced by $\frac{L}{\epsilon^2}$. So, we follow the proof above in Case-I, and instead of (18), we obtain that

$$\begin{aligned} &f(x_{k+1}) - f(x_k) \\ &\leq -h(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2\epsilon^2}(1 - \beta_{1k})) \sum_{i=1}^d |\nabla_i f(x_k)|^2 \\ &= -h(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2\epsilon^2}(1 - \beta_{1k})) \|\nabla f(x_k)\|^2. \end{aligned} \quad (21)$$

Under Assumption 3 and defining $c_{2k} = h(1 - \theta_k)(1 - \beta_{1k})(1 - h \frac{L}{2\epsilon^2}(1 - \beta_{1k})) 2l$, from above we get

$$f(x_{k+1}) - f_* \leq (1 - c_{2k})(f(x_k) - f_*), \forall k \geq 0. \quad (22)$$

The rest of the proof follows the same argument as in the proof of Theorem 1. \square

Remark 1. Bias correction: Since $\mu_0 = \nu_0 = 0$, the moment estimates μ_k and ν_k are biased towards zero at the early stages of iterations. Thus, in practice, Adam is implemented with a bias correction that accounts for this initialization of the moment estimates at zero [11]. When bias correction is active, h in (4c) is essentially replaced by $h \frac{\sqrt{1 - \beta_2^{k+1}}}{1 - \beta_1^k}$. Since $\beta_{1k}, \beta_2 \in [0, 1)$, we have

$\lim_{k \rightarrow \infty} \frac{1 - \beta_1^{k+1}}{\sqrt{1 - \beta_2^{k+1}}} = 1$ and $0 < \frac{1 - \beta_1^{k+1}}{\sqrt{1 - \beta_2^{k+1}}} < \infty$. Thus, in the case of initial bias correction, the result presented in Theorem 2 is valid, upon replacing \bar{h} with the positive quantity $\bar{h} \min_{k \geq 0} \frac{1 - \beta_1^{k+1}}{\sqrt{1 - \beta_2^{k+1}}}$.

Remark 2. Like our proof above, some of the existing analyses of Adam require a decreasing β_{1k} , including [3, 11, 17]. While there are convergence results of Adam with constant β_1 , they do not prove linear convergence.

Remark 3. We do not utilize the explicit update (3a) or (4b) of the denominator y or ν . Our proofs hold as long as $y_{k,i} > 0$ or $\nu_{k,i} > 0$. The specific form of (3a) or (4b) only implicitly appears in our analysis in terms of its bound M . Therefore, Adam's analysis trivially extends to AdaBelief [25] and AMSGrad [17]. In RAdam [14], $\exists T < \infty$ such that $\rho(k)$ is increasing for $k \geq T$ and converges to ρ_∞ . Then, by definition, $0 < r(k) < \infty$. Then, following the argument in Remark 1, our analysis of Adam easily extends to RAdam for $k \geq T$ by replacing \bar{h} with $\bar{h} \min_{k \geq T} r(k)$.

4 AdaGrad with stochastic gradients

In practice, neural networks are trained with stochastic or mini-batch gradients at each iteration of the optimizer. So, in this section, we present our analysis of AdaGrad in the stochastic setting. For simplicity, we consider only one data point in each iteration. However, our result is applicable to mini-batch with more than one data points. For each data point \mathcal{D} , we define the loss function $l : \mathbb{R}^d \rightarrow \mathbb{R}$ as $l(\cdot, \mathcal{D})$ and its gradient $g(x; \mathcal{D}) = \nabla_x l(x; \mathcal{D})$. Our aim is to minimize the empirical risk,

$$\min_{x \in \mathbb{R}^d} f(x) := \mathbb{E}_\zeta [l(x; \mathcal{D}_\zeta)]. \quad (23)$$

At each iteration $k \geq 0$, AdaGrad randomly chooses a data point \mathcal{D}_{ζ_k} , based on the realization ζ_k of the random variable ζ , and computes its stochastic gradient $g_{\zeta_k}(k) = g(x_k; \mathcal{D}_{\zeta_k})$. We let $g_{i, \zeta_k}(k)$ denote the i -th element of $g_{\zeta_k}(k)$, for each $i \in \{1, \dots, d\}$. Instead of (3a)-(3b), x_k and y_k are updated as

$$x_{k+1, i} = x_{k, i} - h \frac{g_{i, \zeta_k}(k)}{\sqrt{y_{k, i}} + \epsilon}, \quad (24a)$$

$$y_{k+1, i} = y_{k, i} + |g_{i, \zeta_k}(k)|^2. \quad (24b)$$

For each iteration $k \geq 0$ we define the following.

- Let $\mathbb{E}_{\zeta_k}[\cdot]$ denote the conditional expectation of a function the random variables ζ_k , given the current x_k and y_k .
- Let $\mathbb{E}_k[\cdot]$ denote the total expectation of a function of the random variables $\{\zeta_0, \dots, \zeta_k\}$ given the initial condition x_0 and y_0 . Specifically, $\mathbb{E}_k[\cdot] = \mathbb{E}_{\zeta_0, \dots, \zeta_k}[\cdot]$, $k \geq 0$.
- For each $i \in \{1, \dots, d\}$, define the conditional variance of $g_{i, \zeta_k}(k)$, which is a function of the random variable ζ_k , given the current x_k and y_k as $\mathbb{V}_{\zeta_k}[g_{i, \zeta_k}(k)] = \mathbb{E}_{\zeta_k}[|g_{i, \zeta_k}(k) - \mathbb{E}_{\zeta_k}[g_{i, \zeta_k}(k)]|^2] = \mathbb{E}_{\zeta_k}[|g_{i, \zeta_k}(k)|^2] - |\mathbb{E}_{\zeta_k}[g_{i, \zeta_k}(k)]|^2$.

We make two additional standard assumptions [2, 18] for stochastic gradients as follows. Assumption 5 is regarding boundedness of coordinate-wise affine noise variance [18].

Assumption 4. At each iteration $k \geq 0$, the stochastic gradient is an unbiased estimate of true gradient, i.e., $\mathbb{E}_{\zeta_k}[g_{\zeta_k}(k)] = \nabla f(x_k)$.

Assumption 5. For each $i \in \{1, \dots, d\}$, there exist two non-negative real scalar values $V_{1, i}$ and $V_{2, i}$ such that, for each $k \geq 0$,

$$\mathbb{V}_{\zeta_k}[g_{i, \zeta_k}(k)] \leq V_{1, i} + V_{2, i} |\nabla_i f(x_k)|^2.$$

We define $M = \max_i V_{1, i}$ and $M_G = \max_i (V_{2, i} + 1)$.

Theorem 3. Consider the AdaGrad algorithm in (24a)-(24b) with initialization $y_0 = 0_d$ and the parameter $\epsilon \in (0, 1)$. If Assumptions 1-5 hold, then there exists $\omega \in (0, \infty)$ such that the following statements are true.

- (i) $\exists \bar{h} > 0$ such that if the step size $0 < h < \bar{h}$, then $\mathbb{E}_{\zeta_k}[f(x_{k+1})] - f(x_k) \leq \rho(f(x_k) - f_*) + \omega$, for all $k \geq 0$, where $\rho \in (0, 1)$ and $\omega = \mathcal{O}(M)$.
- (ii) $\lim_{k \rightarrow \infty} \|\nabla f(x_k)\|^2 > d \frac{LMh}{2 - LM_G h}$.
- (iii) Given arbitrary choices of the initial $x_0 \in \mathbb{R}^d$,

$$\lim_{k \rightarrow \infty} \mathbb{E}_k[f(x_{k+1})] \leq \frac{\omega}{1 - \rho}.$$

Proof. Consider an arbitrary iteration $k \geq 0$. Under Assumption 2, from (2) we have

$$f(x_{k+1}) - f(x_k) \leq (x_{k+1} - x_k)^\top \nabla f(x_k) + \frac{L}{2} \|x_{k+1} - x_k\|^2.$$

Upon taking conditional expectation on both sides and substituting above from (24a),

$$\begin{aligned} & \mathbb{E}_{\zeta_k}[f(x_{k+1})] - f(x_k) \\ & \leq \sum_{i=1}^d \left(\mathbb{E}_{\zeta_k} \left[\frac{-h g_{i, \zeta_k}(k)}{\sqrt{y_{k, i}} + \epsilon} \right] \nabla_i f(x_k) + \frac{Lh^2}{2} \mathbb{E}_{\zeta_k} \left[\frac{|g_{i, \zeta_k}(k)|^2}{|\sqrt{y_{k, i}} + \epsilon|^2} \right] \right) \\ & = \sum_{i=1}^d \left(-h \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k, i}} + \epsilon} + \frac{Lh^2}{2} \frac{\mathbb{E}_{\zeta_k}[|g_{i, \zeta_k}(k)|^2]}{|\sqrt{y_{k, i}} + \epsilon|^2} \right), \end{aligned} \quad (25)$$

where the last equality follows from Assumption 4. From the definition of conditional variance of $g_{i, \zeta_k}(k)$ and Assumptions 4-5,

$$\begin{aligned} \mathbb{E}_{\zeta_k}[|g_{i, \zeta_k}(k)|^2] & \leq V_{1, i} + (V_{2, i} + 1) |\nabla_i f(x_k)|^2 \\ & \leq M + M_G |\nabla_i f(x_k)|^2. \end{aligned}$$

Upon substituting from above in (25),

$$\begin{aligned} & \mathbb{E}_{\zeta_k}[f(x_{k+1})] - f(x_k) \\ & \leq \sum_{i=1}^d \left(-h \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k, i}} + \epsilon} + \frac{Lh^2}{2} \frac{M + M_G |\nabla_i f(x_k)|^2}{|\sqrt{y_{k, i}} + \epsilon|^2} \right). \end{aligned} \quad (26)$$

Case-I: First, we consider the case when $y_{k, i} > (1 - \epsilon)^2$ for all $i \in \{1, \dots, d\}$ for all $k \geq T$, where $T < \infty$. Consider any iteration $k \geq T$. Then, $|\sqrt{y_{k, i}} + \epsilon| > \sqrt{y_{k, i}} + \epsilon$. From (26), then we have

$$\begin{aligned} & \mathbb{E}_{\zeta_k}[f(x_{k+1})] - f(x_k) \\ & \leq -h \sum_{i=1}^d \frac{(1 - \frac{LM_G}{2} h) |\nabla_i f(x_k)|^2 - \frac{LM}{2} h}{\sqrt{y_{k, i}} + \epsilon}. \end{aligned} \quad (27)$$

For $h < \frac{2}{LM_G}$, we have $(1 - \frac{LM_G}{2} h) > 0$. Now, except at the trivial point $\nabla f(x_k) = 0_d$, there exists at least one $j \in \{1, \dots, d\}$ for which $\nabla_j f(x_k) \neq 0$. We consider the non-empty set $I = \{i | i \in \{1, \dots, d\}, \nabla_i f(x_k) \neq 0\}$ and its complement I' . Then, we can rewrite the R.H.S. above as

$$\begin{aligned} & \sum_{i=1}^d \frac{(1 - \frac{LM_G}{2} h) |\nabla_i f(x_k)|^2 - \frac{LM}{2} h}{\sqrt{y_{k, i}} + \epsilon} \\ & = \sum_{i \in I} \frac{(1 - \frac{LM_G}{2} h) |\nabla_i f(x_k)|^2 - \frac{LM}{2} h}{\sqrt{y_{k, i}} + \epsilon} + \sum_{j \in I'} \frac{-\frac{LM}{2} h}{\sqrt{y_{k, j}} + \epsilon}. \end{aligned}$$

Upon rearranging the terms above, we have

$$\begin{aligned} & \sum_{i=1}^d \frac{(1 - \frac{LM_G}{2} h) |\nabla_i f(x_k)|^2 - \frac{LM}{2} h}{\sqrt{y_{k, i}} + \epsilon} > 0 \\ & \iff h < \frac{2}{L} \frac{\sum_{i \in I} \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k, i}} + \epsilon}}{\sum_{i \in I} \frac{M_G |\nabla_i f(x_k)|^2 + M}{\sqrt{y_{k, i}} + \epsilon} + \sum_{j \in I'} \frac{M}{\sqrt{y_{k, j}} + \epsilon}}. \end{aligned} \quad (28)$$

Since $\sum_{j \in I'} \frac{M}{\sqrt{y_{k, j}} + \epsilon} < \infty$, both the numerator and denominator in the R.H.S. of (28) are $\mathcal{O}\left(\sum_{i \in I} \frac{|\nabla_i f(x_k)|^2}{\sqrt{y_{k, i}} + \epsilon}\right)$. So, the R.H.S. of (28) is positive. Then, $h > 0$ can be chosen to satisfy (28), for which

$$\sum_{i=1}^d \frac{(1 - \frac{LM_G}{2} h) |\nabla_i f(x_k)|^2 - \frac{LM}{2} h}{\sqrt{y_{k, i}} + \epsilon} > 0. \quad (29)$$

Then, (27) implies that $\mathbb{E}_{\zeta_k} [f(x_{k+1})] < f(x_k)$. Since $\{\mathbb{E}_{\zeta_k} [f(x_{k+1})] : k \geq T\}$ is a decreasing sequence and, under Assumption 1, bounded from below by f_* , the sequence converges with $\lim_{k \rightarrow \infty} \mathbb{E}_{\zeta_k} [f(x_{k+1})] < \infty$. Next, upon summation from $t = T$ to $t = k$ on both sides of (27), due to telescopic cancellation, we obtain the total expectation

$$\begin{aligned} & \mathbb{E}_k [f(x_{k+1})] - f(x_T) \\ & \leq -h \sum_{t=T}^k \sum_{i=1}^d \frac{(1 - \frac{LM_G}{2}h) |\nabla_i f(x_t)|^2 - \frac{LM}{2}h}{\sqrt{y_{t,i}} + \epsilon}. \end{aligned}$$

Combining with (29), we have

$$\begin{aligned} 0 & < \lim_{k \rightarrow \infty} \sum_{t=T}^k \sum_{i=1}^d \frac{(1 - \frac{LM_G}{2}h) |\nabla_i f(x_t)|^2 - \frac{LM}{2}h}{\sqrt{y_{t,i}} + \epsilon} \\ & \leq \lim_{k \rightarrow \infty} (f(x_T) - \mathbb{E}_k [f(x_{k+1})]) \leq f(x_T) - f_*. \end{aligned}$$

Following the argument as in the proof of Theorem 1, the above implies that

$$\lim_{k \rightarrow \infty} \frac{(1 - \frac{LM_G}{2}h) |\nabla_i f(x_k)|^2 - \frac{LM}{2}h}{\sqrt{y_{k,i}} + \epsilon} = 0, \forall i. \quad (30)$$

So, either $\lim_{k \rightarrow \infty} (1 - \frac{LM_G}{2}h) |\nabla_i f(x_k)|^2 - \frac{LM}{2}h = 0$ or $\lim_{k \rightarrow \infty} \sqrt{y_{k,i}} = \infty$. Consider the case $\lim_{k \rightarrow \infty} \sqrt{y_{k,i}} = \infty$. From (30), we have $(1 - \frac{LM_G}{2}h) |\nabla_i f(x_k)|^2 - \frac{LM}{2}h = o(\sqrt{y_{k,i}} + \epsilon)$, which implies $\frac{\sqrt{(1 - \frac{LM_G}{2}h) |\nabla_i f(x_k)|^2 - \frac{LM}{2}h}}{\sqrt{y_{k,i}} + \epsilon} = o((\sqrt{y_{k,i}} + \epsilon)^{-0.5})$. So, if $\lim_{k \rightarrow \infty} \sqrt{y_{k,i}} = \infty$, we have $\lim_{k \rightarrow \infty} \frac{\sqrt{(1 - \frac{LM_G}{2}h) |\nabla_i f(x_k)|^2 - \frac{LM}{2}h}}{\sqrt{y_{k,i}} + \epsilon} = 0$. Following the argument after (7) in the proof of Theorem 1, $\lim_{k \rightarrow \infty} \sqrt{y_{k,i}} = \infty$ implies that $\lim_{k \rightarrow \infty} (1 - \frac{LM_G}{2}h) |\nabla_i f(x_k)|^2 - \frac{LM}{2}h = 0$. So, in either of the cases above, we have

$$\lim_{k \rightarrow \infty} |\nabla_i f(x_k)|^2 = \frac{\frac{LM}{2}h}{(1 - \frac{LM_G}{2}h)} =: g_*, \forall i. \quad (31)$$

From the above argument, it further follows that $\lim_{k \rightarrow \infty} \sqrt{y_{k,i}} = \infty$ is possible only if $\lim_{k \rightarrow \infty} \|\nabla f(x_k)\|^2 = dg_*$. For $\lim_{k \rightarrow \infty} \|\nabla f(x_k)\|^2 > dg_*$, then $\sqrt{y_{k,i}} + \epsilon$ is bounded. Then, for $0 < h < \frac{2}{LM_G}$, h satisfying (29), and for $\lim_{k \rightarrow \infty} \|\nabla f(x_k)\|^2 > dg_*$, $\exists M \in (0, \infty)$ such that $\sqrt{y_{k,i}} + \epsilon \leq B \forall i$. Then, from (27),

$$\begin{aligned} & \mathbb{E}_{\zeta_k} [f(x_{k+1})] - f(x_k) \\ & \leq -\frac{h}{B} \left(1 - \frac{LM_G}{2}h\right) \|\nabla f(x_k)\|^2 + \frac{LMd}{2B} h^2. \end{aligned}$$

Under Assumption 3, from above we have

$$\begin{aligned} & \mathbb{E}_{\zeta_k} [f(x_{k+1})] - f(x_k) \\ & \leq -\frac{h}{B} \left(1 - \frac{LM_G}{2}h\right) 2l(f(x_k) - f_*) + \frac{LMd}{2B} h^2. \end{aligned}$$

Upon defining $c_1 = \frac{h}{B} \left(1 - \frac{LM_G}{2}h\right) 2l$ and $\omega = \frac{LMd}{2B} h^2$, we rewrite the above as

$$\mathbb{E}_{\zeta_k} [f(x_{k+1})] - f_* \leq (1 - c_1)(f(x_k) - f_*) + \omega, \forall k \geq T. \quad (32)$$

Moreover, $0 < h < \frac{2}{LM_G}$ implies that $c_1 > 0$ and $0 < h < \frac{B}{2l}$ implies that $c_1 < 1$. So, $(1 - c_1) \in (0, 1)$ for $0 < h < \min\{\frac{2}{LM_G}, \frac{B}{2l}\}$.

Case-II: Next, we consider the case when $y_{k,i} \leq (1 - \epsilon)^2$ for all $i \in \{1, \dots, d\}$ for all $k \geq 0$. Then, $\sqrt{y_{k,i}} + \epsilon \leq 1$. Also, $y_0 = 0_d$ and (24b) implies that $y_{k,i} \geq 0 \forall k, i$. From (26), then we have

$$\begin{aligned} & \mathbb{E}_{\zeta_k} [f(x_{k+1})] - f(x_k) \\ & = -h \left(1 - \frac{LM_G h}{2\epsilon^2}\right) \|\nabla f(x_k)\|^2 + \frac{LMdh^2}{2\epsilon^2}. \end{aligned}$$

Under Assumption 3, from above we have

$$\begin{aligned} & \mathbb{E}_{\zeta_k} [f(x_{k+1})] - f(x_k) \\ & \leq -h \left(1 - \frac{LM_G h}{2\epsilon^2}\right) 2l(f(x_k) - f_*) + \frac{LMdh^2}{2\epsilon^2}. \end{aligned}$$

Upon defining $c_2 = h \left(1 - \frac{LM_G h}{2\epsilon^2}\right) 2l$, we obtain that

$$\mathbb{E}_{\zeta_k} [f(x_{k+1})] - f_* \leq (1 - c_2)(f(x_k) - f_*) + \omega. \quad (33)$$

Moreover, $(1 - c_2) \in (0, 1)$ for $0 < h < \min\{\frac{2\epsilon^2}{LM_G}, \frac{1}{2l}\}$.

Following the same argument as in the proof of Theorem 1, for $0 < h < \min\{\frac{2\epsilon^2}{LM_G}, \frac{M}{2l}, \frac{1}{2l}\}$, h satisfying (29), and for $\lim_{k \rightarrow \infty} \|\nabla f(x_k)\|^2 > dg_*$,

$$\mathbb{E}_{\zeta_k} [f(x_{k+1})] - f(x_k) \leq \rho(f(x_k) - f_*) + \omega, \forall k \geq 0,$$

where $\rho = \max\{(1 - c_1), (1 - c_2)\}$. Since $\rho \in (0, 1)$, upon iterating the above from k to 0, by the law of total expectation, $\lim_{k \rightarrow \infty} \mathbb{E}_k [f(x_{k+1})] \leq \frac{\omega}{1 - \rho}$. The proof is complete. \square

According to Theorem 3, AdaGrad in (24a)-(24b) converges linearly in expectation to a neighborhood of minima f_* , for small enough stepsize h . The neighborhood of f_* to which $\mathbb{E}_k [f(x_{k+1})]$ converges is $\mathcal{O}(M)$, i.e., proportional to the variance of stochastic gradients evaluated at the minimum point. Furthermore, the gradient-norm $\|\nabla f(x_k)\|^2$ converges to a limit greater than a value of $\mathcal{O}(M)$.

Remark 4. Following the steps in Theorem 2 and Theorem 3, linear convergence in expectation to a neighborhood of minima f_* can be proved for Adam with stochastic gradients, i.e., $\lim_{k \rightarrow \infty} \mathbb{E}_k [f(x_{k+1})] \leq \frac{\omega}{1 - \rho}$ and $\lim_{k \rightarrow \infty} \|\nabla f(x_k)\|^2 > 0$. It implies that Adam with stochastic gradients not necessarily converges to the minima. This implication of our analysis is consistent with the result in [17] that Adam has non-zero regret at $k \rightarrow \infty$ for some online optimization problems.

5 Summary

We presented a framework that proves linear convergence of two adaptive gradient methods, namely AdaGrad and Adam in discrete-time, for minimizing smooth objective functions that satisfy the PL inequality. Among the prior works on adaptive gradient methods, only the AdaGrad-Norm algorithm and a continuous-time version of Adam have provable linear convergence, for a class of optimization problems. Thus, our work contributes towards reducing the theoretical gap between vanilla gradient-descent and the more successful adaptive gradient optimizers. The unifying approach in our framework could be applicable in rigorously analyzing other adaptive gradient optimizers.

References

- [1] A. Barakat and P. Bianchi. Convergence and dynamical behavior of the ADAM algorithm for nonconvex stochastic optimization. *SIAM Journal on Optimization*, 31(1):244–274, 2021.
- [2] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT 2010: 19th International Conference on Computational Statistics, Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers*, pages 177–186. Springer, 2010.
- [3] X. Chen, S. Liu, R. Sun, and M. Hong. On the convergence of a class of Adam-type algorithms for non-convex optimization. In *International Conference on Learning Representations*, 2019.
- [4] Z. Chen, Y. Xu, E. Chen, and T. Yang. SadaGrad: Strongly adaptive stochastic gradient methods. In *International Conference on Machine Learning*, pages 913–921. PMLR, 2018.
- [5] A. Défossez, L. Bottou, F. Bach, and N. Usunier. A simple convergence proof of Adam and Adagrad. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856.
- [6] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. In *23rd Conference on Learning Theory, COLT 2010*, pages 257–269, 2010.
- [7] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 2011.
- [8] Z. Guo, Y. Xu, W. Yin, R. Jin, and T. Yang. A novel convergence analysis for algorithms of the Adam family. *arXiv preprint arXiv:2112.03459*, 2021.
- [9] Z. Guo, Y. Xu, W. Yin, R. Jin, and T. Yang. A novel convergence analysis for algorithms of the Adam family and beyond. *arXiv preprint arXiv:2104.14840*, 2022.
- [10] H. Karimi, J. Nutini, and M. Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Lojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*, pages 795–811. Springer, 2016.
- [11] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [12] X. Li and F. 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.
- [13] C. Liu, L. Zhu, and M. Belkin. Toward a theory of optimization for over-parameterized systems of non-linear equations: the lessons of deep learning. *arXiv preprint arXiv:2003.00307*, 7, 2020.
- [14] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*, 2020.
- [15] M. C. Mukkamala and M. Hein. Variants of RMSProp and Adagrad with logarithmic regret bounds. In *International Conference on Machine Learning*, pages 2545–2553. PMLR, 2017.
- [16] I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *Mathematical Programming*, 175:69–107, 2019.
- [17] S. J. Reddi, S. Kale, and S. Kumar. On the convergence of Adam and beyond. In *International Conference on Learning Representations*, 2018.
- [18] B. Wang, H. Zhang, Z. Ma, and W. Chen. Convergence of AdaGrad for non-convex objectives: Simple proofs and relaxed assumptions. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 161–190. PMLR, 2023.
- [19] G. Wang, S. Lu, Q. Cheng, W.-w. Tu, and L. Zhang. SAdam: A variant of Adam for strongly convex functions. In *International Conference on Learning Representations*, 2019.
- [20] R. Ward, X. Wu, and L. Bottou. Adagrad stepsizes: Sharp convergence over nonconvex landscapes. *The Journal of Machine Learning Research*, 21(1):9047–9076, 2020.
- [21] X. Wu, Y. Xie, S. S. Du, and R. Ward. Adaloss: A computationally-efficient and provably convergent adaptive gradient method. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8):8691–8699, 2022.
- [22] Y. Xie, X. Wu, and R. Ward. Linear convergence of adaptive stochastic gradient descent. In *International Conference on Artificial Intelligence and Statistics*, pages 1475–1485. PMLR, 2020.
- [23] M. Zaheer, S. Reddi, D. Sachan, S. Kale, and S. Kumar. Adaptive methods for nonconvex optimization. *Advances in Neural Information Processing Systems*, 31, 2018.
- [24] Y. Zhang, C. Chen, N. Shi, R. Sun, and Z.-Q. Luo. Adam can converge without any modification on update rules. *Advances in Neural Information Processing Systems*, 35:28386–28399, 2022.
- [25] J. Zhuang, T. Tang, Y. Ding, S. C. Tatikonda, N. Dvornek, X. Papademetris, and J. Duncan. Adabelief optimizer: Adapting stepsizes by the belief in observed gradients. *Advances in neural information processing systems*, 33:18795–18806, 2020.