CRITICAL BATCH SIZE MINIMIZES STOCHASTIC FIRST-ORDER ORACLE COMPLEXITY OF DEEP LEARNING OPTIMIZER USING HYPERPARAMETERS CLOSE TO ONE

Anonymous authors
Paper under double-blind review

ABSTRACT

Practical results have shown that deep learning optimizers using small constant learning rates, hyperparameters close to one, and large batch sizes can find the model parameters of deep neural networks that minimize the loss functions. We first show theoretical evidence that the momentum method (Momentum) and adaptive moment estimation (Adam) perform well in the sense that the upper bound of the theoretical performance measure is small with a small constant learning rate, hyperparameters close to one, and a large batch size. Next, we show that there exists a batch size called the critical batch size minimizing the stochastic first-order oracle (SFO) complexity, which is the stochastic gradient computation cost, and that SFO complexity increases once the batch size exceeds the critical batch size. Finally, we provide numerical results that support our theoretical results. That is, the numerical results indicate that Adam using a small constant learning rate, hyperparameters close to one, and the critical batch size minimizing SFO complexity has faster convergence than Momentum and stochastic gradient descent (SGD).

1 INTRODUCTION

1.1 BACKGROUND

Useful deep learning optimizers have been proposed to find the model parameters of the deep neural networks that minimize loss functions called the expected risk and empirical risk, such as stochastic gradient descent (SGD) (Robbins & Monro, 1951; Zinkevich, 2003; Nemirovski et al., 2009; Ghadimi & Lan, 2012; 2013), momentum methods (Polyak, 1964; Nesterov, 1983), and adaptive methods. The various adaptive methods include Adaptive Gradient (AdaGrad) (Duchi et al., 2011), Root Mean Square Propagation (RMSProp) (Tieleman & Hinton, 2012), Adaptive Moment Estimation (Adam) (Kingma & Ba, 2015), Adaptive Mean Square Gradient (AMSGrad) (Reddi et al., 2018), Yogi (Zaheer et al., 2018), Adam with decoupled weight decay (AdamW) (Loshchilov & Hutter, 2019), and AdaBelief (named for adapting stepsizes by the belief in observed gradients) (Zhuang et al., 2020).

Theoretical analyses of adaptive methods for nonconvex optimization were presented in (Zaheer et al., 2018; Zou et al., 2019; Chen et al., 2019; Zhou et al., 2020; Zhuang et al., 2020; Chen et al., 2021; Iiduka, 2021) (see (Fehrman et al., 2020; Chen et al., 2020; Scaman & Malherbe, 2020; Loizou et al., 2021) for convergence analyses of SGD for nonconvex optimization). A particularly interesting feature of adaptive methods is the use of hyperparameters, denoted by β_1 and β_2, that can be set to influence the method performance. Almost all of the previous results for convergence analyses of adaptive methods showed that using β_1 and β_2 close to 0 makes the upper bound of the performance measure $(1/K)\sum_{k=1}^{K}\mathbb{E}[\|\nabla f(\boldsymbol{\theta}_k)\|^2]$ small, where ∇f is the gradient of a loss function $f\colon \mathbb{R}^d \to \mathbb{R}$, $(\boldsymbol{\theta}_k)_{k=1}^{K}$ is the sequence generated by an optimizer, and K is the number of steps.

Meanwhile, practical results for adaptive methods were presented in (Kingma & Ba, 2015; Reddi et al., 2018; Zaheer et al., 2018; Zou et al., 2019; Chen et al., 2019; Zhuang et al., 2020; Chen et al., 2021). These studies have shown that using, for example, $\beta_1 \in \{0.9, 0.99\}$ and $\beta_2 \in \{0.99, 0.999\}$ provides superior performance for training deep neural networks. The practically useful β_1 and β_2 are each close to 1, whereas in contrast, the theoretical results (the above paragraph) show that using β_1 and β_2 close to 0 makes the upper bounds of the performance measures small.

The practical performance of a deep learning optimizer strongly depends on the batch size. In (Smith et al., 2018), it was numerically shown that using an enormous batch size leads to a reduction in the number of parameter updates and model training time. The theoretical results in (Zaheer et al., 2018) showed that using large batch sizes makes the upper bound of $\mathbb{E}[\|\nabla f(\boldsymbol{\theta}_k)\|^2]$ of an adaptive method small. Convergence analyses of SGD in (Chen et al., 2020) indicated that running SGD with a decaying learning rate and large batch size for sufficiently many steps leads to convergence to a local minimizer of a loss function. Accordingly, the practical results for large batch sizes match the theoretical ones.

In (Shallue et al., 2019; Zhang et al., 2019), it was studied how increasing the batch size affects the performances of deep learning optimizers. In both studies, it was numerically shown that increasing batch size tends to decrease the number of steps K needed for training deep neural networks, but with diminishing returns. Moreover, it was shown that momentum methods can exploit larger batches than SGD (Shallue et al., 2019), and that K-FAC and Adam can exploit larger batches than momentum methods (Zhang et al., 2019). Thus, it was shown that momentum and adaptive methods can significantly reduce the number of steps K needed for training deep neural networks (Shallue et al., 2019, Figure 4), (Zhang et al., 2019, Figure 5).

1.2 Motivation

1.2.1 Hyperparameters close to one and constant learning rate

As described in Section 1.1, the practically useful β_1 and β_2 are each close to 1, whereas in contrast, the theoretical results show that using β_1 and β_2 close to 0 makes the upper bounds of the performance measures small. Hence, there is a gap between theory ($\beta_1, \beta_2 \approx 0$) and practice ($\beta_1, \beta_2 \approx 1$) for adaptive methods. As a consequence, the first motivation of this paper is to bridge this gap.

Since using small constant learning rates (e.g., 10^{-2}, 10^{-3}, and 10^{-4}) is robust for training deep neural networks (Kingma & Ba, 2015; Reddi et al., 2018; Zaheer et al., 2018; Zou et al., 2019; Chen et al., 2019; Zhuang et al., 2020; Chen et al., 2021), we focus on using a small constant learning rate α. Here, we note that using a learning rate depending on the Lipschitz constant L of the gradient of a loss function f would be unrealistic. This is because computing the Lipschitz constant L is NP-hard (Virmaux & Scaman, 2018, Theorem 2).

1.2.2 Critical batch size

The second motivation of this paper is to clarify theoretically the relationship between the diminishing returns reported in (Shallue et al., 2019; Zhang et al., 2019) and batch size. Numerical evaluations in (Shallue et al., 2019; Zhang et al., 2019) have definitively shown that, for deep learning optimizers, the number of steps K needed to train a deep neural network halves for each doubling of the batch size b and that there is a region of diminishing returns beyond the *critical batch size* $b^\star$. This implies that there is a positive number C such that

$$Kb \approx 2^C \text{ for } b \leq b^\star \text{ and } Kb \geq 2^C \text{ for } b \geq b^\star, \tag{1}$$

where K and b are defined for $i, j \in \mathbb{N}$ by $K = 2^i$ and $b = 2^j$ (For example, Figure 1 in Section 4 shows $C \approx 20$ and $b^\star \approx 2^{11}$ for Adam used to train ResNet-20 on the CIFAR-10 dataset). We define the *stochastic first-order oracle (SFO) complexity* of a deep learning optimizer as Kb on the basis of the number of steps K needed for training the deep neural network and the batch size b used in the optimizer. Let $b^\star$ be a critical batch size such that there are diminishing returns for all batch sizes beyond $b^\star$, as asserted in (Shallue et al., 2019; Zhang et al., 2019). This fact, expressed in (1), implies that, while SFO complexity Kb initially almost does not change (i.e., K halves for each doubling of b), Kb is minimized at critical batch size $b^\star$, and there are diminishing returns once the batch size exceeds $b^\star$.

1.3 CONTRIBUTION

Our results are summarized in Table 1. Our goal is to find a local minimizer of a loss function f over $\mathbb{R}^d$, i.e., a stationary point $\boldsymbol{\theta}^\star \in \mathbb{R}^d$ satisfying $\nabla f(\boldsymbol{\theta}^\star) = \mathbf{0}$, which is equivalent to the variational inequality (VI) defined for all $\boldsymbol{\theta} \in \mathbb{R}^d$ by $\nabla f(\boldsymbol{\theta}^\star)^\top (\boldsymbol{\theta}^\star - \boldsymbol{\theta}) \leq 0$. Hence, we use an ϵ-approximation $\mathrm{VI}(K, \boldsymbol{\theta}) \leq \epsilon$ (Table 1) as the performance measure of the sequence $(\boldsymbol{\theta}_k)_{k\in\mathbb{N}}$ generated by a deep learning optimizer, where $\epsilon > 0$ is the precision (see (5) for other performance measure of $(\boldsymbol{\theta}_k)_{k\in\mathbb{N}}$).

Table 1: Relationship between batch size b and the number of steps K to achieve an ϵ-approximation of an optimizer using a constant learning rate α and hyperparameters β_1, β_2. The critical batch size $b^\star$ minimizes SFO complexity Kb (G, σ^2, M, and v_* are positive constants, $D(\boldsymbol{\theta})$ is a positive real number depending on $\boldsymbol{\theta} \in \mathbb{R}^d$, and h is monotone decreasing with respect to β_1)

Optimizer	SGD	Momentum	Adam
C_1	$\dfrac{\mathbb{E}[\lVert\boldsymbol{\theta}_1 - \boldsymbol{\theta}\rVert^2]}{2\alpha}$	$\dfrac{\mathbb{E}[\lVert\boldsymbol{\theta}_1 - \boldsymbol{\theta}\rVert^2]}{2\alpha\beta_1}$	$\dfrac{dD(\boldsymbol{\theta})\sqrt{M}}{2\alpha\beta_1\sqrt{1-\beta_2}}$
C_2	$\dfrac{\sigma^2\alpha}{2}$	$\dfrac{\sigma^2\alpha}{2\beta_1}$	$\dfrac{\sigma^2\alpha}{2\sqrt{v_*}\beta_1(1-\beta_1)}$
C_3	$\dfrac{G^2\alpha}{2}$	$\dfrac{G^2\alpha}{2\beta_1} + h(\beta_1)$	$\dfrac{G^2\alpha}{2\sqrt{v_*}\beta_1(1-\beta_1)} + h(\beta_1)$
Upper Bound of VI	$\mathrm{VI}(K,\boldsymbol{\theta}) := \dfrac{1}{K}\sum_{k=1}^{K}\mathbb{E}\left[\nabla f(\boldsymbol{\theta}_k)^\top(\boldsymbol{\theta}_k - \boldsymbol{\theta})\right] \leq \dfrac{C_1}{K} + \dfrac{C_2}{b} + C_3 = \epsilon$		
Steps K and SFO Kb	$K = \dfrac{C_1 b}{(\epsilon - C_3)b - C_2} \quad Kb = \dfrac{C_1 b^2}{(\epsilon - C_3)b - C_2}$		
Critical Batch	$b^\star = \dfrac{2C_2}{\epsilon - C_3}$		

1.3.1 ADVANTAGE OF SETTING A SMALL CONSTANT LEARNING RATE AND HYPERPARAMETERS CLOSE TO ONE

We can show that the upper bound $C_1/K + C_2/b + C_3$ of $\mathrm{VI}(K, \boldsymbol{\theta})$ becomes small when α is small, β_1 and β_2 are close to 1, and K is large. This implies that Momentum and Adam perform well when α is small and β_1 and β_2 are each set close to 1. Section 3.1 shows this result in detail.

1.3.2 CRITICAL BATCH SIZE

We can check that the upper bound $C_1/K + C_2/b + C_3$ of $\mathrm{VI}(K, \boldsymbol{\theta})$ becomes small when a batch size b is large. Motivated by the results in (Shallue et al., 2019; Zhang et al., 2019) and Section 1.2.2, we use SFO complexity Kb as the performance measure of a deep learning optimizer. We first show that the number of steps K to satisfy $\mathrm{VI}(K, \boldsymbol{\theta}) \leq \epsilon$ can be defined as in Table 1. As a function, K is convex and monotone decreasing with respect to batch size b. Next, we show that SFO complexity Kb defined as in Table 1 is convex with respect to batch size b. This result agrees with the fact of (1). Moreover, SFO complexity Kb is minimized at

$$b^\star = \frac{2C_2}{\epsilon - C_3}. \tag{2}$$

Section 3.2 shows the above results in detail. From the above discussion, we conclude that a deep learning optimizer should use the following parameters:

- a small constant learning rate α (e.g., $\alpha = 10^{-3}$) and hyperparameters β_1 and β_2 close to 1 (e.g., $\beta_1, \beta_2 \in \{0.9, 0.99, 0.999\}$)
- the critical batch size $b^\star$ defined by (2)

The accurate setting of the critical batch size $b^\star$ defined by (2) would be difficult since $b^\star$ in (2) involves unknown parameters, such as G and $D(\boldsymbol{\theta})$ (see Section 2.2.3). Hence, we would like to estimate appropriate batch sizes using the formula (2) for $b^\star$ before implementing deep learning optimizers. Section 4 will discuss estimation of appropriate batch sizes in detail.

2 NONCONVEX OPTIMIZATION AND DEEP LEARNING OPTIMIZERS

This section gives a nonconvex optimization problem in deep neural networks and optimizers for solving the problem under standard assumptions.

2.1 NONCONVEX OPTIMIZATION IN DEEP LEARNING

Let $\mathbb{R}^d$ be a d-dimensional Euclidean space with inner product $\langle \boldsymbol{x}, \boldsymbol{y} \rangle := \boldsymbol{x}^\top \boldsymbol{y}$ inducing the norm $\|\boldsymbol{x}\|$ and $\mathbb{N}$ be the set of nonnegative integers. Define $[n] := \{1, 2, \ldots, n\}$ for $n \geq 1$. Given a parameter $\boldsymbol{\theta} \in \mathbb{R}^d$ and a data point z in a data domain Z, a machine learning model provides a prediction whose quality is measured by a differentiable nonconvex loss function $\ell(\boldsymbol{\theta}; z)$. We aim to minimize the expected loss defined for all $\boldsymbol{\theta} \in \mathbb{R}^d$ by

$$f(\boldsymbol{\theta}) = \mathbb{E}_{z \sim \mathcal{D}}[\ell(\boldsymbol{\theta}; z)] = \mathbb{E}[\ell_\xi(\boldsymbol{\theta})], \tag{3}$$

where $\mathcal{D}$ is a probability distribution over Z, ξ denotes a random variable with distribution function P, and $\mathbb{E}[\cdot]$ denotes the expectation taken with respect to ξ. A particularly interesting example of (3) is the empirical average loss defined for all $\boldsymbol{\theta} \in \mathbb{R}^d$ by

$$f(\boldsymbol{\theta}; S) = \frac{1}{n} \sum_{i \in [n]} \ell(\boldsymbol{\theta}; z_i) = \frac{1}{n} \sum_{i \in [n]} \ell_i(\boldsymbol{\theta}), \tag{4}$$

where $S = (z_1, z_2, \ldots, z_n)$ denotes the training set and $\ell_i(\cdot) := \ell(\cdot; z_i)$ denotes the loss function corresponding to the i-th training data z_i.

We would like to find a stationary point $\boldsymbol{\theta}^\star \in \mathbb{R}^d$ of f defined by (3). A point $\boldsymbol{\theta}^\star \in \mathbb{R}^d$ is a stationary point of f if and only if $\boldsymbol{\theta}^\star \in \mathbb{R}^d$ satisfies the following variational inequality (VI): for all $\boldsymbol{\theta} \in \mathbb{R}^d$, $\nabla f(\boldsymbol{\theta}^\star)^\top (\boldsymbol{\theta}^\star - \boldsymbol{\theta}) \leq 0$. Hence, we use $\mathrm{VI}(K, \boldsymbol{\theta}) := (1/K) \sum_{k=1}^K \mathbb{E}[\nabla f(\boldsymbol{\theta}_k)^\top (\boldsymbol{\theta}_k - \boldsymbol{\theta})] \leq \epsilon$ as the performance measure of the sequence $(\boldsymbol{\theta}_k)_{k \in \mathbb{N}}$ generated by a deep learning optimizer, where $\epsilon > 0$ is the precision and K denotes the number of steps of the optimizer (see also Table 1). We also use the following performance measure:

$$\mathrm{M}(K) := \frac{1}{K} \sum_{k=1}^K \mathbb{E}\left[\boldsymbol{m}_k^\top (\boldsymbol{\theta}_k - \boldsymbol{\theta}^\star)\right] \leq \epsilon. \tag{5}$$

The performance measure (5) is an ϵ–approximation of the sequence $(\boldsymbol{\theta}_k)_{k \in \mathbb{N}}$ in the sense that the inner product of the vector $\boldsymbol{\theta}_k - \boldsymbol{\theta}^\star$ and the inverse direction $\boldsymbol{m}_k$ of the search direction $-\boldsymbol{m}_k$ is less than or equal to ϵ. Thanks to the previous results in (Zaheer et al., 2018; Zou et al., 2019; Chen et al., 2019; Zhou et al., 2020; Zhuang et al., 2020; Chen et al., 2021; Iiduka, 2021), it is guaranteed that Adam and its variants can find $\boldsymbol{\theta}^\star$, which implies that $0 \leq \mathrm{M}(K)$. If $\mathrm{M}(K) \leq \epsilon$, then the definition of the inner product implies that $\boldsymbol{\theta}_K$ approximates $\boldsymbol{\theta}^\star$. Hence, it is sufficient to check whether or not a deep learning optimizer satisfies (5).

2.2 DEEP LEARNING OPTIMIZERS

2.2.1 CONDITIONS

We assume that a stochastic first-order oracle (SFO) exists such that, for a given $\boldsymbol{\theta} \in \mathbb{R}^d$, it returns a stochastic gradient $\mathsf{G}_\xi(\boldsymbol{\theta})$ of the function f defined by (3), where a random variable ξ is supported on Ξ independently of $\boldsymbol{\theta}$. The following are standard conditions when considering a deep learning optimizer.

(C1) $f \colon \mathbb{R}^d \to \mathbb{R}$ defined by (3) is continuously differentiable.

(C2) Let $(\boldsymbol{\theta}_k)_{k \in \mathbb{N}} \subset \mathbb{R}^d$ be the sequence generated by a deep learning optimizer. For each iteration k,

$$\mathbb{E}_{\xi_k}[\mathsf{G}_{\xi_k}(\boldsymbol{\theta}_k)] = \nabla f(\boldsymbol{\theta}_k), \tag{6}$$

where $\xi_0, \xi_1, \ldots$ are independent samples and the random variable ξ_k is independent of $(\boldsymbol{\theta}_l)_{l=0}^k$. There exists a nonnegative constant σ^2 such that

$$\mathbb{E}_{\xi_k}\left[\|\mathsf{G}_{\xi_k}(\boldsymbol{\theta}_k) - \nabla f(\boldsymbol{\theta}_k)\|^2\right] \leq \sigma^2. \tag{7}$$

(C3) For each iteration k, the optimizer samples a batch B_k of size b independently of k and estimates the full gradient ∇f as

$$\nabla f_{B_k}(\boldsymbol{\theta}_k) := \frac{1}{b} \sum_{i \in [b]} \mathsf{G}_{\xi_{k,i}}(\boldsymbol{\theta}_k),$$

where $\xi_{k,i}$ is a random variable generated by the i-th sampling in the k-th iteration.

In the case that f is defined by (4), we have that, for each k, $B_k \subset [n]$ and $\nabla f_{B_k}(\boldsymbol{\theta}_k) = (1/b) \sum_{i \in [b]} \nabla \ell_{\xi_{k,i}}(\boldsymbol{\theta}_k)$.

2.2.2 ADAM

Algorithm 1 is the Adam optimizer (Kingma & Ba, 2015) under (C1)–(C3). The symbol $\odot$ in step 6 is defined for all $\boldsymbol{x} = (x_i)_{i=1}^d \in \mathbb{R}^d$, $\boldsymbol{x} \odot \boldsymbol{x} := (x_i^2)_{i=1}^d \in \mathbb{R}^d$, and $\mathsf{diag}(x_i)$ in step 8 is a diagonal matrix with diagonal components $x_1, x_2, \ldots, x_d$.

Algorithm 1 Adam (Kingma & Ba, 2015)

Require: $\alpha \in (0, +\infty)$, $b \in (0, +\infty)$, $\beta_1 \in (0, 1)$, $\beta_2 \in [0, 1)$

1: $k \leftarrow 0$, $\boldsymbol{\theta}_0 \in \mathbb{R}^d$, $\boldsymbol{m}_{-1} := \mathbf{0}$, $\boldsymbol{v}_{-1} := \mathbf{0}$
2: **loop**
3: $\nabla f_{B_k}(\boldsymbol{\theta}_k) := b^{-1} \sum_{i \in [b]} \mathsf{G}_{\xi_{k,i}}(\boldsymbol{\theta}_k)$
4: $\boldsymbol{m}_k := \beta_1 \boldsymbol{m}_{k-1} + (1 - \beta_1) \nabla f_{B_k}(\boldsymbol{\theta}_k)$
5: $\hat{\boldsymbol{m}}_k := (1 - \beta_1^{k+1})^{-1} \boldsymbol{m}_k$
6: $\boldsymbol{v}_k := \beta_2 \boldsymbol{v}_{k-1} + (1 - \beta_2) \nabla f_{B_k}(\boldsymbol{\theta}_k) \odot \nabla f_{B_k}(\boldsymbol{\theta}_k)$
7: $\hat{\boldsymbol{v}}_k := (1 - \beta_2^{k+1})^{-1} \boldsymbol{v}_k$
8: $\mathsf{H}_k := \mathsf{diag}(\sqrt{\hat{v}_{k,i}})$
9: $\boldsymbol{\theta}_{k+1} := \boldsymbol{\theta}_k - \alpha \mathsf{H}_k^{-1} \hat{\boldsymbol{m}}_k$
10: $k \leftarrow k + 1$
11: **end loop**

The SGD optimizer under (C1)–(C3) is Algorithm 1 when $\beta_1 = 0$ and H_k is the identity matrix. The Momentum optimizer under (C1)–(C3) is defined for all $k \in \mathbb{N}$ by $\boldsymbol{\theta}_{k+1} := \boldsymbol{\theta}_k - \alpha \boldsymbol{m}_k$.

2.2.3 ASSUMPTIONS

We assume the following conditions that were used in (Kingma & Ba, 2015, Theorem 4.1):

(A1) There exist positive numbers G and B such that, for all $k \in \mathbb{N}$, $\|\nabla f(\boldsymbol{\theta}_k)\| \leq G$ and $\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\| \leq B$.

(A2) For all $\boldsymbol{\theta} \in \mathbb{R}^d$, there exists a positive number $\mathrm{Dist}(\boldsymbol{\theta})$ such that, for all $k \in \mathbb{N}$, $\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\| \leq \mathrm{Dist}(\boldsymbol{\theta})$.

Let $(g_{k,i}^2)_{i=1}^d := \nabla f_{B_k}(\boldsymbol{\theta}_k) \odot \nabla f_{B_k}(\boldsymbol{\theta}_k)$ $(k \in \mathbb{N})$. Assumption (A1) implies that $M := \sup\{\max_{i \in [d]} g_{k,i}^2 \colon k \in \mathbb{N}\} < +\infty$. Assumption (A2) implies that $D(\boldsymbol{\theta}) := \sup\{\max_{i \in [d]} (\theta_{k,i} - \theta_i)^2 \colon k \in \mathbb{N}\} < +\infty$. We define $v_* := \inf\{\min_{i \in [d]} v_{k,i} \colon k \in \mathbb{N}\}$. Theorem 3 in (Reddi et al., 2018) shows that there exists a stochastic convex optimization problem such that Adam using $\beta_1 < \sqrt{\beta_2}$ (e.g., $\beta_1 = 0.9$ and $\beta_2 = 0.999$) does not converge to the optimal solution. If for all $k \in \mathbb{N}$ and all $i \in [d]$, $\hat{v}_{k,i}$ in Adam satisfies

$$\hat{v}_{k+1,i} \geq \hat{v}_{k,i}, \tag{8}$$

then Adam (AMSGrad) with a decaying learning rate $\alpha_k = \mathcal{O}(1/\sqrt{k})$ and β_1 and β_2 satisfying $\beta_1 < \sqrt{\beta_2}$ can solve the stochastic convex optimization problem (Reddi et al., 2018, (2), Theorem 4). We thus assume condition (8) for Adam to guarantee the convergence of Adam.

3 OUR RESULTS

This section states our theoretical results (Theorem 3.1) in Table 1 and our contribution (Sections 3.1 and 3.2) in detail. The proof of Theorem 3.1 is given in Appendix A.

Theorem 3.1. *The sequence $(\boldsymbol{\theta}_k)_{k\in\mathbb{N}}$ generated by each of SGD, Momentum, and Adam with (8) under (C1)–(C3) and (A1) and (A2) satisfies the following:*

(i) [Upper bound of $\mathrm{VI}(K,\boldsymbol{\theta})$] *For all $K \geq 1$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,*

$$\mathrm{VI}(K,\boldsymbol{\theta}) := \frac{1}{K}\sum_{k=1}^{K}\mathbb{E}\left[\nabla f(\boldsymbol{\theta}_k)^\top(\boldsymbol{\theta}_k - \boldsymbol{\theta})\right] \leq \frac{C_1}{K} + \frac{C_2}{b} + C_3,$$

where C_i $(i = 1,2,3)$ for SGD are

$$C_1 := \frac{\mathbb{E}[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2]}{2\alpha},\ C_2 := \frac{\sigma^2\alpha}{2},\ C_3 := \frac{G^2\alpha}{2},$$

C_i $(i = 1,2,3)$ for Momentum are

$$C_1 := \frac{\mathbb{E}[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2]}{2\alpha\beta_1},\ C_2 := \frac{\sigma^2\alpha}{2\beta_1},$$
$$C_3 := \frac{G^2\alpha}{2\beta_1} + \mathrm{Dist}(\boldsymbol{\theta})\left\{\frac{G(1-\beta_1)}{\beta_1} + 2\sqrt{\sigma^2+G^2}\left(\frac{1}{\beta_1} + 2(1-\beta_1)\right)\right\},$$

C_i $(i = 1,2,3)$ for Adam with (8) are

$$C_1 := \frac{dD(\boldsymbol{\theta})\sqrt{M}}{2\alpha\beta_1\sqrt{1-\beta_2}},\ C_2 := \frac{\sigma^2\alpha}{2\sqrt{v_*}\beta_1(1-\beta_1)},$$
$$C_3 := \frac{G^2\alpha}{2\sqrt{v_*}\beta_1(1-\beta_1)} + \mathrm{Dist}(\boldsymbol{\theta})\left\{\frac{G(1-\beta_1)}{\beta_1} + 2\sqrt{\sigma^2+G^2}\left(\frac{1}{\beta_1} + 2(1-\beta_1)\right)\right\},$$

and the parameters are defined as in Section 2.2.

(ii) [Steps to satisfy $\mathrm{VI}(K,\boldsymbol{\theta}) \leq \epsilon$] *The number of steps K defined by*

$$K(b) = \frac{C_1 b}{(\epsilon - C_3)b - C_2} \tag{9}$$

satisfies $\mathrm{VI}(K,\boldsymbol{\theta}) \leq \epsilon$ and the function $K(b)$ defined by (9) is convex and monotone decreasing with respect to batch size b $(> C_2/(\epsilon - C_3) > 0)$.

(iii) [Minimization of SFO complexity] *The SFO complexity defined by*

$$K(b)b = \frac{C_1 b^2}{(\epsilon - C_3)b - C_2} \tag{10}$$

is convex with respect to batch size b $(> C_2/(\epsilon - C_3) > 0)$. The batch size

$$b^\star := \frac{2C_2}{\epsilon - C_3} \tag{11}$$

attains the minimum value $K(b^\star)b^\star = 2C_1C_2/(\epsilon - C_3)^2$ of $K(b)b$.

3.1 ADVANTAGE OF SETTING A SMALL CONSTANT LEARNING RATE AND HYPERPARAMETERS CLOSE TO ONE

We first show theoretical evidence that Adam using a small constant learning rate α, β_1 and β_2 close to 1, and a large number of steps K performs well. Theorem 3.1(i) indicates that the upper bound of $\mathrm{VI}(K,\boldsymbol{\theta})$ for Adam is

$$\mathrm{VI}(K,\boldsymbol{\theta}) \leq \frac{dD(\boldsymbol{\theta})\sqrt{M}}{2\alpha\beta_1\sqrt{1-\beta_2}K} + \frac{\alpha(\sigma^2 b^{-1} + G^2)}{2\sqrt{v_*}\beta_1(1-\beta_1)K}\sum_{k=1}^{K}\sqrt{1-\beta_2^{k+1}} + h(\beta_1) \tag{12}$$

$$\leq \underbrace{\frac{dD(\boldsymbol{\theta})\sqrt{M}}{2\alpha\beta_1\sqrt{1-\beta_2}K} + \frac{\sigma^2\alpha}{2\sqrt{v_*}\beta_1(1-\beta_1)b}}_{A(K,\boldsymbol{\theta})} + \underbrace{\frac{G^2\alpha}{2\sqrt{v_*}\beta_1(1-\beta_1)} + h(\beta_1)}_{C_3}, \tag{13}$$

where $\beta_1 \in (0,1)$, $\beta_2 \in [0,1)$, and

$$h(\beta_1) := \mathrm{Dist}(\boldsymbol{\theta}) \left\{ \frac{G(1-\beta_1)}{\beta_1} + 2\sqrt{\sigma^2 + G^2} \left(\frac{1}{\beta_1} + 2(1-\beta_1) \right) \right\} \tag{14}$$

(while Theorem 3.1(i) shows (13), the strict evaluation (12) comes from (37) in Appendix A). Since the function $h(\beta_1)$ defined by (14) is monotone decreasing, using β_1 close to 1 makes $h(\beta_1)$ small. Since $1/(\beta_1(1-\beta_1))$ is monotone increasing for $\beta_1 \geq 1/2$, using β_1 close to 1 makes $1/(\beta_1(1-\beta_1))$ large. Hence, we need to set a small α to make $\alpha(\sigma^2/b+G^2)/(2\sqrt{v_*}\beta_1(1-\beta_1))$ small. The function $\sqrt{1-\beta_2^{k+1}}$ is monotone decreasing with respect to β_2, while using β_2 close to 1 makes $1/\sqrt{1-\beta_2}$ large. When β_2 close to 1 and a small learning rate α are used, we need to use a large number of steps K to make $dD(\boldsymbol{\theta})\sqrt{M}/(2\alpha\beta_1\sqrt{1-\beta_2}K)$ small. Let $\boldsymbol{\theta}^\star$ be a stationary point of f. We can also check that the upper bound of $\mathrm{M}(K)$ defined by (5) for Adam is

$$\mathrm{M}(K) \leq A(K, \boldsymbol{\theta}^\star) + \underbrace{\frac{G^2\alpha}{2\sqrt{v_*}\beta_1(1-\beta_1)} + \overbrace{\mathrm{Dist}(\boldsymbol{\theta}^\star)(1-\beta_1)\left(\frac{G}{\beta_1} + 2\sqrt{\sigma^2+G^2}\right)}^{\tilde{h}(\beta_1)}}_{\tilde{C}_3}, \tag{15}$$

where $A(K, \boldsymbol{\theta}^\star)$ is defined as in (13) (see (39)). Accordingly, Adam performs well in the sense that the upper bound of $\mathrm{M}(K)$ is small with a small constant learning rate α, β_1 and β_2 close to 1, and a large number of steps K. See (27) in Appendix A for the upper bound of $\mathrm{M}(K)$ for Momentum.

3.2 Critical batch size

Here, we can see that the upper bound of $\mathrm{VI}(K, \boldsymbol{\theta}^\star)$ for SGD is

$$\mathrm{M}(K) = \mathrm{VI}(K, \boldsymbol{\theta}^\star) \leq \frac{\mathbb{E}[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}^\star\|^2]}{2\alpha K} + \frac{\sigma^2\alpha}{2b} + \frac{G^2\alpha}{2}. \tag{16}$$

Compared with (13) and (15), it seems that, for a fixed small constant learning rate α (e.g., $\alpha = 10^{-3}$), SGD performs better than Momentum and Adam in terms of the upper bounds of $\mathrm{VI}(K, \boldsymbol{\theta})$ and $\mathrm{M}(K)$. Meanwhile, we can check that the upper bounds of $\mathrm{VI}(K, \boldsymbol{\theta}^\star)$ and $\mathrm{M}(K)$ in (13), (15), and (16) are small when batch size b is large. Accordingly, we need to check carefully how to set batch sizes to improve theoretical and practical performances of deep learning optimizers. Motivated by the results in (Shallue et al., 2019; Zhang et al., 2019) and Section 1.2.2, we use SFO complexity Kb as the performance measure of a deep learning optimizer.

Theorem 3.1(ii) indicates that the number of steps K to satisfy $\mathrm{VI}(K, \boldsymbol{\theta}) \leq \epsilon$ can be expressed as (9) (K to satisfy $\mathrm{M}(K) \leq \epsilon$ can be obtained by replacing C_3 with $\tilde{C}_3$ in (15)). The function $K(b)$ defined by (9) is convex and monotone decreasing. Hence, the form of K defined by (9) supports theoretically the relationship between K and b shown in (Shallue et al., 2019; Zhang et al., 2019) (see also Figures 1 and 3 in this paper). Theorem 3.1(iii) indicates that SFO complexity defined by (10) is convex with respect to batch size b. This result agrees with the fact of (1) (see also Figures 2 and 4 in this paper). Moreover, SFO complexity Kb is minimized at $b^\star$ defined by (11); e.g., $b_{\mathrm{S}}^\star$ for SGD, $b_{\mathrm{M}}^\star$ for Momentum, and $b_{\mathrm{A}}^\star$ for Adam are respectively

$$b_{\mathrm{S}}^\star = \frac{\sigma_{\mathrm{S}}^2\alpha}{\epsilon - C_{3,\mathrm{S}}},\ b_{\mathrm{M}}^\star = \frac{\sigma_{\mathrm{M}}^2\alpha}{\beta_1(\epsilon - C_{3,\mathrm{M}})},\ b_{\mathrm{A}}^\star = \frac{\sigma_{\mathrm{A}}^2\alpha}{\sqrt{v_*}\beta_1(1-\beta_1)(\epsilon - C_{3,\mathrm{A}})}, \tag{17}$$

where σ_{S}, σ_{M}, and σ_{A} are positive constants depending on SGD, Momentum, and Adam and $C_3 = C_{3,\mathrm{S}}, C_{3,\mathrm{M}}, C_{3,\mathrm{A}}$ are positive constants defined as in Theorem 3.1 (see also $\tilde{C}_3$ in (15)). From (17), the lower bounds b^* of $b_{\mathrm{S}}^\star$, $b_{\mathrm{M}}^\star$, and $b_{\mathrm{A}}^\star$ are respectively

$$b_{\mathrm{S}}^\star > b_{\mathrm{S}}^* := \frac{\sigma_{\mathrm{S}}^2\alpha}{\epsilon},\ b_{\mathrm{M}}^\star > b_{\mathrm{M}}^* := \frac{\sigma_{\mathrm{M}}^2\alpha}{\beta_1\epsilon},\ b_{\mathrm{A}}^\star > b_{\mathrm{A}}^* := \frac{\sigma_{\mathrm{A}}^2\alpha}{\sqrt{v_*}\beta_1(1-\beta_1)\epsilon}.$$

If $\sigma_{\mathrm{S}} = \sigma_{\mathrm{M}} = \sigma_{\mathrm{A}}$, α and β_1 are fixed (e.g., $\alpha = 10^{-3}$ and $\beta_1 = 0.9$), and $\sqrt{v_*}(1-\beta_1) < 1$ (e.g., $0.1\sqrt{v_*} < 1$ when $\beta_1 = 0.9$), then

$$b_{\mathrm{S}}^* := \frac{\sigma_{\mathrm{S}}^2\alpha}{\epsilon} < b_{\mathrm{M}}^* := \frac{\sigma_{\mathrm{M}}^2\alpha}{\beta_1\epsilon} < b_{\mathrm{A}}^* := \frac{\sigma_{\mathrm{A}}^2\alpha}{\sqrt{v_*}\beta_1(1-\beta_1)\epsilon}. \tag{18}$$

Hence, b_{A}^* is larger than b_{S}^* and b_{M}^*. Section 4 discusses estimation of appropriate batch sizes in detail.

4 NUMERICAL RESULTS

We evaluated the performances of SGD, Momentum, and Adam with different batch sizes. The metrics were the number of steps K and the SFO complexity Kb satisfying $f(\boldsymbol{\theta}_K) \leq 10^{-1}$, where $\boldsymbol{\theta}_K$ is generated for each of SGD, Momentum, and Adam using batch size b. The stopping condition was 200 epochs. The experimental environment consisted of two Intel(R) Xeon(R) Gold 6148 2.4-GHz CPUs with 20 cores each, a 16-GB NVIDIA Tesla V100 900-Gbps GPU, and the Red Hat Enterprise Linux 7.6 OS. The code was written in Python 3.8.2 using the NumPy 1.17.3 and PyTorch 1.3.0 packages. A constant learning rate $\alpha = 10^{-3}$ was commonly used. Momentum used $\beta_1 = 0.9$. Adam used $\beta_1 = 0.9$ and $\beta_2 = 0.999$ (Kingma & Ba, 2015).

Here, we use (18) and estimate appropriate batch sizes in the sense that SFO complexity is minimized. The definitions of v_* and $v_{k,i}$ (see also (32) in Appendix A) imply that, for $k \in [K]$ and all $i \in [d]$, $v_* \leq v_{k,i} \leq \max_{k\in[K]} \max_{i\in[d]} g_{k,i}^2 =: g_{k^*,i^*}^2 \leq \sum_{i=1}^d g_{k^*,i}^2 = \|\nabla f_{B_{k^*}}(\boldsymbol{\theta}_{k^*})\|^2$. Condition (C2)(6) implies that $\mathbb{E}[\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\|^2] \leq \sigma^2/b + \mathbb{E}[\|\nabla f(\boldsymbol{\theta}_k)\|^2]$ (see also (21) in Appendix A). Conditions (C2)(7) and (C3) imply that, if b is large, then σ is small. Hence, assuming that $\sigma^2/b \approx 0$ implies that $v_* \leq \|\nabla f(\boldsymbol{\theta}_{k^*})\|^2 = \|(1/n)\sum_{i\in[n]} \nabla \ell_i(\boldsymbol{\theta}_{k^*})\|^2 = (1/n^2)\|\sum_{i\in[n]} \nabla \ell_i(\boldsymbol{\theta}_{k^*})\|^2 =: (1/n^2)\|\boldsymbol{G}_{k^*}\|^2 = (1/n^2)\sum_{i\in[d]} G_{k^*,i}^2 \leq (d/n^2)\max_{i\in[d]} G_{k^*,i}^2$. Since deep learning optimizers can approximate stationary points of f, we assume that, for example, $G_{k^*,i} \approx \epsilon$. Then, (18) implies that

$$b_{\mathrm{S}}^* := \frac{\sigma_{\mathrm{S}}^2}{10^3\epsilon},\ b_{\mathrm{M}}^* := \frac{\sigma_{\mathrm{M}}^2}{9\cdot 10^2\epsilon},\ b_{\mathrm{A}}^* := \frac{\sigma_{\mathrm{A}}^2}{9\cdot 10\sqrt{v_*}\epsilon} > \frac{\sigma_{\mathrm{A}}^2 n}{9\cdot 10\sqrt{d}\epsilon^2} =: b_{\mathrm{A}}^{**}. \tag{19}$$

Conditions (C2)(7) and (C3) imply that, if b is large, then σ is small. If SGD and Momentum can exploit large batch sizes, then σ_{S} and σ_{M} are small. However, (19) implies that b_{S}^* and b_{M}^* must be small when σ_{S} and σ_{M} are small. Accordingly, SGD and Momentum would not be able to use large batch sizes. Meanwhile, from (19), we expect that Adam can exploit a large batch size $b_{\mathrm{A}}^\star > b_{\mathrm{A}}^* > b_{\mathrm{A}}^{**}$, for example, when $b_{\mathrm{A}}^{**} < 2^{11}$ (CIFAR-10; $n = 50000$, ResNet-20; $d \approx 1.1 \times 10^7$, $b_{\mathrm{A}}^{**} \approx 1600$) and $2^{10} < b_{\mathrm{A}}^{**} < 2^{11}$ (MNIST; $n = 60000$, ResNet-18; $d \approx 1.0 \times 10^7$, $b_{\mathrm{A}}^{**} \approx 2000$) (He et al., 2016; Leong et al., 2020), where $\sigma_{\mathrm{A}}^2 \approx 10^{-2}$ and $\epsilon \approx 10^{-3}$ are used.

4.1 RESNET-20 ON THE CIFAR-10 DATASET

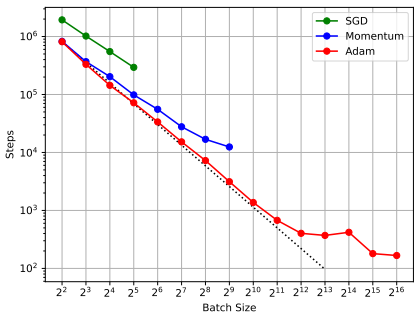

Figure 1: Number of steps for SGD, Momentum, and Adam versus batch size needed to train ResNet-20 on CIFAR-10. There is an initial period of perfect scaling (indicated by dashed line) such that the number of steps K for Adam is inversely proportional to batch size b. Adam has critical batch size $b^\star = 2^{11}$.

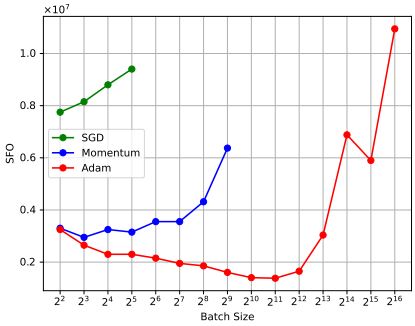

Figure 2: SFO complexities for SGD, Momentum, and Adam versus batch size needed to train ResNet-20 on CIFAR-10. SFO complexity of Adam (resp. Momentum) is minimized at critical batch size $b^\star = 2^{11}$ (resp. $b^\star = 2^3$), whereas SFO complexity for SGD tends to increase with batch size.

Let us consider training ResNet-20 on the CIFAR-10 dataset with $n = 50000$. Figure 1 shows the number of steps for SGD, Momentum, and Adam versus batch size. For SGD and Momentum, the number of steps K needed for $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ initially decreased. However, SGD with $b \geq 2^6$ and

Momentum with $b \geq 2^{10}$ did not satisfy $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ before the stopping condition was reached. Adam had an initial period of perfect scaling (indicated by dashed line) such that the number of steps K needed for $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ was inversely proportional to batch size b, and critical batch size $b^\star = 2^{11}$ such that K was not inversely proportional to the batch size beyond $b^\star$; i.e., there were diminishing returns.

Figure 2 plots the SFO complexities for SGD, Momentum, and Adam versus batch size. For SGD, SFO complexity was minimum at $b^\star = 2^2$; for Momentum, it was minimum at $b^\star = 2^3$. This implies that SGD and Momentum could not use large batch sizes, as shown in (19). For Adam, SFO complexity was minimum at the critical batch size $b^\star = 2^{11}$ that was close to estimation of critical batch $b^\star_{\mathrm{A}} > b^{**}_{\mathrm{A}}$ with $1600 \approx b^{**}_{\mathrm{A}} < b^\star_{\mathrm{A}} = 2^{11} = 2048$, as shown in (19). We also checked that the elapsed time for Adam monotonically decreased for $b \leq 2^{11}$ and that the elapsed time for critical batch size $b^\star = 2^{11}$ was the shortest. The elapsed time for $b \geq 2^{12}$ increased with the SFO complexity, as shown in Figure 2 (see Tables 2, 3, 4, and 5 in Appendix A).

4.2 ResNet-18 on the MNIST dataset

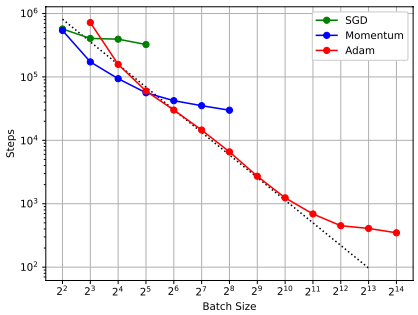

Figure 3: Number of steps for SGD, Momentum, and Adam versus batch size needed to train ResNet-18 on MNIST. There is an initial period of perfect scaling (indicated by dashed line) such that the number of steps K for Adam is inversely proportional to batch size b. Adam has critical batch size $b^\star = 2^{10}$.

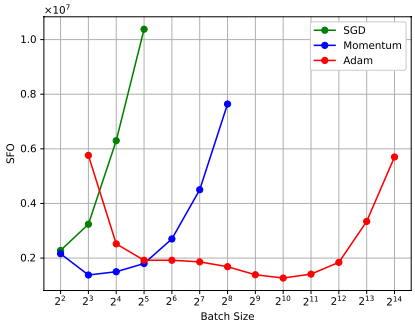

Figure 4: SFO complexities for SGD, Momentum, and Adam versus batch size needed to train ResNet-18 on MNIST. SFO complexity of Adam (resp. Momentum) is minimized at critical batch size $b^\star = 2^{10}$ (resp. $b^\star = 2^3$), whereas SFO complexity for SGD tends to increase with batch size.

Let us consider training ResNet-18 on the MNIST dataset with $n = 60000$. Figures 3 and 4 indicate that Adam could exploit larger batch sizes than SGD and Momentum. Moreover, these figures indicate that Momentum minimized SFO complexity at the critical batch size $b^\star = 2^3$ and Adam minimized SFO complexity at the critical batch size $b^\star = 2^{10}$ that was close to estimation of critical batch $b^\star_{\mathrm{A}} > b^{**}_{\mathrm{A}}$ with $2^{10} < b^{**}_{\mathrm{A}} \approx 2000 < 2^{11}$, as shown in (19). We can also check that the elapsed time for critical batch size $b^\star = 2^{10}$ was the shortest (see Tables 6, 7, 8, and 9 in Appendix A).

We also estimated appropriate batch sizes of Adam for (i) CNN on MNIST ($n = 60000$, $d \approx 7.7 \times 10^6$) and (ii) ResNet-32 on CIFAR-10 ($n = 50000$, $d \approx 2.0 \times 10^7$) from b^{**}_{A} in (19) and checked that actual critical batch sizes in (Zhang et al., 2019, Figure 5(a), (e)) are close to estimated batch sizes ((i) $b^\star_{\mathrm{A}} = 2^{11} = 2048 > b^{**}_{\mathrm{A}} \approx 2000$, (ii) $b^\star_{\mathrm{A}} = 2^{10} = 1024 \approx b^{**}_{\mathrm{A}} \approx 1200$).

5 Conclusion

This paper showed the relationship between batch size b and the number of steps K to achieve an ϵ-approximation of deep learning optimizers using a small constant learning rate α and hyperparameters β_1 and β_2 close to 1. From the convexity of SFO complexity Kb, there exists a global minimizer $b^\star$ of Kb that is the critical batch size. We also gave numerical results indicating that Adam using a small constant learning rate, hyperparameters close to one, and the critical batch size has faster convergence than Momentum and SGD. Moreover, we estimated appropriate batch sizes from our formula for $b^\star$ and showed that actual critical batch sizes are close to estimated batch sizes.

REFERENCES

Hao Chen, Lili Zheng, Raed AL Kontar, and Garvesh Raskutti. Stochastic gradient descent in correlated settings: A study on Gaussian processes. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

Jinghui Chen, Dongruo Zhou, Yiqi Tang, Ziyan Yang, Yuan Cao, and Quanquan Gu. Closing the generalization gap of adaptive gradient methods in training deep neural network. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, volume 452, pp. 3267–3275, 2021.

Xiangyi Chen, Sijia Liu, Ruoyu Sun, and Mingyi Hong. On the convergence of a class of Adam-type algorithms for non-convex optimization. In *Proceedings of The International Conference on Learning Representations*, 2019.

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.

Benjamin Fehrman, Benjamin Gess, and Arnulf Jentzen. Convergence rates for the stochastic gradient descent method for non-convex objective functions. *Journal of Machine Learning Research*, 21:1–48, 2020.

Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization I: A generic algorithmic framework. *SIAM Journal on Optimization*, 22:1469–1492, 2012.

Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization II: Shrinking procedures and optimal algorithms. *SIAM Journal on Optimization*, 23:2061–2089, 2013.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 1985.

Hideaki Iiduka. Appropriate learning rates of adaptive learning rate optimization algorithms for training deep neural networks. *IEEE Transactions on Cybernetics*, DOI: 10.1109/TCYB.2021.3107415, 2021.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of The International Conference on Learning Representations*, 2015.

Mei Chee Leong, Dilip K. Prasad, Yong Tsui Lee, and Feng Lin. Semi-CNN architecture for effective spatio-temporal learning in action recognition. *Applied Sciences*, 10(2), 2020.

Nicolas Loizou, Sharan Vaswani, Issam Laradji, and Simon Lacoste-Julien. Stochastic polyak stepsize for SGD: An adaptive learning rate for fast convergence. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130, 2021.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of The International Conference on Learning Representations*, 2019.

Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19:1574–1609, 2009.

Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. *Doklady AN USSR*, 269:543–547, 1983.

Boris T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4:1–17, 1964.

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

Herbert Robbins and Herbert Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22:400–407, 1951.

Kevin Scaman and Cédric Malherbe. Robustness analysis of non-convex stochastic gradient descent using biased expectations. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

Christopher J. Shallue, Jaehoon Lee, Joseph Antognini, Jascha Sohl-Dickstein, Roy Frostig, and George E. Dahl. Measuring the effects of data parallelism on neural network training. *Journal of Machine Learning Research*, 20:1–49, 2019.

Samuel L. Smith, Pieter-Jan Kindermans, and Quoc V. Le. Don't decay the learning rate, increase the batch size. In *Proceedings of The International Conference on Learning Representations*, 2018.

Tijmen Tieleman and Geoffrey Hinton. RMSProp: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4:26–31, 2012.

Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Manzil Zaheer, Sashank Reddi, Devendra Sachan, Satyen Kale, and Sanjiv Kumar. Adaptive methods for nonconvex optimization. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Guodong Zhang, Lala Li, Zachary Nado, James Martens, Sushant Sachdeva, George E. Dahl, Christopher J. Shallue, and Roger Grosse. Which algorithmic choices matter at which batch sizes? Insights from a noisy quadratic model. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Dongruo Zhou, Jinghui Chen, Yuan Cao, Yiqi Tang, Ziyan Yang, and Quanquan Gu. On the convergence of adaptive gradient methods for nonconvex optimization. In *12th Annual Workshop on Optimization for Machine Learning*, 2020.

Juntang Zhuang, Tommy Tang, Yifan Ding, Sekhar Tatikonda, Nicha Dvornek, Xenophon Papademetris, and James S. Duncan. AdaBelief optimizer: Adapting stepsizes by the belief in observed gradients. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, pp. 928–936, 2003.

Fangyu Zou, Li Shen, Zequn Jie, Weizhong Zhang, and Wei Liu. A sufficient condition for convergences of Adam and RMSProp. In *Computer Vision and Pattern Recognition Conference*, pp. 11127–11135, 2019.

A Appendix

Unless stated otherwise, all relationships between random variables are supposed to hold almost surely. Let S be a positive definite matrix, which is denoted by $S \in \mathbb{S}^d_{++}$. The S-inner product of $\mathbb{R}^d$ is defined for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^d$ by $\langle \boldsymbol{x}, \boldsymbol{y} \rangle_S := \langle \boldsymbol{x}, S\boldsymbol{y} \rangle = \boldsymbol{x}^\top (S\boldsymbol{y})$, and the S-norm is defined by $\|\boldsymbol{x}\|_S := \sqrt{\langle \boldsymbol{x}, S\boldsymbol{x} \rangle}$.

A.1 Lemmas

Lemma A.1. *Suppose that (C1), (C2)(6), and (C3) hold. Then, Adam satisfies the following: for all $k \in \mathbb{N}$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,*

$$\mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2_{\mathsf{H}_k}\right] = \mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2_{\mathsf{H}_k}\right] + \alpha^2 \mathbb{E}\left[\|\mathbf{d}_k\|^2_{\mathsf{H}_k}\right] + 2\alpha \left\{ \frac{\beta_1}{\tilde{\beta}_{1k}} \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \boldsymbol{m}_{k-1}\right] + \frac{\hat{\beta}_1}{\tilde{\beta}_{1k}} \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f(\boldsymbol{\theta}_k)\right] \right\},$$

where $\mathbf{d}_k := -\mathsf{H}_k^{-1} \hat{\boldsymbol{m}}_k$, $\hat{\beta}_1 := 1 - \beta_1$, and $\tilde{\beta}_{1k} := 1 - \beta_1^{k+1}$.

Proof. Let $\boldsymbol{\theta} \in \mathbb{R}^d$ and $k \in \mathbb{N}$. The definition $\boldsymbol{\theta}_{k+1} := \boldsymbol{\theta}_k + \alpha \mathbf{d}_k$ implies that

$$\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|_{\mathsf{H}_k}^2 = \|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|_{\mathsf{H}_k}^2 + 2\alpha \langle \boldsymbol{\theta}_k - \boldsymbol{\theta}, \mathbf{d}_k \rangle_{\mathsf{H}_k} + \alpha^2 \|\mathbf{d}_k\|_{\mathsf{H}_k}^2.$$

Moreover, the definitions of $\mathbf{d}_k$, $\boldsymbol{m}_k$, and $\hat{\boldsymbol{m}}_k$ ensure that

$$\begin{aligned}\langle \boldsymbol{\theta}_k - \boldsymbol{\theta}, \mathbf{d}_k \rangle_{\mathsf{H}_k} &= \langle \boldsymbol{\theta}_k - \boldsymbol{\theta}, \mathsf{H}_k \mathbf{d}_k \rangle = \langle \boldsymbol{\theta} - \boldsymbol{\theta}_k, \hat{\boldsymbol{m}}_k \rangle = \frac{1}{\tilde{\beta}_{1k}} (\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \boldsymbol{m}_k \\ &= \frac{\beta_1}{\tilde{\beta}_{1k}} (\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \boldsymbol{m}_{k-1} + \frac{\hat{\beta}_1}{\tilde{\beta}_{1k}} (\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f_{B_k}(\boldsymbol{\theta}_k).\end{aligned}$$

Hence,

$$\begin{aligned}\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|_{\mathsf{H}_k}^2 &= \|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|_{\mathsf{H}_k}^2 + \alpha^2 \|\mathbf{d}_k\|_{\mathsf{H}_k}^2 \\ &\quad + 2\alpha \left\{ \frac{\beta_1}{\tilde{\beta}_{1k}} (\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \boldsymbol{m}_{k-1} + \frac{\hat{\beta}_1}{\tilde{\beta}_{1k}} (\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f_{B_k}(\boldsymbol{\theta}_k) \right\}.\end{aligned} \tag{20}$$

Conditions (6) and (C3) guarantee that

$$\mathbb{E}\left[\mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f_{B_k}(\boldsymbol{\theta}_k) \middle| \boldsymbol{\theta}_k\right]\right] = \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \mathbb{E}\left[\nabla f_{B_k}(\boldsymbol{\theta}_k) \middle| \boldsymbol{\theta}_k\right]\right] = \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f(\boldsymbol{\theta}_k)\right].$$

Therefore, the lemma follows from taking the expectation on both sides of (20). This completes the proof. □

The discussion in the proof of Lemma A.1 also gives the following lemma.

Lemma A.2. *Suppose that (C1), (C2)(6), and (C3) hold. Then, SGD satisfies the following: for all $k \in \mathbb{N}$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,*

$$\mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2\right] = \mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2\right] + \alpha^2 \mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\|^2\right] + 2\alpha \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f(\boldsymbol{\theta}_k)\right].$$

Moreover, Momentum satisfies the following: for all $k \in \mathbb{N}$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\begin{aligned}\mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2\right] &= \mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2\right] + \alpha^2 \mathbb{E}\left[\|\boldsymbol{m}_k\|^2\right] \\ &\quad + 2\alpha \left\{ \beta_1 \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \boldsymbol{m}_{k-1}\right] + \hat{\beta}_1 \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f(\boldsymbol{\theta}_k)\right] \right\},\end{aligned}$$

where $\hat{\beta}_1 := 1 - \beta_1$.

The following lemma indicates the bounds on $(\mathbb{E}[\|\boldsymbol{m}_k\|^2])_{k \in \mathbb{N}}$ and $(\mathbb{E}[\|\mathbf{d}_k\|_{\mathsf{H}_k}^2])_{k \in \mathbb{N}}$.

Lemma A.3. *Adam under (C2)(6), (7), and (A1), for all $k \in \mathbb{N}$ satisfies*

$$\mathbb{E}\left[\|\boldsymbol{m}_k\|^2\right] \leq \frac{\sigma^2}{b} + G^2, \quad \mathbb{E}\left[\|\mathbf{d}_k\|_{\mathsf{H}_k}^2\right] \leq \frac{\sqrt{\tilde{\beta}_{2k}}}{\tilde{\beta}_{1k}^2 \sqrt{v_*}} \left(\frac{\sigma^2}{b} + G^2\right),$$

where $v_* := \inf\{\min_{i \in [d]} v_{k,i} \colon k \in \mathbb{N}\}$, $\tilde{\beta}_{1k} := 1 - \beta_{1k}^{k+1}$, and $\tilde{\beta}_{2k} := 1 - \beta_{2k}^{k+1}$.

Proof. Condition (C2)(6) implies that

$$\begin{aligned}\mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\|^2 \middle| \boldsymbol{\theta}_k\right] &= \mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k) - \nabla f(\boldsymbol{\theta}_k) + \nabla f(\boldsymbol{\theta}_k)\|^2 \middle| \boldsymbol{\theta}_k\right] \\ &= \mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k) - \nabla f(\boldsymbol{\theta}_k)\|^2 \middle| \boldsymbol{\theta}_k\right] + \mathbb{E}\left[\|\nabla f(\boldsymbol{\theta}_k)\|^2 \middle| \boldsymbol{\theta}_k\right] \\ &\quad + 2\mathbb{E}\left[(\nabla f_{B_k}(\boldsymbol{\theta}_k) - \nabla f(\boldsymbol{\theta}_k))^\top \nabla f(\boldsymbol{\theta}_k) \middle| \boldsymbol{\theta}_k\right] \\ &= \mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k) - \nabla f(\boldsymbol{\theta}_k)\|^2 \middle| \boldsymbol{\theta}_k\right] + \|\nabla f(\boldsymbol{\theta}_k)\|^2,\end{aligned} \tag{21}$$

which, together with (C2)(7) and (A1), in turn implies that

$$\mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\|^2\right] \leq \frac{\sigma^2}{b} + G^2. \tag{22}$$

The convexity of $\|\cdot\|^2$, together with the definition of $\boldsymbol{m}_k$ and (22), guarantees that, for all $k \in \mathbb{N}$,

$$\begin{aligned}\mathbb{E}\left[\|\boldsymbol{m}_k\|^2\right] &\leq \beta_1 \mathbb{E}\left[\|\boldsymbol{m}_{k-1}\|^2\right] + \hat{\beta}_1 \mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\|^2\right] \\ &\leq \beta_1 \mathbb{E}\left[\|\boldsymbol{m}_{k-1}\|^2\right] + \hat{\beta}_1 \left(\frac{\sigma^2}{b} + G^2\right).\end{aligned}$$

Induction thus ensures that, for all $k \in \mathbb{N}$,

$$\mathbb{E}\left[\|\boldsymbol{m}_k\|^2\right] \leq \max\left\{\|\boldsymbol{m}_{-1}\|^2, \frac{\sigma^2}{b} + G^2\right\} = \frac{\sigma^2}{b} + G^2, \tag{23}$$

where $\boldsymbol{m}_{-1} = \mathbf{0}$. For $k \in \mathbb{N}$, $\mathsf{H}_k \in \mathbb{S}^d_{++}$ guarantees the existence of a unique matrix $\overline{\mathsf{H}}_k \in \mathbb{S}^d_{++}$ such that $\mathsf{H}_k = \overline{\mathsf{H}}_k^2$ (Horn & Johnson, 1985, Theorem 7.2.6). We have that, for all $\boldsymbol{x} \in \mathbb{R}^d$, $\|\boldsymbol{x}\|^2_{\mathsf{H}_k} = \|\overline{\mathsf{H}}_k \boldsymbol{x}\|^2$. Accordingly, the definitions of $\mathbf{d}_k$ and $\hat{\boldsymbol{m}}_k$ imply that, for all $k \in \mathbb{N}$,

$$\mathbb{E}\left[\|\mathbf{d}_k\|^2_{\mathsf{H}_k}\right] = \mathbb{E}\left[\left\|\overline{\mathsf{H}}_k^{-1}\mathsf{H}_k \mathbf{d}_k\right\|^2\right] \leq \frac{1}{\tilde{\beta}_{1k}^2}\mathbb{E}\left[\left\|\overline{\mathsf{H}}_k^{-1}\right\|^2 \|\boldsymbol{m}_k\|^2\right],$$

where

$$\left\|\overline{\mathsf{H}}_k^{-1}\right\| = \left\|\mathsf{diag}\left(\hat{v}_{k,i}^{-\frac{1}{4}}\right)\right\| = \max_{i \in [d]} \hat{v}_{k,i}^{-\frac{1}{4}} = \max_{i \in [d]} \left(\frac{v_{k,i}}{\tilde{\beta}_{2k}}\right)^{-\frac{1}{4}} =: \left(\frac{v_{k,i^*}}{\tilde{\beta}_{2k}}\right)^{-\frac{1}{4}}.$$

Moreover, the definition of

$$v_* := \inf\{v_{k,i^*} \colon k \in \mathbb{N}\}$$

and (23) imply that, for all $k \in \mathbb{N}$,

$$\mathbb{E}\left[\|\mathbf{d}_k\|^2_{\mathsf{H}_k}\right] \leq \frac{\tilde{\beta}_{2k}^{\frac{1}{2}}}{\tilde{\beta}_{1k}^2 v_*^{\frac{1}{2}}}\left(\frac{\sigma^2}{b} + G^2\right),$$

completing the proof. □

A.2 SGD

We show Theorem 3.1 for SGD.

Proof. (i) Lemma A.2 and (22) imply that, for all $k \in \mathbb{N}$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\begin{aligned}2\alpha \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f(\boldsymbol{\theta}_k)\right] &= \mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2\right] - \mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2\right] + \alpha^2 \mathbb{E}\left[\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\|^2\right] \\ &\leq \mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2\right] - \mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2\right] + \alpha^2\left(\frac{\sigma^2}{b} + G^2\right).\end{aligned}$$

Summing the above inequality from $k = 1$ to $k = K$ leads to the finding that, for all $K \geq 1$,

$$2\alpha \sum_{k=1}^{K} \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f(\boldsymbol{\theta}_k)\right] \leq \mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2\right] - \mathbb{E}\left[\|\boldsymbol{\theta}_{K+1} - \boldsymbol{\theta}\|^2\right] + \alpha^2\left(\frac{\sigma^2}{b} + G^2\right) K,$$

which implies that, for all $K \geq 1$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\begin{aligned}\mathrm{VI}(K, \boldsymbol{\theta}) := \frac{1}{K}\sum_{k=1}^{K} \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f(\boldsymbol{\theta}_k)\right] &\leq \frac{\mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2\right]}{2\alpha K} + \frac{\alpha}{2}\left(\frac{\sigma^2}{b} + G^2\right) \\ &= \underbrace{\frac{\mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2\right]}{2\alpha}}_{C_1}\frac{1}{K} + \underbrace{\frac{\sigma^2 \alpha}{2}}_{C_2}\frac{1}{b} + \underbrace{\frac{G^2 \alpha}{2}}_{C_3}.\end{aligned} \tag{24}$$

(ii) Let $\boldsymbol{\theta} \in \mathbb{R}^d$ and $\epsilon > 0$. Condition $C_1/K + C_2/b + C_3 = \epsilon$ is equivalent to

$$K = K(b) = \frac{C_1 b}{(\epsilon - C_3)b - C_2}. \tag{25}$$

Since $\epsilon = C_1/K + C_2/b + C_3 > C_3$, we consider the case $b > C_2/(\epsilon - C_3) > 0$ to guarantee that $K > 0$. From (24), the function K defined by (25) satisfies $\mathrm{VI}(K, \boldsymbol{\theta}) \leq C_1/K + C_2/b + C_3 = \epsilon$. Moreover, from (25),

$$\frac{\mathrm{d}K(b)}{\mathrm{d}b} = \frac{-C_1 C_2}{\{(\epsilon - C_3)b - C_2\}^2} \leq 0, \ \frac{\mathrm{d}^2 K(b)}{\mathrm{d}b^2} = \frac{2C_1 C_2(\epsilon - C_3)}{\{(\epsilon - C_3)b - C_2\}^3} \geq 0,$$

which implies that K is convex and monotone decreasing with respect to b.

(iii) We have that

$$Kb = K(b)b = \frac{C_1 b^2}{(\epsilon - C_3)b - C_2}.$$

Accordingly,

$$\frac{\mathrm{d}K(b)b}{\mathrm{d}b} = \frac{C_1 b\{(\epsilon - C_3)b - 2C_2\}}{\{(\epsilon - C_3)b - C_2\}^2}, \ \frac{\mathrm{d}^2 K(b)b}{\mathrm{d}b^2} = \frac{2C_1 C_2^2}{\{(\epsilon - C_3)b - C_2\}^3} \geq 0,$$

which implies that $K(b)b$ is convex with respect to b and

$$\frac{\mathrm{d}K(b)b}{\mathrm{d}b} \begin{cases} < 0 & \text{if } b < b^\star, \\ = 0 & \text{if } b = b^\star = \frac{2C_2}{\epsilon - C_3}, \\ > 0 & \text{if } b > b^\star. \end{cases}$$

The point $b^\star$ attains the minimum value $K(b^\star)b^\star = 2C_1 C_2/(\epsilon - C_3)^2$ of $K(b)b$. This completes the proof. □

A.3 Momentum

We show Theorem 3.1 for Momentum.

Proof. (i) Lemma A.2 ensures that, for all $k \in \mathbb{N}$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\begin{aligned} \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1}\right] = \frac{1}{2\alpha\beta_1}\left\{\mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2\right] - \mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2\right]\right\} + \frac{\alpha}{2\beta_1}\mathbb{E}\left[\|\boldsymbol{m}_k\|^2\right] \\ + \frac{\hat{\beta}_1}{\beta_1}\mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f(\boldsymbol{\theta}_k)\right], \end{aligned}$$

which, together with Lemma A.3, the Cauchy–Schwarz inequality, and (A1) and (A2), implies that

$$\begin{aligned} \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1}\right] \leq \frac{1}{2\alpha\beta_1}\left\{\mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2\right] - \mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2\right]\right\} + \frac{\alpha}{2\beta_1}\left(\frac{\sigma^2}{b} + G^2\right) \\ + \frac{\hat{\beta}_1}{\beta_1}\mathrm{Dist}(\boldsymbol{\theta})G. \end{aligned}$$

Summing the above inequality from $k = 1$ to $k = K$ gives a relation that implies that

$$\begin{aligned} \sum_{k=1}^{K}\mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1}\right] \leq \frac{1}{2\alpha\beta_1}\left\{\mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2\right] - \mathbb{E}\left[\|\boldsymbol{\theta}_{K+1} - \boldsymbol{\theta}\|^2\right]\right\} + \frac{\alpha}{2\beta_1}\left(\frac{\sigma^2}{b} + G^2\right)K \\ + \frac{\hat{\beta}_1}{\beta_1}\mathrm{Dist}(\boldsymbol{\theta})GK, \end{aligned}$$

and hence,

$$\frac{1}{K}\sum_{k=1}^{K}\mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1}\right] \leq \frac{\mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2\right]}{2\alpha\beta_1 K} + \frac{\alpha}{2\beta_1}\left(\frac{\sigma^2}{b} + G^2\right) + \frac{\hat{\beta}_1}{\beta_1}\mathrm{Dist}(\boldsymbol{\theta})G.$$

Moreover, we have that, for all $k \in \mathbb{N}$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\begin{aligned}(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k &= (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1} + \hat{\beta}_1 (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top (\nabla f_{B_k}(\boldsymbol{\theta}_k) - \boldsymbol{m}_{k-1}) \\ &\leq (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1} + \hat{\beta}_1 \mathrm{Dist}(\boldsymbol{\theta}) \left(\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\| + \|\boldsymbol{m}_{k-1}\| \right),\end{aligned}$$

where the first equality comes from the definition of $\boldsymbol{m}_k$ and the first inequality comes from the Cauchy–Schwarz inequality, the triangle inequality, and (A2). Hence, from Lemma A.3, (22), Jensen's inequality, and $b \geq 1$,

$$\begin{aligned}\mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k\right] &\leq (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1} + 2\hat{\beta}_1 \mathrm{Dist}(\boldsymbol{\theta}) \sqrt{\frac{\sigma^2}{b} + G^2} \\ &\leq (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1} + 2\hat{\beta}_1 \mathrm{Dist}(\boldsymbol{\theta}) \sqrt{\sigma^2 + G^2}.\end{aligned} \tag{26}$$

Therefore, for all $K \geq 1$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\begin{aligned}&\frac{1}{K} \sum_{k=1}^{K} \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k\right] \\ &\leq \frac{\mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2\right]}{2\alpha\beta_1 K} + \frac{\sigma^2 \alpha}{2\beta_1 b} + \frac{G^2 \alpha}{2\beta_1} + \hat{\beta}_1 \mathrm{Dist}(\boldsymbol{\theta}) \left(\frac{G}{\beta_1} + 2\sqrt{\sigma^2 + G^2} \right).\end{aligned} \tag{27}$$

The definition of $\boldsymbol{m}_k$ ensures that

$$\begin{aligned}&(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f_{B_k}(\boldsymbol{\theta}_k) \\ &= (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k + (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top (\nabla f_{B_k}(\boldsymbol{\theta}_k) - \boldsymbol{m}_{k-1}) + (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top (\boldsymbol{m}_{k-1} - \boldsymbol{m}_k) \\ &= (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k + \frac{1}{\beta_1} (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top (\nabla f_{B_k}(\boldsymbol{\theta}_k) - \boldsymbol{m}_k) + \hat{\beta}_1 (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top (\boldsymbol{m}_{k-1} - \nabla f_{B_k}(\boldsymbol{\theta}_k)),\end{aligned}$$

which, together with the Cauchy–Schwarz inequality, the triangle inequality, and (A2), implies that

$$\begin{aligned}&(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f_{B_k}(\boldsymbol{\theta}_k) \\ &\leq (\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k + \frac{1}{\beta_1} \mathrm{Dist}(\boldsymbol{\theta}) (\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\| + \|\boldsymbol{m}_k\|) + \hat{\beta}_1 \mathrm{Dist}(\boldsymbol{\theta}) (\|\nabla f_{B_k}(\boldsymbol{\theta}_k)\| + \|\boldsymbol{m}_{k-1}\|).\end{aligned}$$

Lemma A.3, (22), Jensen's inequality, and $b \geq 1$ guarantee that

$$\begin{aligned}\mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f(\boldsymbol{\theta}_k)\right] &\leq \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k\right] + 2\left(\frac{1}{\beta_1} + \hat{\beta}_1\right) \mathrm{Dist}(\boldsymbol{\theta}) \sqrt{\frac{\sigma^2}{b} + G^2} \\ &\leq \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k\right] + 2\left(\frac{1}{\beta_1} + \hat{\beta}_1\right) \mathrm{Dist}(\boldsymbol{\theta}) \sqrt{\sigma^2 + G^2}.\end{aligned} \tag{28}$$

Therefore, (27) ensures that, for all $K \geq 1$ and all $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\begin{aligned}&\frac{1}{K} \sum_{k=1}^{K} \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f(\boldsymbol{\theta}_k)\right] \\ &\leq \underbrace{\frac{\mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2\right]}{2\alpha\beta_1}}_{C_1} \frac{1}{K} + \underbrace{\frac{\sigma^2 \alpha}{2\beta_1}}_{C_2} \frac{1}{b} + \underbrace{\frac{G^2 \alpha}{2\beta_1} + \mathrm{Dist}(\boldsymbol{\theta}) \left\{ \frac{G\hat{\beta}_1}{\beta_1} + 2\sqrt{\sigma^2 + G^2} \left(\frac{1}{\beta_1} + 2\hat{\beta}_1 \right) \right\}}_{C_3}.\end{aligned}$$

(ii) A discussion similar to the one showing (ii) in Theorem 3.1 for SGD would show (ii) in Theorem 3.1 for Momentum.

(iii) An argument similar to that which obtained (iii) in Theorem 3.1 for SGD would prove (iii) in Theorem 3.1 for Momentum. □

A.4 ADAM

We show Theorem 3.1 for Adam.

Proof. (i) Let $\boldsymbol{\theta} \in \mathbb{R}^d$. Lemma A.1 guarantees that for all $k \in \mathbb{N}$,

$$\mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1}\right] = \underbrace{\frac{\tilde{\beta}_{1k}}{2\alpha\beta_1}\left\{\mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2_{\mathsf{H}_k}\right] - \mathbb{E}\left[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}\|^2_{\mathsf{H}_k}\right]\right\}}_{a_k} + \underbrace{\frac{\alpha\tilde{\beta}_{1k}}{2\beta_1}\mathbb{E}\left[\|\mathbf{d}_k\|^2_{\mathsf{H}_k}\right]}_{b_k}$$
$$+ \underbrace{\frac{\hat{\beta}_1}{\beta_1}\mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f(\boldsymbol{\theta}_k)\right]}_{c_k}. \tag{29}$$

We define $\gamma_k := \tilde{\beta}_{1k}/(2\beta_1\alpha)$ $(k \in \mathbb{N})$. Then, for all $K \geq 1$,

$$\sum_{k=1}^{K} a_k = \gamma_1 \mathbb{E}\left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|^2_{\mathsf{H}_1}\right] + \underbrace{\sum_{k=2}^{K}\left\{\gamma_k \mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2_{\mathsf{H}_k}\right] - \gamma_{k-1}\mathbb{E}\left[\|\boldsymbol{\theta}_k - \boldsymbol{\theta}\|^2_{\mathsf{H}_{k-1}}\right]\right\}}_{\Gamma_K}$$
$$- \gamma_K \mathbb{E}\left[\|\boldsymbol{\theta}_{K+1} - \boldsymbol{\theta}\|^2_{\mathsf{H}_K}\right]. \tag{30}$$

Since $\overline{\mathsf{H}}_k \in \mathbb{S}^d_{++}$ exists such that $\mathsf{H}_k = \overline{\mathsf{H}}_k^2$, we have $\|\boldsymbol{x}\|^2_{\mathsf{H}_k} = \|\overline{\mathsf{H}}_k \boldsymbol{x}\|^2$ for all $\boldsymbol{x} \in \mathbb{R}^d$. Accordingly, we also have

$$\Gamma_K = \mathbb{E}\left[\sum_{k=2}^{K}\left\{\gamma_k \left\|\overline{\mathsf{H}}_k(\boldsymbol{\theta}_k - \boldsymbol{\theta})\right\|^2 - \gamma_{k-1}\left\|\overline{\mathsf{H}}_{k-1}(\boldsymbol{\theta}_k - \boldsymbol{\theta})\right\|^2\right\}\right].$$

From $\overline{\mathsf{H}}_k = \mathsf{diag}(\hat{v}_{k,i}^{1/4})$, we have that, for all $\boldsymbol{x} = (x_i)_{i=1}^d \in \mathbb{R}^d$, $\|\overline{\mathsf{H}}_k \boldsymbol{x}\|^2 = \sum_{i=1}^d \sqrt{\hat{v}_{k,i}} x_i^2$. Hence, for all $K \geq 2$,

$$\Gamma_K = \mathbb{E}\left[\sum_{k=2}^{K}\sum_{i=1}^{d}\left(\gamma_k\sqrt{\hat{v}_{k,i}} - \gamma_{k-1}\sqrt{\hat{v}_{k-1,i}}\right)(\theta_{k,i} - \theta_i)^2\right]. \tag{31}$$

Condition (8) and $\gamma_k \geq \gamma_{k-1}$ $(k \geq 1)$ imply that, for all $k \geq 1$ and all $i \in [d]$,

$$\gamma_k\sqrt{\hat{v}_{k,i}} - \gamma_{k-1}\sqrt{\hat{v}_{k-1,i}} \geq 0.$$

Moreover, (A2) ensures that $D(\boldsymbol{\theta}) := \sup\{\max_{i\in[d]}(\theta_{k,i} - \theta_i)^2 \colon k \in \mathbb{N}\} < +\infty$. Accordingly, for all $K \geq 2$,

$$\Gamma_K \leq D(\boldsymbol{\theta})\mathbb{E}\left[\sum_{k=2}^{K}\sum_{i=1}^{d}\left(\gamma_k\sqrt{\hat{v}_{k,i}} - \gamma_{k-1}\sqrt{\hat{v}_{k-1,i}}\right)\right] = D(\boldsymbol{\theta})\mathbb{E}\left[\sum_{i=1}^{d}\left(\gamma_K\sqrt{\hat{v}_{K,i}} - \gamma_1\sqrt{\hat{v}_{1,i}}\right)\right].$$

Let $\nabla f_{B_k}(\boldsymbol{\theta}_k) \odot \nabla f_{B_k}(\boldsymbol{\theta}_k) := (g_{k,i}^2) \in \mathbb{R}^d_+$. Assumption (A1) ensures that there exists $M \in \mathbb{R}$ such that, for all $k \in \mathbb{N}$, $\max_{i\in[d]} g_{k,i}^2 \leq M$. The definition of $\boldsymbol{v}_k$ guarantees that, for all $i \in [d]$ and all $k \in \mathbb{N}$,

$$v_{k,i} = \beta_2 v_{k-1,i} + \hat{\beta}_2 g_{k,i}^2.$$

Induction thus ensures that, for all $i \in [d]$ and all $k \in \mathbb{N}$,

$$v_{k,i} \leq \max\{v_{0,i}, M\} = M, \tag{32}$$

where $\boldsymbol{v}_0 = (v_{0,i}) = \mathbf{0}$. From the definition of $\hat{\boldsymbol{v}}_k$, we have that, for all $i \in [d]$ and all $k \in \mathbb{N}$,

$$\hat{v}_{k,i} = \frac{v_{k,i}}{\tilde{\beta}_{2k}} \leq \frac{M}{\tilde{\beta}_{2k}}. \tag{33}$$

Therefore, (30), $\mathbb{E}[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}\|_{\mathsf{H}_1}^2] \leq D(\boldsymbol{\theta})\mathbb{E}[\sum_{i=1}^d \sqrt{\hat{v}_{1,i}}]$, and (33) imply, for all $K \geq 1$,

$$\begin{aligned}\sum_{k=1}^K a_k &\leq \gamma_1 D(\boldsymbol{\theta})\mathbb{E}\left[\sum_{i=1}^d \sqrt{\hat{v}_{1,i}}\right] + D(\boldsymbol{\theta})\mathbb{E}\left[\sum_{i=1}^d \left(\gamma_K \sqrt{\hat{v}_{K,i}} - \gamma_1 \sqrt{\hat{v}_{1,i}}\right)\right] \\ &= \gamma_K D(\boldsymbol{\theta})\mathbb{E}\left[\sum_{i=1}^d \sqrt{\hat{v}_{K,i}}\right] \\ &\leq \gamma_K D(\boldsymbol{\theta}) \sum_{i=1}^d \sqrt{\frac{M}{\tilde{\beta}_{2K}}} \\ &\leq \frac{dD(\boldsymbol{\theta})\sqrt{M}\tilde{\beta}_{1K}}{2\beta_1\alpha\sqrt{\tilde{\beta}_{2K}}}.\end{aligned} \tag{34}$$

Inequality (34) with $\tilde{\beta}_{1K} := 1 - \beta_1^{K+1} \leq 1$ and $\tilde{\beta}_{2K} := 1 - \beta_2^{K+1} \geq 1 - \beta_2 =: \hat{\beta}_2$ implies that

$$\sum_{k=1}^K a_k \leq \frac{dD(\boldsymbol{\theta})\sqrt{M}\tilde{\beta}_{1K}}{2\beta_1\alpha\sqrt{\tilde{\beta}_{2K}}} \leq \frac{dD(\boldsymbol{\theta})\sqrt{M}}{2\beta_1\alpha\sqrt{\hat{\beta}_2}}. \tag{35}$$

Lemma A.3 guarantees that, for all $k \in \mathbb{N}$,

$$b_k = \frac{\alpha\tilde{\beta}_{1k}}{2\beta_1}\mathbb{E}\left[\|\mathbf{d}_k\|_{\mathsf{H}_k}^2\right] \leq \frac{\alpha\tilde{\beta}_{1k}}{2\beta_1}\frac{\sqrt{\tilde{\beta}_{2k}}}{\tilde{\beta}_{1k}^2\sqrt{v_*}}\left(\frac{\sigma^2}{b} + G^2\right) = \frac{\alpha\sqrt{\tilde{\beta}_{2k}}}{2\sqrt{v_*}\beta_1\tilde{\beta}_{1k}}\left(\frac{\sigma^2}{b} + G^2\right). \tag{36}$$

Inequality (36) with $\tilde{\beta}_{1k} := 1 - \beta_1^{k+1} \geq 1 - \beta_1 =: \hat{\beta}_1$ and $\tilde{\beta}_{2k} := 1 - \beta_2^{k+1} \leq 1$ implies that

$$b_k \leq \frac{\alpha\sqrt{\tilde{\beta}_{2k}}}{2\sqrt{v_*}\beta_1\tilde{\beta}_{1k}}\left(\frac{\sigma^2}{b} + G^2\right) \leq \frac{\alpha}{2\sqrt{v_*}\beta_1\hat{\beta}_1}\left(\frac{\sigma^2}{b} + G^2\right). \tag{37}$$

The Cauchy–Schwarz inequality and (A2) imply that, for all $k \in \mathbb{N}$,

$$c_k = \frac{\hat{\beta}_1}{\beta_1}\mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\theta}_k)^\top \nabla f(\boldsymbol{\theta}_k)\right] \leq \mathrm{Dist}(\boldsymbol{\theta})G\frac{\hat{\beta}_1}{\beta_1}. \tag{38}$$

Hence, (29), (35), (37), and (38) ensure that, for all $K \geq 1$,

$$\frac{1}{K}\sum_{k=1}^K \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_{k-1}\right] \leq \frac{dD(\boldsymbol{\theta})\sqrt{M}}{2\beta_1\alpha\sqrt{\hat{\beta}_2}K} + \frac{\alpha(\sigma^2 b^{-1} + G^2)}{2\sqrt{v_*}\beta_1\hat{\beta}_1} + \mathrm{Dist}(\boldsymbol{\theta})G\frac{\hat{\beta}_1}{\beta_1}.$$

Therefore, from (26), for all $K \geq 1$,

$$\begin{aligned}&\frac{1}{K}\sum_{k=1}^K \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \boldsymbol{m}_k\right] \\ &\leq \frac{dD(\boldsymbol{\theta})\sqrt{M}}{2\beta_1\alpha\sqrt{\hat{\beta}_2}K} + \frac{\sigma^2\alpha}{2\sqrt{v_*}\beta_1\hat{\beta}_1 b} + \frac{G^2\alpha}{2\sqrt{v_*}\beta_1\hat{\beta}_1} + \hat{\beta}_1\mathrm{Dist}(\boldsymbol{\theta})\left(\frac{G}{\beta_1} + 2\sqrt{\sigma^2 + G^2}\right).\end{aligned} \tag{39}$$

From (28) and (39), for all $K \geq 1$,

$$\begin{aligned}&\frac{1}{K}\sum_{k=1}^K \mathbb{E}\left[(\boldsymbol{\theta}_k - \boldsymbol{\theta})^\top \nabla f(\boldsymbol{\theta}_k)\right] \\ &\leq \underbrace{\frac{dD(\boldsymbol{\theta})\sqrt{M}}{2\alpha\beta_1\sqrt{\hat{\beta}_2}}}_{C_1}\frac{1}{K} + \underbrace{\frac{\sigma^2\alpha}{2\sqrt{v_*}\beta_1\hat{\beta}_1}}_{C_2}\frac{1}{b} + \underbrace{\frac{G^2\alpha}{2\sqrt{v_*}\beta_1\hat{\beta}_1} + \mathrm{Dist}(\boldsymbol{\theta})\left\{\frac{G\hat{\beta}_1}{\beta_1} + 2\sqrt{\sigma^2 + G^2}\left(\frac{1}{\beta_1} + 2\hat{\beta}_1\right)\right\}}_{C_3}.\end{aligned}$$

(ii) A discussion similar to the one showing (ii) in Theorem 3.1 for SGD would show (ii) in Theorem 3.1 for Adam.

(iii) An argument similar to that which obtained (iii) in Theorem 3.1 for SGD would prove (iii) in Theorem 3.1 for Adam. □

A.5 ADDITIONAL NUMERICAL RESULTS

Table 2: Elapsed time and training accuracy of SGD when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-20 on CIFAR-10

	SGD							
Batch Size	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
Time (s)	16983.64	9103.76	6176.19	3759.25	—	—	—	—
Accuracy (%)	96.75	96.69	96.66	96.88	—	—	—	—

Table 3: Elapsed time and training accuracy of Momentum when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-20 on CIFAR-10

	Momentum							
Batch Size	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
Time (s)	7978.90	3837.72	2520.82	1458.70	887.01	678.66	625.10	866.65
Accuracy (%)	96.49	96.79	96.51	96.72	96.70	96.94	96.94	98.34

Table 4: Elapsed time and training accuracy of Adam when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-20 on CIFAR-10

	Adam							
Batch Size	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
Time (s)	10601.78	4405.73	2410.28	1314.01	617.14	487.75	281.74	225.03
Accuracy (%)	96.46	96.38	96.65	96.53	96.43	96.68	96.58	96.72

Table 5: Elapsed time and training accuracy of Adam when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-20 on CIFAR-10

	Adam						
Batch Size	2^{10}	$\mathbf{2^{11}}$	2^{12}	2^{13}	2^{14}	2^{15}	2^{16}
Time (s)	197.78	195.40	233.70	349.81	691.04	644.19	1148.68
Accuracy (%)	96.74	97.21	97.54	97.75	97.51	99.05	99.03

Table 6: Elapsed time and training accuracy of SGD when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-18 on MNIST

	SGD							
Batch Size	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
Time (s)	1620.98	1002.89	1140.47	1298.52	—	—	—	—
Accuracy (%)	99.70	99.69	99.70	99.70	—	—	—	—

Table 7: Elapsed time and training accuracy of Momentum when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-18 on MNIST

	Momentum							
Batch Size	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
Time (s)	1949.66	550.28	298.72	238.66	268.98	362.89	567.54	—
Accuracy (%)	99.68	99.65	99.68	99.69	99.71	99.70	99.72	—

Table 8: Elapsed time and training accuracy of Adam when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-18 on MNIST

	Adam							
Batch Size	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
Time (s)	—	7514.43	614.27	297.97	216.10	160.93	127.97	102.98
Accuracy (%)	—	99.74	99.66	99.64	99.66	99.66	99.69	99.71

Table 9: Elapsed time and training accuracy of Adam when $f(\boldsymbol{\theta}_K) \leq 10^{-1}$ for training ResNet-18 on MNIST

	Adam				
Batch Size	$\mathbf{2^{10}}$	2^{11}	2^{12}	2^{13}	2^{14}
Time (s)	93.53	100.22	129.29	217.23	375.87
Accuracy (%)	99.68	99.67	99.67	99.70	99.69