

On the Last Iterate Convergence of Momentum Methods

Xiaoyu Li

XIAOYULI@BU.EDU

Division of System Engineering, Boston University, Boston, MA 02215, USA

Mingrui Liu

MINGRUIL@GMU.EDU

Department of Computer Science, George Mason University, Fairfax, VA 22030, USA

Francesco Orabona

FRANCESCO@ORABONA.COM

Electrical & Computer Engineering, Boston University, Boston, MA 02215, USA

Editors: Sanjoy Dasgupta and Nika Haghtalab

Abstract

SGD with Momentum (SGDM) is a widely used family of algorithms for large-scale optimization of machine learning problems. Yet, when optimizing generic convex functions, no advantage is known for any SGDM algorithm over plain SGD. Moreover, even the most recent results require changes to the SGDM algorithms, like averaging of the iterates and a projection onto a bounded domain, which are rarely used in practice. In this paper, we focus on the convergence rate of the last iterate of SGDM. For the first time, we prove that for any constant momentum factor, there exists a Lipschitz and convex function for which the last iterate of SGDM suffers from a suboptimal convergence rate of $\Omega(\frac{\ln T}{\sqrt{T}})$ after T iterations. Based on this fact, we study a class of (both adaptive and non-adaptive) Follow-The-Regularized-Leader-based SGDM algorithms with *increasing momentum* and *shrinking updates*. For these algorithms, we show that the last iterate has optimal convergence $O(\frac{1}{\sqrt{T}})$ for unconstrained convex stochastic optimization problems without projections onto bounded domains nor knowledge of T . Further, we show a variety of results for FTRL-based SGDM when used with adaptive stepsizes. Empirical results are shown as well.

Keywords: Convex Optimization, Momentum methods, Stochastic Optimization

1. Introduction

Momentum methods have become one of the most used first-order optimization algorithms in machine learning applications. When momentum is used together with Stochastic Gradient Descent (SGD), there are two main variants considered in the literature: the stochastic version of the Heavy Ball momentum (SHB) (Polyak, 1964) and Nesterov’s momentum (also called Nesterov Accelerate Gradient method) (Nesterov, 1983). Besides these two, there are other variations as well. For example, an exponential moving average of the (stochastic) gradients can be used to replace the gradients in the updates (Kingma and Ba, 2015; Reddi et al., 2016; Alacaoglu et al., 2020; Liu et al., 2020).

In this paper, we denote by Stochastic Gradient Descent with Momentum (SGDM) the general updates

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \mathbf{m}_t, \quad \mathbf{m}_t = \beta_t \mathbf{m}_{t-1} + \nu_t \mathbf{g}_t, \quad (1)$$

where $\nu_t > 0$ and $0 \leq \beta_t \leq 1$. In particular, when $\nu_t = 1$, it recovers the updates of SHB. Instead, when $\nu_t = 1 - \beta_t$, it recovers the variant with exponential moving average of the stochastic gradients.

Despite this zoo of variants, due to the presence of noise, it is well-known that SGDM *does not* guarantee an accelerated rate of convergence of noise nor any real advantage over plain SGD

on generic convex problems. For example, recent works have proved that a variant of SGD with momentum improves only the non-dominant terms in the convergence rate on some specific stochastic problems (Dieuleveut et al., 2017; Jain et al., 2018). Moreover, often an idealized version of SGDM is used in the theoretical analysis rather than the actual SGDM people use in practice. For example, projections onto bounded domains at each step, averaging of the iterates (e.g., Alacaoglu et al., 2020), and knowledge of the total number of iterations (Ghadimi and Lan, 2012) are often assumed. Overall, recent analyses seem unable to pinpoint any advantage of using a momentum term in SGD in the stochastic optimization of general convex functions.

In this paper, to show a discriminant difference between SGD and SGDM, we focus on the convergence of the last iterate. Hence, we study the convergence of the *last iterate* of SGDM for unconstrained optimization of convex functions. Unfortunately, our first result is a negative one: We show that the last iterate of SGDM can have a suboptimal convergence rate for *any constant momentum setting*.

Hence, motivated by the above result, we analyze yet another variant of SGDM. We start from the very recent observation (Defazio, 2020) that SGDM can be seen as a primal averaging procedure (Nesterov and Shikhman, 2015; Tao et al., 2018; Cutkosky, 2019) applied to the iterates of Online Mirror Descent (OMD) (Nemirovsky and Yudin, 1983; Warmuth and Jagota, 1997). Based on this fact, we analyze SGDM algorithms based on the Follow-the-Regularized-Leader (FTRL) framework¹ (Shalev-Shwartz, 2007; Abernethy et al., 2008) and the primal averaging. The use of FTRL instead of OMD removes the necessity of projections onto bounded domains, while the primal averaging acts as a momentum term and guarantees the optimal convergence of the last iterate. The resulting algorithm has an *increasing momentum* and *shrinking updates* that precisely allow to avoid our lower bound.

More in detail, we prove that the expected suboptimality gap of the last iterate of FTRL-based SGDM converges at the optimal rate of $O(1/\sqrt{T})$ on convex functions, without assuming bounded domains nor the knowledge of the total number of iterations. This also disproves a more general conjecture than the one in (Jain et al., 2019, 2021), removing the bounded assumption. Moreover, we show that our construction is general enough to allow for an entire family of FTRL-based SGDM methods, both adaptive and non-adaptive. For example, we show that “adaptive” learning rates give rise to convergence rates that are adaptive to gradients, noise, and to the interpolation regime.

The rest of the paper is organized as follows: We discuss the related work in Section 2 and the setting and assumptions in Section 3. We then present our main results: the lower bound (Section 4) and the new FTRL-based SGDM (Section 5). Finally, in Section 6 we present an empirical evaluation of our algorithms and in Section 7 we outline a future work direction.

2. Related Work

Stochastic Momentum Methods SGDM has become a popular tool in deep learning and its importance has been discussed by recent studies (Sutskever et al., 2013). Polyak (1964) first proposed the use of momentum in gradient descent, calling it the Heavy-Ball method. In the stochastic setting, there are multiple work analyzing the use of momentum in SGD. In particular, Yang et al. (2016) prove a convergence rate of $O(1/\sqrt{T})$ for the averaged iterate in the convex setting, and for an iterate taken uniformly at random in the nonconvex setting. Liu et al. (2020) provide a convergence analysis

1. FTRL is known in the offline optimization literature as Dual Averaging (DA) (Nesterov, 2009), but in reality, DA is a special case of FTRL when the functions are linearized.

for SGDM and Multistage SGDM for smooth functions in the strongly convex and nonconvex settings. Also, adaptive variants of momentum methods (Kingma and Ba, 2015; Reddi et al., 2018; Luo et al., 2018) are very popular in the deep learning literature, even if their guarantees are only for the online convex optimization setting assuming a decreasing momentum factor and projections onto bounded domains. Alacaoglu et al. (2020) recently removed the assumption of a vanishing momentum factor, but they still require projections over a bounded domain. In the non-convex and smooth case, Cutkosky and Orabona (2019) introduce a variant of SGDM with a variance-reduction effect and a faster convergence rate than SGD on non-convex functions, but it requires two stochastic gradients per step.

Lower Bound Harvey et al. (2019) prove the tight convergence bound $O(\ln T/\sqrt{T})$ of the last iterate of SGD for convex and Lipschitz functions. Kidambi et al. (2018) provide a lower bound for the Heavy Ball method for least square regression problems. To the best of our knowledge, there is no lower bound for the last iterate of SGDM in the general non-smooth non-strongly-convex setting.

Last Iterate Convergence of SGDM Nesterov and Shikhman (2015) introduces a quasi-monotone subgradient method, which uses double averaging (both in Primal and Dual) based on Dual Averaging, to achieve the optimal convergence of the last iterate for the convex and Lipschitz functions. However, they just considered the batch case. This approach was then rediscovered and extended by Cutkosky (2019). Our FTRL-based SGDM is a generalization of the approach in Nesterov and Shikhman (2015) with generic regularizers and stochastic gradients. Tao et al. (2018) extends Nesterov and Shikhman (2015)’s method to Mirror Descent, calling it stochastic primal averaging. They recover the same bound for convex functions, again with a bounded domain assumption. Defazio (2020) points out that the sequence generated by the stochastic primal averaging (Tao et al., 2018) can be identical to that of stochastic gradient descent with momentum for specific choices of the hyper-parameters. Accordingly, they give a Lyapunov analysis in the nonconvex and smooth case. Based on this work, Jelassi and Defazio (2020) introduce “Modernized dual averaging method”, which is actually equal to the one by Nesterov and Shikhman (2015). They also give a similar Lyapunov analysis as in Defazio (2020) with specific choices of hyper-parameters in the non-convex and smooth optimization setting, where they assume a bounded domain and get a convergence bound $O(\ln T/\sqrt{T})$. Recently, Tao et al. (2021) propose the very same algorithm as in Tao et al. (2018) and analyze it as a modified Polyak’s Heavy-ball method (already pointed out by Defazio (2020)). They give an analysis in the convex cases and extend it to an adaptive version, obtaining in both cases an optimal convergence of the last iterate. However, they still assume the use of projections onto bounded domains.

Last iterate convergence rate $O(\frac{1}{\sqrt{T}})$ Ghadimi and Lan (2012) present the last iterate of AC-SA (Nemirovski et al., 2009; Lan, 2012) for convex functions in the unconstrained setting, that in the Euclidean case reduces to SGD with an increasing Nesterov momentum, showing that it can achieve a convergence rate $O(\frac{1}{\sqrt{T}})$ if the number of iterations T is known in advance. Sebbouh et al. (2021) analyze Stochastic Heavy Ball-Iterave Moving Average method (SHB-IMA), which is equal to the Stochastic Heavy Ball method (SHB) with a specific choice of hyper-parameters. They prove a convergence rate for the last iterate of $O(\frac{1}{\sqrt{T}})$ if T is given in advance, and is $O(\frac{\ln T}{\sqrt{T}})$ if T is unknown. Jain et al. (2019, 2021) conjecture that under assumption (H3’) (see next Section) “for any-time algorithm (i.e., without apriori knowledge of T) expected error rate of $\frac{DG \ln T}{\sqrt{T}}$ is information-theoretically optimal”, where D is the diameter of the bounded domain. This was already disproved by the results in Tao et al. (2021), but here we disprove it even in the more challenging unconstrained setting.

Table 1: Last iterate convergence of momentum methods in convex setting

Algorithm	Assumption	Bounded Domain	Requires T	Rate	Reference
Adaptive-HB	(H3')	Yes	No	$O(\frac{1}{\sqrt{T}})$	Tao et al. (2021)
SHB-IMA	(H1) + (H2)	No	Yes	$O(\frac{1}{\sqrt{T}})$	Sebbouh et al. (2021)
			No	$O(\frac{\ln T}{\sqrt{T}})$	
AC-SA	(H2) + (H1) or (H2) + Lipschitz	No	Yes	$O(\frac{1}{\sqrt{T}})$	Ghadimi and Lan (2012)
FTRL-SGDM	(H3)	No	No	$O(\frac{1}{\sqrt{T}})$	This paper, Corollary 1
	(H1)+(H2)+(H3')	No	No	$O(\frac{\ln T}{T} + \frac{\sigma}{\sqrt{T}})$	This paper, Corollary 4

We summarize the results on the last iterate convergence for convex optimization and their assumptions in Table 1. The assumptions are defined in the next section.

3. Problem Set-up

Notation We denote vectors by bold letters, e.g. $\mathbf{x} \in \mathbb{R}^d$. All standard operations on the vectors, e.g., $\mathbf{x}\mathbf{y}$, \mathbf{x}/\mathbf{y} , $\sqrt{\mathbf{x}}$ and $\mathbf{x} < \mathbf{y}$, are to be considered element-wise. We denote by $\mathbb{E}[\cdot]$ the expectation with respect to the underlying probability space and by $\mathbb{E}_t[\cdot]$ the conditional expectation with respect to the past. Any norm without particular notation in this work is the ℓ_2 norm.

Setting We consider the unconstrained optimization problem $\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$, where $f(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex function and we denote its infimum by f^* . We also assume to have access to a first-order black-box optimization oracle that returns a stochastic subgradient in any point $\mathbf{x} \in \mathbb{R}^d$. In particular, we assume that we receive a vector $\mathbf{g}(\mathbf{x}, \xi)$ such that $\mathbb{E}_\xi[\mathbf{g}(\mathbf{x}, \xi)] = \nabla f(\mathbf{x})$ for any $\mathbf{x} \in \mathbb{R}^d$. To make the notation concise, we let $\mathbf{g}_t \triangleq \mathbf{g}(\mathbf{x}_t, \xi_t)$ and $\mathbb{E}_t[\mathbf{g}_t] = \nabla f(\mathbf{x}_t), \forall t$.

We will make different assumptions on the objective function f . Sometimes, we will assume that

- **(H1)** f is L -smooth, that is, f is continuously differentiable and its gradient is L -Lipschitz, i.e., $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|$.

We also use one or more of the following assumptions on the stochastic gradients \mathbf{g}_t .

- **(H2)** bounded variance: $\mathbb{E}_t\|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2 \leq \sigma^2$.
- **(H3)** bounded in expectation: $\mathbb{E}\|\mathbf{g}_t\|^2 \leq G^2$.
- **(H3')** ℓ_2 bounded: $\|\mathbf{g}_t\| \leq G$.
- **(H3'')** ℓ_∞ bounded: $\|\mathbf{g}_t\|_\infty \leq G_\infty$.

4. Lower bound for SGDM

First of all, as we discussed in the related work, most of the analyses of SGDM assume a vanishing momentum or a constant one. However, is constant momentum the best setting for stochastic optimization of convex functions, especially for the convergence of the last iterate? For this question, it is worth remembering that the use of a constant momentum term is mainly motivated by the empirical evidence in the deep learning literature. However, deep learning objective functions are

non-convex and the convex setting might be different. Also, the deep learning literature offers no theoretical explanations.

In this section, we show the surprising result that for SGD with any constant momentum, there exists a function for which the lower bound of the last iterate is $\Omega\left(\ln T/\sqrt{T}\right)$. Our proof extends the one in [Harvey et al. \(2019\)](#) to SGD with momentum.

We consider SGDM with constant momentum factors β and $\nu > 0$ in (1), where $\mathbf{g}_t \in \partial f(\mathbf{x}_t)$ and a polynomial stepsize $\eta_t = c \cdot t^{-\alpha}$, $0 \leq \alpha \leq \frac{1}{2}$.

Let \mathcal{X} denote the Euclidean ball with radius $\frac{2c}{1-\beta}$ in \mathbb{R}^T . For any fixed β, ν and α , we introduce the following function. Define $f: \mathcal{X} \rightarrow \mathbb{R}$ and $\mathbf{h}_i \in \mathbb{R}^T$ for $i \in [T+1]$ by

$$f(\mathbf{x}) = \max_{i \in [T+1]} \mathbf{h}_i^T \mathbf{x}, \quad h_{i,j} = \begin{cases} a_j, & 1 \leq j < i \\ -b_j, & i = j < T \\ 0, & i < j \leq T \end{cases} \quad (2)$$

where $b_j = \frac{2j^\alpha}{\nu T^\alpha}$ and $a_j = \frac{1-\beta}{2\nu(T-j+1)}$. We have that $\partial f(\mathbf{x}_t)$ is the convex hull of $\mathbf{h}_i : i \in \mathcal{I}(\mathbf{x})$ where $\mathcal{I}(\mathbf{x}) = \{i : \mathbf{h}_i^T \mathbf{x} = f(\mathbf{x})\}$.

Note that f is Lipschitz over \mathcal{X} since

$$\|\mathbf{h}_i\|^2 \leq \sum_{i=1}^T a_i^2 + \frac{4}{\nu^2} \leq \frac{1}{4\nu^2} \sum_{i=1}^T \frac{1}{i^2} + \frac{4}{\nu^2} \leq \frac{5}{\nu^2}.$$

Also, $\inf_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ is non-positive since $f(0) = 0$.

Theorem 1 (Lower bound of SGDM) *Fix a polynomial stepsize sequence $\eta_t = c \cdot t^{-\alpha}$, where $0 \leq \alpha \leq \frac{1}{2}$, momentum factors β in $[0, 1)$ and $\nu > 0$, and a number of iterations $T \geq 1$. Then, there exists a function f in (2) such that the T -th iterate of SGDM with stepsizes η_t and momentum factors β and ν satisfies*

$$f(\mathbf{x}_T) - f^* \geq \frac{c(1-\beta) \ln T}{2\nu T^\alpha}.$$

Before proving it, we stress that $\ln T$ cannot be cancelled by any setting of β . Indeed, the above lower bound can be instantiated by any β and any T . Hence, for a given β , there exists T large enough such that $\ln T$ is constant-times bigger than $\frac{1}{1-\beta}$.

Proof Define a sequence \mathbf{z}_t for $t \in [T+1]$ as follows: $\mathbf{z}_1 = 0$ and

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \nu \eta_t \sum_{i=1}^t \beta^{t-i} \mathbf{h}_i. \quad (3)$$

We will show that \mathbf{z}_t are exactly the updates of SGDM and $f(\mathbf{z}_{T+1}) \geq \Omega\left(\frac{\ln T}{T^\alpha}\right)$. We will use the following two lemmas.

Lemma 2 *Let $b_j = \frac{2j^\alpha}{\nu T^\alpha}$, $a_j = \frac{1-\beta}{2\nu(T-j+1)}$, and $\eta_j = c \cdot j^{-\alpha}$. \mathbf{z}_t is defined as in (3). Then, for $t \leq j$, $z_{t,j} = 0$, and for $t > j$, $z_{t,j} \geq \frac{c}{T^\alpha}$ and $z_{t,j} \leq \frac{2c}{(1-\beta)T^\alpha}$.*

Proof [Proof of Lemma 2] We first prove by induction that when $t \leq j$, $z_{t,j} = 0$. First, $z_1 = 0$. Also, suppose it holds for t . Then, in the case of $t + 1$, for any $j \geq t + 1$,

$$z_{t+1,j} = z_{t,j} - \nu\eta_t \sum_{i=1}^t \beta^{t-i} h_{i,j} = 0 - 0 = 0,$$

which implies $t \leq j$, $z_{t,j} = 0$ holds. Next, we claim that z_t satisfies

$$z_{t,j} \geq \nu \left(b_j \eta_j - \frac{a_j}{1-\beta} \sum_{k=j+1}^t \eta_k \right), \quad j < t \leq T. \quad (4)$$

We prove (4) by induction. For any t , $z_{t,t-1}$ satisfies (4) since

$$z_{t+1,t} = z_{t,t} - \nu\eta_t \sum_{i=1}^t \beta^{t-i} h_{i,t} = -\nu\eta_t h_{t,t} = \nu\eta_t b_t.$$

Then, suppose (4) holds for any $j < t$. We show that it holds for any $j < t + 1$. We already proved for $j = t$. For $j < t$,

$$\begin{aligned} z_{t+1,j} &= z_{t,j} - \nu\eta_t \sum_{i=1}^t \beta^{t-i} h_{i,j} = z_{t,j} - \nu\eta_t \sum_{i=j}^t \beta^{t-i} h_{i,j} \\ &= z_{t,j} + \nu\eta_t \beta^{t-j} b_j - \nu\eta_t \sum_{i=j+1}^t \beta^{t-i} h_{i,j} \geq z_{t,j} - \nu\eta_t \sum_{i=j+1}^t \beta^{t-i} a_j \\ &\geq \nu \left(b_j \eta_j - \frac{a_j}{1-\beta} \sum_{k=j+1}^t \eta_k \right) - \nu\eta_t \sum_{i=j+1}^t \beta^{t-i} a_j \\ &\geq \nu \left(b_j \eta_j - \frac{a_j}{1-\beta} \sum_{k=j+1}^t \eta_k - a_j \sum_{i=j+1}^t \beta^{t-i} \sum_{k=j+1}^t \eta_k \right) \\ &\geq \nu \left(b_j \eta_j - \frac{a_j}{1-\beta} \sum_{k=j+1}^t \eta_k \right), \end{aligned} \quad (5)$$

where in the second inequality we used the induction hypothesis.

Using that $b_j = \frac{2j^\alpha}{\nu T^\alpha}$, $a_j = \frac{1-\beta}{2\nu(T-j+1)}$, and $\eta_j = \frac{c}{j^\alpha}$, we have

$$(5) = \frac{2c}{T^\alpha} - \frac{c}{2(T-j+1)} \sum_{k=j+1}^t \frac{1}{k^\alpha}. \quad (6)$$

By Lemma 8 in the Appendix A, we have that for $0 < \alpha \leq \frac{1}{2}$, (6) $\geq \frac{2c}{T^\alpha} - \frac{c}{T^\alpha} \geq \frac{c}{T^\alpha}$, and for $\alpha = 0$, (6) $\geq 2c - \frac{(t-j-1)c}{(T-j+1)} \geq c$. Thus, we have $z_{t,j} \geq \frac{c}{T^\alpha}$.

For the upper bound, from the above analysis we have

$$z_{t+1,j} \leq z_{t,j} + \nu\eta_t \beta^{t-j} b_j \leq \nu b_j \sum_{i=j}^t \eta_i \beta^{i-j} \leq \nu b_j \eta_j \sum_{i=j}^t \beta^{i-j} \leq \frac{2c}{(1-\beta)T^\alpha}.$$

■

Lemma 3 $f(\mathbf{z}_t) = \mathbf{h}_t^T \mathbf{z}_t$ for any $t \in [T + 1]$. The subgradient oracle for f at \mathbf{z}_t returns \mathbf{h}_t .

Proof [Proof of Lemma 3] We claim that $\mathbf{h}_t^T \mathbf{z}_t = \mathbf{h}_i^T \mathbf{z}_t$ for all $i < t$ and $\mathbf{h}_t^T \mathbf{z}_t > \mathbf{h}_i^T \mathbf{z}_t$ for all $1 \leq i < t$.

When $i < t$, \mathbf{z}_t is supported on the first $t - 1$ coordinates, while \mathbf{h}_t and \mathbf{h}_i agree on the first $t - 1$ coordinates.

In the case of $i > t$, by the definition of \mathbf{z}_t and \mathbf{h}_t , we have

$$\mathbf{z}_t^T (\mathbf{h}_t - \mathbf{h}_i) = \sum_{j=1}^{t-1} z_{t,j} (h_{t,j} - h_{i,j}) = \sum_{j=i}^{t-1} z_{t,j} (h_{t,j} - h_{i,j}) = z_{t,j} (a_i + b_i) + \sum_{j=i+1}^{t-1} z_{t,j} a_j > 0.$$

Thus, we have proved $f(\mathbf{z}_t) = \mathbf{h}_t^T \mathbf{z}_t$ by the definition. Moreover, $\mathcal{I}(\mathbf{z}_t) = \{i : \mathbf{h}_i^T \mathbf{z}_t = f(\mathbf{z}_t)\} = \{t, \dots, T + 1\}$. So the subgradient evaluated at \mathbf{z}_t is \mathbf{h}_t . \blacksquare

Now, we first get a lower bound and an upper bound of \mathbf{z}_t using Lemma 2. Then, by Lemma 3, we have shown that \mathbf{z}_t are exactly the updates of SGDM.

Thus, we have

$$f(\mathbf{z}_{T+1}) = \mathbf{h}_{T+1}^T \mathbf{z}_{T+1} = \sum_{j=1}^T h_{T+1,j} z_{T+1,j} \geq \frac{c}{T^\alpha} \sum_{j=1}^T \frac{1 - \beta}{2\nu(T - j + 1)} \geq \frac{c(1 - \beta) \ln T}{2\nu T^\alpha}.$$

\blacksquare

5. FTRL-based SGDM

The lower bound for the last iterate in the previous section motivates us to study a different variant of SGDM. In particular, we aim to find a way to remove the $\ln T$ term from the convergence rate.

Defazio (2020) points out that the stochastic primal averaging method (Tao et al., 2018) (which is also an instance of Algorithm 1 in Cutkosky (2019) with OMD):

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \gamma_t \mathbf{g}_t, \quad \mathbf{x}_{t+1} = s_t \mathbf{x}_t + (1 - s_t) \mathbf{z}_t$$

could be one-to-one mapped to the momentum method

$$\mathbf{m}_{t+1} = \beta_t \mathbf{m}_t + \mathbf{g}_t, \quad \mathbf{x}_{t+1} = \mathbf{x}_t - \alpha_t \mathbf{m}_t$$

by setting $\gamma_{t+1} = \frac{\gamma_t - \alpha_t}{\beta_{t+1}}$. While this is true, the convergence rate depends on the convergence rate of OMD with time-varying stepsizes, that in turn requires to assume that $\|\mathbf{x}_t - \mathbf{x}^*\|^2 \leq D^2$. This is possible only by using a projection onto a bounded domain in each step.

Thus, to go beyond bounded domains, we propose to study a new variant of SGDM which has the following form (details in Algorithm 1),

$$\mathbf{m}_{t+1} = \beta_t \mathbf{m}_t + (1 - \beta_t) \mathbf{g}_t, \quad \mathbf{x}_{t+1} = s_t \mathbf{x}_t - \alpha_t \mathbf{m}_t.$$

Note the presence of a shrinking factor $s_t \leq 1$ in the iterates in each step. This variant comes naturally when using the primal averaging scheme with FTRL rather than OMD. Hence, we just denote it by FTRL-based SGDM. Now, this momentum variant inherits all the good properties of FTRL. In particular, we no longer need the bounded domain assumption. Moreover, we will show that it guarantees the optimal convergence $O(\frac{1}{\sqrt{T}})$ (Agarwal et al., 2012) of the last iterate for convex and Lipschitz functions.

Algorithm 1 FTRL-based SGDM

- 1: **Input:** A sequence $\alpha_1, \dots, \alpha_T$, with $\alpha_1 > 0$. Non-increasing sequence $\gamma_1, \dots, \gamma_{T-1}$. $\mathbf{m}_0 = 0$. $\mathbf{x}_1 \in \mathbb{R}^d$.
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: Get \mathbf{g}_t at \mathbf{x}_t such that $\mathbb{E}_t[\mathbf{g}_t] = \nabla f(\mathbf{x}_t)$
 - 4: $\beta_t = \frac{\sum_{i=1}^{t-1} \alpha_i}{\sum_{i=1}^t \alpha_i}$ (Define $\sum_{i=1}^0 \alpha_i = 0$)
 - 5: $\mathbf{m}_t = \beta_t \mathbf{m}_{t-1} + (1 - \beta_t) \mathbf{g}_t$
 - 6: $\boldsymbol{\eta}_t = \frac{\alpha_{t+1} \sum_{i=1}^t \alpha_i}{\sum_{i=1}^{t+1} \alpha_i} \gamma_t$
 - 7: $\mathbf{x}_{t+1} = \frac{\sum_{i=1}^t \alpha_i}{\sum_{i=1}^{t+1} \alpha_i} \mathbf{x}_t + \frac{\alpha_{t+1}}{\sum_{i=1}^{t+1} \alpha_i} \mathbf{x}_1 - \boldsymbol{\eta}_t \mathbf{m}_t$
 - 8: **end for**
-

5.1. Convergence Rates for FTRL-based SGDM

We first present a very general theorem for FTRL-based SGDM.

Theorem 4 *Under the assumption in Section 3, Algorithm 1 guarantees*

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \frac{1}{\sum_{t=1}^T \alpha_t} \mathbb{E} \left[\left\| \frac{\mathbf{x}_1 - \mathbf{x}^*}{\sqrt{\gamma_{T-1}}} \right\|^2 + \sum_{t=1}^T \langle \gamma_{t-1}, \alpha_t^2 \mathbf{g}_t^2 \rangle \right].$$

The above theorem is very general and it gives rise to a number of different variations of the FTRL-based SGDM. In particular, we can instantiate it with the following choices.

First, we consider the most used polynomial stepsize $\frac{c}{\sqrt{t}}$ for convex and Lipschitz function, and the constant stepsize $\frac{c}{\sqrt{T}}$ if T is given in advance.

Corollary 1 *Assume (H3) and set $\alpha_t = 1$ for all t . Algorithm 1 with either $\gamma_{t-1} = \frac{c}{G\sqrt{t}} \cdot \mathbf{1}$ or $\gamma_{t-1} = \frac{c}{G\sqrt{T}} \cdot \mathbf{1}$ guarantees*

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2 G}{c\sqrt{T}} + \frac{2cG}{\sqrt{T}}.$$

The above corollary tells that both of these two stepsizes give the optimal bound $O(\frac{1}{\sqrt{T}})$ for the last iterate. Next, we will show that if we use an adaptive² stepsize, Algorithm 1 gives a data-dependent convergence rate for the last iterate. We first consider a global version of the AdaGrad stepsize as in [Streeter and McMahan \(2010\)](#); [Li and Orabona \(2019\)](#); [Ward et al. \(2019\)](#).

Corollary 2 *Assume (H3') and take $\gamma_t = \frac{\alpha \cdot \mathbf{1}}{\sqrt{\epsilon + \sum_{i=1}^t \alpha_i^2 \|\mathbf{g}_i\|^2}}$, $\epsilon > 0$, $1 \leq t \leq T$ and $\alpha_t = 1$. Then, Algorithm 1 guarantees*

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \frac{1}{T} \left[\left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \sqrt{\mathbb{E} \sum_{t=1}^T \|\mathbf{g}_t\|^2} + \epsilon + \frac{\alpha G^2}{\sqrt{\epsilon}} \right].$$

2. Even if widely used in the literature, it is a misnomer to call these stepsize “adaptive”: an algorithm can be adaptive to some unknown quantities (if proved so), not the stepsizes.

We also state a result for the coordinate-wise AdaGrad stepsizes (McMahan and Streeter, 2010; Duchi et al., 2010).

Corollary 3 Assume (H3'') and set $\gamma_t = \frac{\alpha}{\sqrt{\epsilon + \sum_{i=1}^t \alpha_i^2 g_i^2}}$, $\epsilon > 0$, $1 \leq t \leq T$ and $\alpha_t = 1$. Then, Algorithm 1 guarantees

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \frac{1}{T} \left[\left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \sum_{j=1}^d \sqrt{\mathbb{E} \sum_{t=1}^T \mathbf{g}_{t,j}^2} + \epsilon + \frac{\alpha d G_\infty^2}{\sqrt{\epsilon}} \right].$$

The above two corollaries show that the convergence bound are adaptive to the stochastic gradients. In words, in the worst case (i.e., $\sum_{t=1}^T \|\mathbf{g}_t\|^2 = O(T)$ and $\sum_{t=1}^{T-1} \mathbf{g}_{t,j}^2 = O(T)$), the convergence rate is $O(\frac{1}{\sqrt{T}})$. However, when the stochastic gradients are small or sparse, the rate could be much faster than $O(\frac{1}{\sqrt{T}})$. Moreover, the above results give very simple ways to obtain optimal convergence for the last iterate of first-order stochastic methods, that was still unclear if it could be obtained as discussed in Jain et al. (2019, 2021).

Also, we now show that if in addition f is smooth, the last iterate of FTRL-based momentum with the global adaptive stepsize of Corollary 2 gives adaptive rates of convergence that interpolate between $O(\frac{1}{\sqrt{T}})$ and $O(\frac{\ln T}{T})$.

Corollary 4 Assume (H1). Then, under the same assumption and parameter setting of Corollary 2, Algorithm 1 guarantees

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \frac{C}{T} \left(\sqrt{\epsilon + 4L^2 C^2 \ln^2 T + 4LC\sqrt{\epsilon} \ln T} + \frac{2\alpha G^2}{\sqrt{\epsilon}} + \frac{\alpha G^2}{\sqrt{\epsilon}} \right) + \frac{\sqrt{2}C\sigma}{\sqrt{T}}.$$

where $C \triangleq \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right)$.

Observe that when $\sigma = 0$, namely when there is no noise on the gradients, the rate of $O(\frac{\ln T}{T})$ is obtained. As far as we know, the above theorems are the first convergence guarantees for the last iterate of momentum algorithms with adaptive learning rates in unconstrained convex optimization.

5.2. Convergence Rate in Interpolation Regime

Now we assume that $F(\mathbf{x}) = \mathbb{E}_\xi[f(\mathbf{x}, \xi)]$ and that the stochastic gradient is calculated drawing one function in each time step and calculating its gradient: $\mathbf{g}_t = \nabla f(\mathbf{x}_t, \xi_t)$. In this scenario, it makes sense to consider the *interpolation* condition (Needell et al., 2015; Ma et al., 2018)

$$\mathbf{x}^* \in \underset{\mathbf{x}}{\operatorname{argmin}} F(\mathbf{x}) \Rightarrow \mathbf{x}^* \in \underset{\mathbf{x}}{\operatorname{argmin}} f(\mathbf{x}, \xi), \forall \xi. \quad (7)$$

This condition says that the problem is ‘‘easy’’, in the sense that all the functions in the expectation share the same minimizer. This case morally corresponds to the case in which there is no noise on the stochastic gradients. However, this condition seems weaker because it says that only in the optimum the gradient is exact and noisy everywhere else. We will also assume that each function $f(\mathbf{x}, \xi)$ is L -smooth in the first argument.

Algorithm 2 Anytime Online-to-Batch (Cutkosky, 2019)

- 1: **Input:** Online learning algorithm \mathcal{A} with convex domain D , $\alpha_1, \dots, \alpha_T$, with $\alpha_1 > 0$.
 - 2: Get Initial point \mathbf{w}_1 from \mathcal{A}
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: $\mathbf{x}_t = \frac{\sum_{i=1}^t \alpha_i \mathbf{w}_i}{\sum_{i=1}^t \alpha_i}$
 - 5: Play \mathbf{x}_t , receive subgradient \mathbf{g}_t
 - 6: Send $\ell_t(\mathbf{x}) = \langle \alpha_t \mathbf{g}_t, \mathbf{x} \rangle$ to \mathcal{A} as the t th loss
 - 7: Get \mathbf{w}_{t+1} from \mathcal{A}
 - 8: **end for**
-

Algorithm 3 Anytime Online-to-Batch with FTRL

- 1: **Input:** $\alpha_1, \dots, \alpha_T$, with $\alpha_t > 0$. $0 < \gamma_{t+1} \leq \gamma_t$.
 - 2: Initialize \mathbf{w}_1
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: $\mathbf{x}_t = \frac{\sum_{i=1}^t \alpha_i \mathbf{w}_i}{\sum_{i=1}^t \alpha_i}$
 - 5: Play \mathbf{x}_t , receive subgradient \mathbf{g}_t
 - 6: $\mathbf{w}_{t+1} = \mathbf{w}_1 - \gamma_t \sum_{i=1}^t \alpha_i \mathbf{g}_i$
 - 7: **end for**
-

Theorem 5 Assume (H1) (H3'). Then, under the interpolation assumption in (7), Algorithm 1 with

$$\gamma_t = \frac{\alpha \cdot \mathbf{1}}{\sqrt{\epsilon + \sum_{i=1}^t \alpha_i^2 \|\mathbf{g}_i\|^2}}, \epsilon > 0 \text{ guarantees}$$

$$\mathbb{E}[F(\mathbf{x}_T)] - F(\mathbf{x}^*) \leq \frac{C}{T} \left(\sqrt{\epsilon + 4L^2 C^2 \ln^2 T + 4LC\sqrt{\epsilon} \ln T} + \frac{2\alpha G^2}{\sqrt{\epsilon}} + \frac{\alpha G^2}{\sqrt{\epsilon}} \right).$$

where $C \triangleq \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right)$.

To the best of our knowledge, this is the first convergence rate for the last iterate of momentum methods in the interpolation setting.

5.3. Proofs

Before presenting the proofs of our convergence rates, we revisit the Online-to-Batch algorithm (Algorithm 2) by Cutkosky (2019), which introduce a modification to any online learning algorithm to obtain a guarantee on the last iterate in the stochastic convex setting.

Lemma 6 (Cutkosky, 2019, Theorem 1) Assume (H2). Then, for all $\mathbf{x}^* \in D$, Algorithm 2 guarantees

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \mathbb{E} \left[\frac{R_T(\mathbf{x}^*)}{\sum_{t=1}^T \alpha_t} \right].$$

Set $\psi_t(\mathbf{x}) = \frac{\|\mathbf{x}_1 - \mathbf{x}\|^2}{\sqrt{\gamma_{t-1}}}$, $1 \leq t \leq T$ as the regularizers of FTRL, where $\gamma_{t+1} \leq \gamma_t$ and $\gamma_0 > 0$. Then, we write FTRL with loss $\ell_t(\mathbf{w}) = \langle \alpha_t \mathbf{g}_t, \mathbf{w} \rangle$ as

$$\mathbf{w}_t \in \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^d} \psi_t(\mathbf{w}) + \sum_{i=1}^{t-1} \langle \alpha_i \mathbf{g}_i, \mathbf{w} \rangle = \mathbf{w}_1 - \gamma_{t-1} \sum_{i=1}^{t-1} \alpha_i \mathbf{g}_i.$$

We then plug FTRL into Algorithm 2 and it gives Algorithm 3. Hence, using the well-known regret upper bound of FTRL (Lemma 9 in the Appendix B), we get the following Lemma.

Lemma 7 *Under the same setting with Lemma 6, Algorithm 3 guarantees*

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \frac{1}{\sum_{t=1}^T \alpha_t} \mathbb{E} \left[\left\| \frac{\mathbf{u} - \mathbf{x}_1}{\sqrt{\gamma_{T-1}}} \right\|^2 + \sum_{t=1}^T \langle \gamma_{t-1}, \alpha_t^2 \mathbf{g}_t^2 \rangle \right].$$

Now we prove the connection between the FTRL-based SGDM and Algorithm 3.

Proof [Proof of Theorem 4] We prove that the updates of \mathbf{x}_t in Algorithm 1 can be one-to-one mapped to the updates of \mathbf{x}_t Algorithm 3 when $\mathbf{w}_1 = \mathbf{x}_1$.

The update of \mathbf{x}_t in Algorithm 3 can be written as following:

$$\mathbf{x}_{t+1} = \frac{\sum_{i=1}^t \alpha_i}{\sum_{i=1}^{t+1} \alpha_i} \mathbf{x}_t + \frac{\alpha_{t+1}}{\sum_{i=1}^{t+1} \alpha_i} \mathbf{w}_{t+1} = \frac{\sum_{i=1}^t \alpha_i}{\sum_{i=1}^{t+1} \alpha_i} \mathbf{x}_t + \frac{\alpha_{t+1}}{\sum_{i=1}^{t+1} \alpha_i} \left(\mathbf{w}_1 - \gamma_t \sum_{i=1}^t \alpha_i \mathbf{g}_i \right).$$

It is enough to prove that for any t , $\eta_t \mathbf{m}_t = \frac{\alpha_{t+1}}{\sum_{i=1}^{t+1} \alpha_i} (\gamma_t \sum_{i=1}^t \alpha_i \mathbf{g}_i)$. We claim it is true and prove it by induction.

When $t = 1$, it holds that $\eta_1 \mathbf{m}_1 = \frac{\alpha_2 \alpha_1}{\alpha_1 + \alpha_2} \gamma_1 \mathbf{g}_1$. Suppose it holds for $t = k - 1, k \geq 2$. Then in the case of $t = k$, we have

$$\begin{aligned} \eta_k \mathbf{m}_k &= \left(\frac{\sum_{i=1}^{k-1} \alpha_i}{\sum_{i=1}^k \alpha_i} \mathbf{m}_{k-1} + \frac{\alpha_k}{\sum_{i=1}^k \alpha_i} \mathbf{g}_k \right) \cdot \frac{\alpha_{k+1} \sum_{i=1}^k \alpha_i}{\sum_{i=1}^{k+1} \alpha_i} \gamma_k \\ &= \left(\frac{\sum_{i=1}^{k-1} \alpha_i}{\sum_{i=1}^k \alpha_i} \left(\frac{1}{\eta_{k-1}} \frac{\alpha_k}{\sum_{i=1}^k \alpha_i} \gamma_{k-1} \sum_{i=1}^{k-1} \alpha_i \mathbf{g}_i \right) + \frac{\alpha_k}{\sum_{i=1}^k \alpha_i} \mathbf{g}_k \right) \cdot \frac{\alpha_{k+1} \sum_{i=1}^k \alpha_i}{\sum_{i=1}^{k+1} \alpha_i} \gamma_k \\ &= \frac{\alpha_{k+1} \sum_{i=1}^k \alpha_i}{\sum_{i=1}^{k+1} \alpha_i} \gamma_k \cdot \left(\frac{\sum_{i=1}^{k-1} \alpha_i}{\sum_{i=1}^k \alpha_i} \left(\frac{\sum_{i=1}^{k-1} \alpha_i \mathbf{g}_i}{\sum_{i=1}^{k-1} \alpha_i} \right) + \frac{\alpha_k}{\sum_{i=1}^k \alpha_i} \mathbf{g}_k \right) = \frac{\alpha_{k+1}}{\sum_{i=1}^{k+1} \alpha_i} \gamma_k \sum_{i=1}^k \alpha_i \mathbf{g}_i. \end{aligned}$$

where in the first equation we used the definitions of η_k and \mathbf{m}_k and in the second equality we used the induction step. So we proved the above claim. Thus, we can directly use Lemma 7. \blacksquare

The proof of Corollary 1 is immediate and we omit it, while the proofs of Corollaries 2-4 are standard and they are presented in the Appendix C. Instead, here we show the proof of Theorem 5.

Proof [Proof of Theorem 5] By Theorem 4, we have

$$\mathbb{E}[F(\mathbf{x}_T)] - F(\mathbf{x}^*) \leq \frac{2}{T} \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \sqrt{\mathbb{E} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t, \xi_t)\|^2} + \epsilon + \frac{\alpha G^2}{\sqrt{\epsilon}}. \quad (8)$$

Under the interpolation condition and L -smoothness of the functions f , it satisfies that

$$\mathbb{E} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t, \xi_t)\|^2 \leq 2L \mathbb{E} \left[\sum_{t=1}^T (f(\mathbf{x}_t, \xi_t) - f(\mathbf{x}^*, \xi_t)) \right] \leq 2L \sum_{t=1}^T \mathbb{E}[F(\mathbf{x}_t)] - F(\mathbf{x}^*).$$




Figure 1: Squared hinge loss for classification, objective value vs number of epoch.




Figure 2: Hinge loss for classification, objective value vs number of epoch.

Use (8) on each t to get

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[F(\mathbf{x}_t)] - F(\mathbf{x}^*) &\leq \sum_{t=1}^T \frac{1}{t} \left[\left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \sqrt{\mathbb{E} \sum_{i=1}^t \|\nabla f(\mathbf{x}_i, \xi_i)\|^2} + \epsilon + \frac{\alpha G^2}{\sqrt{\epsilon}} \right] \\ &\leq \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \cdot \left(\sqrt{\mathbb{E} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t, \xi_t)\|^2} + \epsilon + \frac{\alpha G^2}{\sqrt{\epsilon}} \right) \ln T. \end{aligned}$$

Then, we solve for $\mathbb{E} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t, \xi_t)\|^2$ and get

$$\begin{aligned} &\mathbb{E} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t, \xi_t)\|^2 \\ &\leq 4L^2 \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right)^2 \ln^2 T + 4L\sqrt{\epsilon} \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \ln T + \frac{2\alpha G^2}{\sqrt{\epsilon}}. \end{aligned}$$

Using this expression in (8), we have the stated bound. ■

6. Empirical Results

We have presented a family of FTRL-based SGDM algorithms, that exhibit optimal convergence of the last iteration. These algorithms are motivated by a new lower bound that shows that constant

momentum SGDM is provably suboptimal to minimize convex Lipschitz functions. However, the theory guarantees only an improvement of $\ln T$, so it is unlikely to make a difference in practical applications. Yet, we also perform some experiments to show that FTRL-based momentum methods have also interesting empirical properties.

We compare FTRL-M (Algorithm 1, $\gamma_t = \frac{c-1}{\sqrt{t}}$), AdaFTRL-M (Algorithm 1, $\gamma_t = \frac{\alpha}{\sqrt{\epsilon + \sum_{i=1}^t \alpha_i^2 g_i^2}}$) with classic SGDM ($\beta = 0.9$), SGDM-AVG (averaged iterates of SGDM, $\beta = 0.9$), and AdaGrad (McMahan and Streeter, 2010; Duchi et al., 2010). The initial stepsizes for all the algorithms were tuned with a fine grid-searching procedure.

Synthetic Data For the first experiment, we generate synthetic data and test the algorithms following the protocol in Vaswani et al. (2019). We generate a synthetic binary classification dataset with $n = 8000$ and the dimension $d = 100$. We make the data linearly separable with a margin, in which case the interpolation condition is satisfied. We train linear classifiers with the squared hinge loss: $f(\mathbf{w}) = \sum_{i=1}^n (\max(0, 1 - \mathbf{y}_i \mathbf{w}^T \mathbf{x}_i))^2$. Note that the loss function is smooth and $f(\mathbf{w}^*) = 0$. In this case, the optimal convergence rate is at least as fast as $1/T$.

We plot the suboptimality gap versus the number of epochs with different margin values in Figure 1, in loglog plots. Also, we add a line to fit the curves, where the slopes represent the power of t . From Figure 1, we observe that two adaptive algorithms AdaGrad and AdaFTRL-M bring faster convergence and AdaFTRL-M has the biggest slope in all the cases. Also, the performance of FTRL-M is on par with SGDM and SGDM-AVG.

Real Data We also test the algorithms on real datasets. We use classification datasets from the LIBSVM website (Chang and Lin, 2001); *real-sim*, *w8a*, and *phishing*. The details of the datasets are in Appendix D.

We train linear classifiers with the hinge loss and no regularization: $f(\mathbf{w}) = \sum_{i=1}^n (\max(0, 1 - \mathbf{y}_i \mathbf{w}^T \mathbf{x}_i))$. The stochastic gradients are obtained evaluating the subgradient on one example at the time. We repeat the experiments for 5 times for each algorithm and report the average of 5 repetitions. We show the objective value versus the number of epochs in Figure 2.

The results show that the algorithms with non-adaptive stepsizes tend to perform worse than the ones with adaptive stepsizes. Moreover, the performance of AdaFTRL-M is close to the last iterate of AdaGrad and sometimes outperforms all the other algorithms, especially in the last iterations.

7. Conclusion

We have presented an analysis of the convergence of the last iterate of SGDM in the convex setting. We prove for the first time through a lower bound the suboptimal convergence rate for the last iterate of SGDM with constant momentum after T iterations. Moreover, we study a class of FTRL-based SGDM algorithms with increasing momentum and shrinking updates, of which the last iterate has optimal convergence rate without projections onto bounded domain nor knowledge of T . Furthermore, we present empirical results showing that FTRL-based SGDM with adaptive stepsize matches or outperforms the other similar algorithms in the last iterations.

In the future, we plan on studying the convergence in high probability of FTRL-based SGDM, similarly to the analysis in Li and Orabona (2020).

Acknowledgments

This material is based upon work supported by the National Science Foundation under the grants no. 1925930 “Collaborative Research: TRIPODS Institute for Optimization and Learning” and no. 2046096 “CAREER: Parameter-free Optimization Algorithms for Machine Learning”.

References

- J. D. Abernethy, E. Hazan, and A. Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. In Rocco A. Servedio and Tong Zhang, editors, *Proc. of Conference on Learning Theory (COLT)*, pages 263–274. Omnipress, 2008.
- A. Agarwal, P. L. Bartlett, P. Ravikumar, and M. J. Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, 2012.
- A. Alacaoglu, Y. Malitsky, P. Mertikopoulos, and V. Cevher. A new regret analysis for Adam-type algorithms. In *International Conference on Machine Learning*, pages 202–210. PMLR, 2020.
- C.-C. Chang and C.-J. Lin. *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- A. Cutkosky. Anytime online-to-batch, optimism and acceleration. In K. Chaudhuri and R. Salakhutdinov, editors, *Proc. of the 36th International Conference on Machine Learning*, volume 97 of *Proc. of Machine Learning Research*, pages 1446–1454, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- A. Cutkosky and F. Orabona. Momentum-based variance reduction in non-convex SGD. In *Advances in Neural Information Processing Systems*, pages 15236–15245, 2019.
- A. Defazio. Understanding the role of momentum in non-convex optimization: Practical insights from a lyapunov analysis. *arXiv preprint arXiv:2010.00406*, 2020.
- A. Dieuleveut, N. Flammarion, and F. Bach. Harder, better, faster, stronger convergence rates for least-squares regression. *J. Mach. Learn. Res.*, 18(1):3520–3570, January 2017.
- J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. In *COLT*, 2010.
- P. Gaillard, G. Stoltz, and T. Van Erven. A second-order bound with excess losses. In *Conference on Learning Theory*, pages 176–196. PMLR, 2014.
- S. Ghadimi and G. Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal on Optimization*, 22(4):1469–1492, 2012.
- N. J. Harvey, C. Liaw, Y. Plan, and S. Randhawa. Tight analyses for non-smooth stochastic gradient descent. In *Conference on Learning Theory*, pages 1579–1613, 2019.

- P. Jain, S. M. Kakade, R. Kidambi, P. Netrapalli, and A. Sidford. Accelerating stochastic gradient descent for least squares regression. In S. Bubeck, V. Perchet, and P. Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 545–604. PMLR, 06–09 Jul 2018.
- P. Jain, D. Nagaraj, and P. Netrapalli. Making the last iterate of SGD information theoretically optimal. In A. Beygelzimer and D. Hsu, editors, *Proc. of the Conference on Learning Theory (COLT)*, volume 99 of *Proceedings of Machine Learning Research*, pages 1752–1755, Phoenix, USA, 25–28 Jun 2019. PMLR.
- P. Jain, D. M. Nagaraj, and P. Netrapalli. Making the last iterate of sgd information theoretically optimal. *SIAM Journal on Optimization*, 31(2):1108–1130, 2021. doi: 10.1137/19M128908X.
- S. Jelassi and A. Defazio. Dual averaging is surprisingly effective for deep learning optimization. *arXiv preprint arXiv:2010.10502*, 2020.
- R. Kidambi, P. Netrapalli, P. Jain, and S. Kakade. On the insufficiency of existing momentum schemes for stochastic optimization. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–9. IEEE, 2018.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- G. Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, 2012.
- X. Li and F. Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In *Proc. of the 22nd International Conference on Artificial Intelligence and Statistics, AISTATS*, 2019.
- X. Li and F. Orabona. A high probability analysis of adaptive SGD with momentum. In *ICML 2020 Workshop on Beyond First Order Methods in ML Systems*, 2020.
- Y. Liu, Y. Gao, and W. Yin. An improved analysis of stochastic gradient descent with momentum. *Advances in Neural Information Processing Systems*, 33, 2020.
- L. Luo, Y. Xiong, Y. Liu, and X. Sun. Adaptive gradient methods with dynamic bound of learning rate. In *International Conference on Learning Representations*, 2018.
- S. Ma, R. Bassily, and M. Belkin. The power of interpolation: Understanding the effectiveness of SGD in modern over-parametrized learning. In *International Conference on Machine Learning*, pages 3325–3334. PMLR, 2018.
- H. B. McMahan and M. J. Streeter. Adaptive bound optimization for online convex optimization. In *COLT*, 2010.
- D. Needell, N. Srebro, and R. Ward. Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm. *Mathematical Programming*, 155:549–573, 2015.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.

- A. S. Nemirovsky and D. Yudin. *Problem complexity and method efficiency in optimization*. Wiley, New York, NY, USA, 1983.
- Y. Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In *Doklady AN SSSR (translated as Soviet. Math. Docl.)*, volume 269, pages 543–547, 1983.
- Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical programming*, 120(1):221–259, 2009.
- Y. Nesterov and V. Shikhman. Quasi-monotone subgradient methods for nonsmooth convex minimization. *Journal of Optimization Theory and Applications*, 165(3):917–940, 2015.
- F. Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- B. T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- S. J. Reddi, A. Hefny, S. Sra, B. Póczos, and A. Smola. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*, pages 314–323, 2016.
- S. J. Reddi, S. Kale, and S. Kumar. On the convergence of Adam and beyond. In *International Conference on Learning Representations*, 2018.
- O. Sebbouh, R. M Gower, and A. Defazio. Almost sure convergence rates for stochastic gradient descent and stochastic heavy ball. In *Conference on Learning Theory*, pages 3935–3971. PMLR, 2021.
- S. Shalev-Shwartz. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, The Hebrew University, 2007.
- M. Streeter and H. B. McMahan. Less regret via online conditioning, 2010. arXiv:1002.4862.
- I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147, 2013.
- W. Tao, Z. Pan, G. Wu, and Q. Tao. Primal averaging: A new gradient evaluation step to attain the optimal individual convergence. *IEEE transactions on cybernetics*, 50(2):835–845, 2018.
- W. Tao, S. Long, G. Wu, and Q. Tao. The role of momentum parameters in the optimal convergence of adaptive Polyak’s heavy-ball methods. In *International Conference on Learning Representations*, 2021.
- S. Vaswani, F. Bach, and M. Schmidt. Fast and faster convergence of SGD for over-parameterized models and an accelerated Perceptron. In K. Chaudhuri and M. Sugiyama, editors, *Proc. of the 22nd International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proc. of Machine Learning Research*, pages 1195–1204. PMLR, 16–18 Apr 2019.
- R. Ward, X. Wu, and L. Bottou. AdaGrad stepsizes: Sharp convergence over nonconvex landscapes. In *International Conference on Machine Learning*, pages 6677–6686. PMLR, 2019.

M. K. Warmuth and A. K. Jagota. Continuous and discrete-time nonlinear gradient descent: Relative loss bounds and convergence. In *Electronic proceedings of the 5th International Symposium on Artificial Intelligence and Mathematics*, volume 326, 1997.

T. Yang, Q. Lin, and Z. Li. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. *arXiv preprint arXiv:1604.03257*, 2016.

Appendix A. Lemma for the Proof of Theorem 1

Lemma 8 For any $1 \leq j \leq t \leq T$ and $0 < \alpha \leq \frac{1}{2}$, we have $\frac{1}{T-j+1} \sum_{k=j+1}^t \frac{1}{k^\alpha} \leq \frac{2}{T^\alpha}$.

Proof First, we observe that

$$\begin{aligned} \sum_{k=j+1}^t \frac{1}{k^\alpha} &\leq \int_j^t \frac{1}{x^\alpha} dx = \frac{t^{1-\alpha} - j^{1-\alpha}}{1-\alpha} = \frac{1}{1-\alpha} \frac{t^{2-2\alpha} - j^{2-2\alpha}}{t^{1-\alpha} + j^{1-\alpha}} \\ &\leq \frac{1}{1-\alpha} \frac{(2-2\alpha)t^{1-2\alpha}(t-j)}{t^{1-\alpha} + j^{1-\alpha}} \leq \frac{2(t-j)}{t^\alpha}, \end{aligned}$$

where in the second inequality we used the convexity of $f(x) = x^{2-2\alpha}$, $0 < \alpha \leq \frac{1}{2}$.

Then, we claim $\frac{1}{T-j+1} \frac{t-j}{t^\alpha} \leq \frac{1}{T^\alpha}$.

Let $g(x) = \frac{x-j}{x^\alpha}$. The derivative $g'(x) = \frac{1-\alpha+\frac{j}{x}}{x^\alpha}$ is positive for all $x > 0$ and $j \geq 0$. So it satisfies that $\frac{t-j}{t^\alpha} \leq \frac{T-j}{T^\alpha}$, which implies the claim. \blacksquare

Appendix B. Lemma for the Proof of Theorem 4

Algorithm 4 Follow-the-Regularized-Leader on Linearized Losses

- 1: **Input:** Regularizers $\psi_1, \dots, \psi_T : \mathbb{R}^d \rightarrow (-\infty, \infty]$.
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: $\mathbf{w}_t \in \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^d} \psi_t(\mathbf{w}) + \sum_{i=1}^{t-1} \langle \mathbf{g}_i, \mathbf{w} \rangle$
 - 4: Receive $\ell_t : \mathbb{R}^d \rightarrow (-\infty, \infty]$ and pay $\ell_t(\mathbf{w}_t)$
 - 5: Set $\mathbf{g}_t \in \partial \ell_t(\mathbf{w}_t)$
 - 6: **end for**
-

The following lemma is a well-known result for FTRL (see, e.g., [Orabona, 2019](#)).

Lemma 9 Let ℓ_t a sequence of convex loss functions. Set the sequence of regularizers as $\psi_t(\mathbf{x}) = \left\| \frac{\mathbf{x}_1 - \mathbf{u}}{\sqrt{\gamma_{t-1}}} \right\|^2$, where $\gamma_{t+1} \leq \gamma_t$, $t = 1, \dots, T$. Then, FTRL (Algorithm 4) guarantees

$$\sum_{t=1}^T \ell_t(\mathbf{x}_t) - \ell_t(\mathbf{u}) \leq \left\| \frac{\mathbf{u} - \mathbf{x}_1}{\sqrt{\gamma_{T-1}}} \right\|^2 + \frac{1}{2} \sum_{t=1}^T \langle \gamma_{t-1}, \mathbf{g}_t^2 \rangle.$$

Appendix C. Proofs of Corollaries 2-4

First, we state some technical lemmas.

Lemma 10 (*Li and Orabona, 2019, Lemma 4*) Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be M -smooth and bounded from below, then for all $\mathbf{x} \in \mathbb{R}^d$

$$\|\nabla f(\mathbf{x})\|^2 \leq 2M(f(\mathbf{x}) - \inf_{\mathbf{y} \in \mathbb{R}^d} f(\mathbf{y})) .$$

Lemma 11 (*Gaillard et al., 2014, Lemma 14*) Let $a_0 > 0$ and $a_1, \dots, a_m \in [0, A]$ be real numbers and let $f : (0, +\infty) \rightarrow [0, +\infty)$ nonincreasing function. Then

$$\sum_{i=1}^m a_i f(a_0 + \dots + a_{i-1}) \leq \int_{a_0}^{\sum_{i=0}^m a_i} f(u) du + Af(a_0) .$$

Proof Denote by $s_t = \sum_{i=0}^t a_i$.

$$\begin{aligned} \sum_{i=1}^m a_i f(s_{i-1}) &= \sum_{i=1}^m a_i f(s_i) + \sum_{i=1}^m a_i (f(s_{i-1}) - f(s_i)) \\ &\leq \sum_{i=1}^m a_i f(s_i) + A \sum_{i=1}^m (f(s_{i-1}) - f(s_i)) \\ &\leq \sum_{i=1}^m \int_{s_{i-1}}^{s_i} f(x) dx + A \sum_{i=1}^m (f(s_{i-1}) - f(s_i)) \\ &\leq \int_{a_0}^{\sum_{i=0}^m a_i} f(u) du + Af(a_0) , \end{aligned}$$

where the first inequality holds because $f(x_{i-1}) \geq f(s_i)$ and $a_i \leq A$, while the second inequality uses the fact that f is nonincreasing together with $s_i - s_{i-1} = a_i$. \blacksquare

We can now present the proofs of the Corollaries 2-4.

Proof [Proof of Corollary 2 and Corollary 3] By Lemma 11, for adaptive stepsize $\gamma_t = \frac{\alpha \cdot 1}{\sqrt{\epsilon + \sum_{i=1}^t \alpha_i^2 \|\mathbf{g}_i\|^2}}$, we have

$$\sum_{t=1}^T \gamma_{t-1} \|\mathbf{g}_t\|^2 = \sum_{t=1}^T \frac{\alpha \|\mathbf{g}_t\|^2}{\sqrt{\epsilon + \sum_{i=1}^{t-1} \|\mathbf{g}_i\|^2}} \leq 2\alpha \sqrt{\sum_{t=1}^T \|\mathbf{g}_t\|^2} + \frac{\alpha G^2}{\sqrt{\epsilon}} .$$

Similarly for $\gamma_t = \frac{\alpha}{\sqrt{\epsilon + \sum_{i=1}^t \alpha_i^2 \mathbf{g}_i^2}}$, we have

$$\sum_{t=1}^T \langle \gamma_{t-1}, \mathbf{g}_t^2 \rangle = \sum_{j=1}^d \sum_{t=1}^T \frac{\alpha \mathbf{g}_{t,j}^2}{\sqrt{\epsilon + \sum_{i=1}^{t-1} \mathbf{g}_{i,j}^2}} \leq 2\alpha \sum_{j=1}^d \sqrt{\sum_{t=1}^T \mathbf{g}_{t,j}^2} + \frac{\alpha d G_\infty^2}{\sqrt{\epsilon}} .$$

\blacksquare

Proof [Proof of Corollary 4] By Corollary 2, we have

$$\mathbb{E}[f(\mathbf{x}_T)] - f^* \leq \frac{1}{T} \left[\left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \sqrt{\epsilon + \mathbb{E} \sum_{t=1}^T \|\mathbf{g}_t\|^2} + \frac{\alpha G^2}{\sqrt{\epsilon}} \right]. \quad (9)$$

From the unbiasedness of the gradients, we have

$$\mathbb{E} \sum_{t=1}^T \|\mathbf{g}_t\|^2 \leq \mathbb{E} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2 + T\sigma^2,$$

and

$$\begin{aligned} \mathbb{E} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2 &\leq 2L \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}_t)] - f^* \\ &\leq 2L \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \sum_{t=1}^T \frac{\sqrt{\epsilon + \mathbb{E} \sum_{i=1}^t \|\mathbf{g}_i\|^2}}{t} \\ &\leq 2L \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \cdot \left(\sqrt{\mathbb{E} \sum_{t=1}^T \|\mathbf{g}_t\|^2} + \epsilon + \frac{\alpha G^2}{\sqrt{\epsilon}} \right) \ln T, \end{aligned}$$

where in the second inequality we used Lemma 10 and Holder's and Jensen's inequalities in the third inequality.

Solve for $\mathbb{E} \sum_{t=1}^T \|\mathbf{g}_t\|^2$ to have

$$\begin{aligned} &\mathbb{E} \sum_{t=1}^T \|\mathbf{g}_t\|^2 \\ &\leq 4L^2 \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right)^2 \ln^2 T + 4L\sqrt{\epsilon} \left(\frac{\|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\alpha} + 2\alpha \right) \ln T + \frac{2\alpha G^2}{\sqrt{\epsilon}} + 2T\sigma^2. \end{aligned}$$

Putting it back to (9), we get the stated bound. ■

Appendix D. Details of Experiments

Table 2: Real datasets

Name	# of Samples	# of Features
real-sim	72,309	20,958
w8a	49,749	300
phishing	11,055	68