

Almost sure convergence rates for Stochastic Gradient Descent and Stochastic Heavy Ball

Othmane Sebbouh

CNRS, ENS, PSL University, CREST, ENSAE, Paris, France

OTHMANE.SEBBOUH@ENS.FR

Robert M. Gower

LTCI, Télécom Paris, Institut Polytechnique de Paris, Paris, France

ROBERT.GOWER@TELECOM-PARIS.FR

Aaron Defazio

Facebook AI Research, New York, USA

ADEFAZIO@FB.COM

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

We study stochastic gradient descent (SGD) and the stochastic heavy ball method (SHB, otherwise known as the momentum method) for the general stochastic approximation problem. For SGD, in the convex and smooth setting, we provide the first *almost sure* asymptotic convergence rates for a weighted average of the iterates. More precisely, we show that the convergence rate of the function values is arbitrarily close to $o(1/\sqrt{k})$, and is exactly $o(1/k)$ in the so-called overparametrized case. We show that these results still hold when using a decreasing step size version of stochastic line search and stochastic Polyak stepsizes, thereby giving the first proof of convergence of these methods in the non-overparametrized regime. Using a substantially different analysis, we show that these rates hold for SHB as well, but at the last iterate. This distinction is important because it is the last iterate of SGD and SHB which is used in practice. We also show that the last iterate of SHB converges to a minimizer *almost surely*. Additionally, we prove that the function values of the deterministic HB converge at a $o(1/k)$ rate, which is faster than the previously known $O(1/k)$. Finally, in the nonconvex setting, we prove similar rates on the lowest gradient norm along the trajectory of SGD.

Keywords: Stochastic Gradient Descent, Stochastic Heavy Ball, Momentum Method, Almost sure convergence, convex, nonconvex.

1. Introduction

Consider the stochastic approximation problem

$$x_* \in \operatorname{argmin}_{x \in \mathbb{R}^d} f(x) \stackrel{\text{def}}{=} \mathbb{E}_{v \sim \mathcal{D}} [f_v(x)], \quad (1)$$

where \mathcal{D} is a distribution on an arbitrary space Ω and f_v is a real-valued function. Let $\mathcal{X}_* \subset \mathbb{R}^d$ be the set of solutions of (1) (which we assume to be nonempty) and $f_* = f(x_*)$ for any solution $x_* \in \mathcal{X}_*$. The stochastic approximation problem (1) encompasses several problems in machine learning, including Online Learning and Empirical Risk Minimization (ERM). In these settings, when the function f can be accessed only through sampling or when the size of the datasets is very high, first-order stochastic gradient methods have proven to be very effective thanks to their low iteration complexity. The methods we analyze, Stochastic Gradient descent (SGD, (Robbins and Monro, 1951)) and Stochastic Heavy Ball (SHB, (Polyak, 1964)), are among the most popular such methods.

1.1. Contributions and Background

Here we summarize the relevant background and our contributions. All of our rates of convergence are also given succinctly in Table 1.

Almost sure convergence rates for SGD. The *almost sure* convergence of the iterates of SGD is a well-studied question (Bottou, 2003; Zhou et al., 2017; Nguyen et al., 2018). For functions satisfying $\forall(x, x_*) \in \mathbb{R}^d \times \mathcal{X}_*, \langle \nabla f(x), x - x_* \rangle \geq 0$, called *variationally coherent*, the convergence was shown in Bottou (2003) by assuming that the minimizer is unique. Recently in Zhou et al. (2017), the uniqueness assumption of the minimizer was dropped for variationally coherent functions by assuming bounded gradients. The easier question of the *almost sure* convergence of the norm of the gradients of SGD in the nonconvex setting, and of the objective values in the convex setting, has also been positively answered by several works, see Bertsekas and Tsitsiklis (2000) and references therein, or more recently Mertikopoulos et al. (2020); Orabona (2020a). In this work, we aim to quantify this convergence. Indeed, while convergence rates are commonplace for convergence in expectation (Nemirovski et al. (2009); Bach and Moulines (2011); Ghadimi and Lan (2013) for example), the literature on the convergence rates of SGD in the *almost sure* sense is sparse. For an adaptive SGD method, Li and Orabona (2019) prove the convergence of a subsequence of the squared gradient at a rate arbitrarily close to $o(1/\sqrt{k})$. More precisely, they show that $\liminf_k k^{\frac{1}{2}-\epsilon} \|\nabla f(x_k)\|^2 = 0$ for all $\epsilon > 0$, where x_k is the k th iterate of SGD. Godichon-Baggioni (2016) proves that, for locally strongly convex functions, the sequence $(\|x_k - x_*\|^2)_k$, where x_* is the unique minimizer of f , converges *almost surely* at a rate arbitrarily close to $o(1/k)$.

Contributions: 1. In the convex and smooth setting, we show that the function values at a weighted average of the iterates of SGD converge *almost surely* at a rate arbitrarily close to $o(1/\sqrt{k})$. In the so-called overparametrized case, where the stochastic gradients at any minimizer $\nabla f_v(x_*)$ are 0, we show that this rate improves to $o(1/k)$. The proof of these results is surprisingly simple, and relies on a new weighted average of the iterates of SGD and on the classical Robbins-Siegmund supermartingale convergence theorem (Lemma 6). We also complement the well-known Robbins-Monro (Robbins and Monro, 1951) conditions on the stepsizes with new conditions (See Condition 1) that allow us to derive convergence rates in the *almost sure* sense. We also show that our theory still holds in the nonsmooth setting when we assume bounded subgradients (Appendix F). **2.** In the nonconvex setting, under the recently introduced *ABC condition* (Khaled and Richtárik, 2020), we derive *almost sure* convergence rates for the minimum squared gradient norm along the trajectory of SGD which match the rates we derived for the objective values of SGD.

Asymptotic convergence of SGD with adaptive step sizes. One drawback of the theory of SGD in the smooth setting is that it relies on the knowledge of the smoothness constant. Two of the earliest methods which have been proposed to address this issue are Line-Search (LS) (Nocedal and Wright, 2006) and Polyak Stepsizes (PS) (Polyak, 1987). But while their convergence had been established in the deterministic case, it wasn't until recently (Vaswani et al., 2019b, 2020; Loizou et al., 2020) that SGD with LS and with PS has been shown to converge assuming only smoothness and convexity of the functions f_v . For both methods, it has been shown that SGD converges to the minimum at a rate $O(1/k)$ in the overparametrized setting, but converges only to a neighborhood of the minimum when overparametrization does not hold.

Contributions. We show that SGD with LS or PS converges asymptotically at a rate arbitrarily close to $O(1/\sqrt{k})$ in expectation, and to $o(1/\sqrt{k})$ *almost surely* when using decreasing maximum step sizes. Moreover, in the overparametrized setting, using the proof technique we developed for SGD, we show that SGD with LS or PS converges *almost surely* to the minimum at a $o(1/k)$ rate.

Almost sure convergence rates for SHB and $o(1/k)$ convergence for HB. The first local convergence of the deterministic Heavy Ball method was given in Polyak (1964), showing that it converges at an accelerated rate for twice differentiable strongly convex functions. Only recently did Ghadimi et al. (2015) show that the deterministic Heavy Ball method converged globally and sublinearly for smooth and convex functions. SHB has recently been analysed for nonconvex functions and for strongly convex functions in Gadat et al. (2018). For strongly convex functions, they prove a $O(1/t^\beta)$ convergence rate for any $\beta < 1$. Using a similar Lyapunov function to the one in Ghadimi et al. (2015), a $O(1/\sqrt{t})$ convergence rate for SHB in the convex setting was given in Yang et al. (2016) and Orvieto et al. (2019) under the bounded gradient variance assumption. For the specialized setting of minimizing quadratics, it has been shown that the SHB iterates converge linearly at an accelerated rate, but only in expectation rather than in L2 (Loizou and Richtárik, 2018). By using stronger assumptions on the noise as compared to Kidambi et al. (2018), Can et al. (2019) show that by using a specific parameter setting, the SHB applied on quadratics converges at an accelerated rate to a neighborhood of a minimizer. Finally, the *almost sure* convergence of SHB to a minimizer for nonconvex functions was proven in Gadat et al. (2018) under an *elliptic condition* which guarantees that SHB escapes any unstable point. But we are not aware of any convergence rates for the *almost sure* convergence of SHB.

Contributions. **1.** In the smooth and convex setting, we show that the function values at the last iterate of SHB converge *almost surely* at a rate close to $o(1/\sqrt{k})$. Similarly to SGD, this rate can be improved to $o(1/k)$ in the overparametrized setting. Moreover, we show that the last iterate of SHB converges to a minimizer *almost surely*. In the deterministic setting, where we use the gradient ∇f at each iteration, we prove that the function values of the deterministic HB converge at a $o(1/k)$ rate, which is faster than the previously known $O(1/k)$ (Ghadimi et al., 2015) and matches the rate recently derived for Gradient Descent in Lee and Wright (2019). Compared to the SGD analysis we develop, the derivation of *almost sure* convergence rates for SHB is quite involved, and combines tools developed in Attouch and Peypouquet (2016) for the analysis of the (deterministic) Nesterov Accelerated Gradient method and the classical Robbins-Siegmund theorem. **2.** Our results rely on an iterate averaging viewpoint of SHB (Proposition 5), which considerably simplifies our analysis and suggests parameter settings different from the usual settings of the momentum parameter, which is fixed at around 0.9, and often exhibits better empirical performance than SGD (Sutskever et al., 2013). We show through extensive numerical experiments in Figure 1 that our new parameter setting is statistically superior to the standard rule-of-thumb settings on convex problems. **3.** Additionally, we show in Appendix G that the bounded gradients and bounded noise assumptions used in Yang et al. (2016); Orvieto et al. (2019) can be avoided, and prove that SHB at the last iterate converges in expectation at a $O(1/k)$ rate to a neighborhood of the minimum and at a $O(1/\sqrt{k})$ rate to the minimum exactly.

1.2. Assumptions and general consequences

Our theory in the convex setting relies on the following assumption of convexity and smoothness.

Assumption 1 For all $v \sim \mathcal{D}$, there exists $L_v > 0$ such that for every $x, y \in \mathbb{R}^d$ we have that

$$f_v(y) \geq f_v(x) + \langle \nabla f_v(x), y - x \rangle, \quad (2)$$

$$f_v(y) \leq f_v(x) + \langle \nabla f_v(x), y - x \rangle + \frac{L_v}{2} \|y - x\|^2, \quad (3)$$

almost surely. Let $\mathcal{L} \stackrel{\text{def}}{=} \sup_{v \sim \mathcal{D}} L_v$. We assume that $\mathcal{L} < \infty$. Consequently, f is also smooth and we use $L > 0$ to denote its smoothness constant.

Definition 1 Define the residual gradient noise as

$$\sigma^2 \stackrel{\text{def}}{=} \sup_{x \in \mathcal{X}^*} \mathbb{E}_{v \sim \mathcal{D}} \left[\|\nabla f_v(x^*)\|^2 \right]. \quad (4)$$

Assumption 1 has the following simple consequence on the expectation of the gradients.

Lemma 2 If Assumption 1 holds, then

$$\mathbb{E}_{\mathcal{D}} \left[\|\nabla f_v(x)\|^2 \right] \leq 4\mathcal{L} (f(x) - f_*) + 2\sigma^2. \quad (5)$$

In all our results of Sections 2 and 3, we only use convexity and the inequality (5). Thus, Assumption 1 can be slightly relaxed by removing the smoothness condition (3) and re-branding (5) as an assumption, as opposed to a consequence. With this bound (5), we do not need to assume a uniform bound on the squared norm of the gradients or on their variance, as is often done when analyzing SGD (Nemirovski et al., 2009) or SHB (Yang et al., 2016). Note, however, that the analysis carried for SGD and SHB in Nemirovski et al. (2009) and Yang et al. (2016) is more general and applies to the nonsmooth case, for which assuming bounded subgradients is often necessary. As an illustration, we show that our results hold in the nonsmooth case under the bounded subgradients assumption in Appendix F. Note also that all our results still hold with the usual but more restrictive assumption of bounded gradient variance (see for example Ghadimi and Lan (2013)). Indeed, when this assumption holds, (5) holds with L in place of \mathcal{L} , where L is the smoothness constant of f .

Definition 3 (Informal) When $\sigma^2 = 0$, we say that we have an overparametrized model.

When our models have enough parameters to *interpolate the data* (Vaswani et al., 2019a), then $\nabla f_v(x^*) = 0$, $\forall v \sim \mathcal{D}$, and consequently $\sigma^2 = 0$. This property has been observed especially for the training of large neural networks in Empirical Risk Minimization, where f is a finite-sum.

Remark 4 (Finite-sum setting) Let $n \in \mathbb{N}^*$ and define $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$. Let $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$, where each f_i is L_i -smooth and convex, and $L_{\max} = \max_{i \in [n]} L_i < \infty$. If we sample minibatches of size b without replacement, then Gazagnadou et al. (2019); Gower et al. (2019) show that (5) holds with

$$\mathcal{L} \equiv \mathcal{L}(b) \stackrel{\text{def}}{=} \frac{1}{b} \frac{n-b}{n-1} L_{\max} + \frac{n-b-1}{b} \frac{b-1}{n-1} L \quad \text{and} \quad \sigma^2 \equiv \sigma^2(b) = \frac{1}{b} \frac{n-b}{n-1} \sigma_1^2,$$

where $\sigma_1^2 \stackrel{\text{def}}{=} \frac{1}{n} \sup_{x \in \mathcal{X}^*} \sum_{i=1}^n \|\nabla f_i(x_*)\|^2$. Note that $\sigma^2(n) = 0$ and $\mathcal{L}(n) = L$, as expected, since $b = n$ corresponds to full batch gradients, or equivalently to using deterministic GD or HB. Similarly, $\mathcal{L}(1) = L_{\max}$, since $b = 1$ corresponds to sampling one individual f_i function.

1.3. SGD and an iterate-averaging viewpoint of SHB

In Section 2, we will analyze SGD, where we sample at each iteration $v_k \sim \mathcal{D}$, and iterate

$$\boxed{x_{k+1} = x_k - \eta_k \nabla f_{v_k}(x_k)}, \quad (\text{SGD})$$

where η_k is a step size. In Section 3, we will analyze SHB, whose iterates are

$$x_{k+1} = x_k - \alpha_k \nabla f_{v_k}(x_k) + \beta_k (x_k - x_{k-1}), \quad (\text{SHB})$$

where α_k is commonly referred to as the step size and β_k as the momentum parameter. Our forthcoming analysis of (SHB) leverages an *iterate moving-average viewpoint* of (SHB) and particular parameter choices that we present in Proposition 5.

Proposition 5 *Let $z_0 = x_0 \in \mathbb{R}^d$ and $\eta_k, \lambda_k > 0$. Consider the iterate-moving-average (IMA) method:*

$$\boxed{z_{k+1} = z_k - \eta_k \nabla f_{v_k}(x_k), \quad x_{k+1} = \frac{\lambda_{k+1}}{\lambda_{k+1} + 1} x_k + \frac{1}{\lambda_{k+1} + 1} z_{k+1}} \quad (\text{SHB-IMA})$$

$$\text{If} \quad \alpha_k = \frac{\eta_k}{1 + \lambda_{k+1}} \quad \text{and} \quad \beta_k = \frac{\lambda_k}{1 + \lambda_{k+1}}, \quad (6)$$

then the x_k iterates in (SHB-IMA) are equal to the x_k iterates of the method (SHB).

The equivalence between this formulation and the original SHB is proven in the supplementary material (Section B.2). The IMA formulation (SHB-IMA) is crucial in comparing SHB and SGD as it allows to interpret the parameter α_k in SHB as a *scaled* step size and unveils a natural stepsize η_k . In all of our theorems, the parameters η_k and λ_k naturally arise in the recurrences and Lyapunov functions. We determine how to set the parameters η_k and λ_k , which in turn gives settings for α_k and β_k through (6). In the remainder of this work, we will directly analyze the method SHB-IMA.

Having new reformulations often leads to new insights. This is the case for Nesterov's accelerated gradient method, where at least six forms are known (Defazio, 2019) and recent research suggests that iterate-averaged reformulations are the easiest to generalize to the combined proximal and variance-reduced case (Lan and Zhou, 2017).

2. Almost sure convergence rates for SGD and SGD with adaptive stepsizes

We will first present *almost sure* convergence rates for SGD, then for SGD with Line-Search and Polyak Stepsizes.

2.1. SGD: average-iterates almost sure convergence

Our results rely on a classical convergence result (Robbins and Siegmund, 1971).

Lemma 6 *Consider a filtration $(\mathcal{F}_k)_k$, the nonnegative sequences of $(\mathcal{F}_k)_k$ -adapted processes $(V_k)_k$, $(U_k)_k$ and $(Z_k)_k$, and a sequence of positive numbers $(\gamma_k)_k$ such that $\sum_k Z_k < \infty$ almost surely, $\prod_{k=0}^{\infty} (1 + \gamma_k) < \infty$, and*

$$\forall k \in \mathbb{N}, \quad \mathbb{E} [V_{k+1} | \mathcal{F}_k] + U_{k+1} \leq (1 + \gamma_k) V_k + Z_k.$$

Then $(V_k)_k$ converges and $\sum_k U_k < \infty$ almost surely.

Algorithm	σ^2	Stepsize	Conv.	Rate	Iterate	Ref
SGD	$\neq 0$	$O(k^{-1/2-\epsilon})$	<i>a.s.</i>	$o(k^{-1/2+\epsilon})$	average	Cor. 9
SGD	$= 0$	$O(1)$	<i>a.s.</i>	$o(k^{-1})$	average	Cor. 9
SGD-ALS, SGD-PS	$\neq 0$	adaptive*	<i>a.s.</i> , \mathbb{E}	$o(k^{-1/2+\epsilon}), O(k^{-1/2+\epsilon})$	average	Cor. 12
SGD-ALS, SGD-PS	$= 0$	adaptive*	<i>a.s.</i> , \mathbb{E}	$o(k^{-1}), O(k^{-1})$	average	Cor. 12
SHB	$\neq 0$	$O(k^{-1/2+\epsilon}), O(k^{-1/2})$	<i>a.s.</i> , \mathbb{E}	$o(k^{-1/2-\epsilon}), O(k^{-1/2})$	last	Cor. 17, 25
SHB	$= 0$	$O(1)$	<i>a.s.</i> , \mathbb{E}	$o(k^{-1}), O(k^{-1})$	last	Cor. 14, 25
SGD, nonconvex	$C \neq 0$	$O(k^{-1/2+\epsilon})$	<i>a.s.</i>	$o(k^{-1/2+\epsilon})$	$\min \ \nabla\ ^2$	Cor. 18
SGD, nonconvex	$C = 0$	$O(1)$	<i>a.s.</i>	$o(k^{-1})$	$\min \ \nabla\ ^2$	Cor. 18

Table 1: Summary of the rates we obtain. All small-o (resp. big-O) rates are *almost surely* (resp. in expectation). The constants C , \mathcal{L} and σ are defined in (ABC), (5) and (4), respectively. *a.s.*: *almost surely*, \mathbb{E} : in expectation. $\min \|\nabla\|^2$: lowest squared norm of the gradient along the trajectory of SGD. adaptive*: Maximum step sizes need to verify conditions similar to Condition 1, but do not require knowing the smoothness constant \mathcal{L} .

We use the following condition on the step sizes in our *almost sure* convergence results.

Condition 1 *The sequence $(\eta_k)_k$ is decreasing, $\sum_k \eta_k = \infty$, $\sum_k \eta_k^2 \sigma^2 < \infty$ and $\sum_k \frac{\eta_k}{\sum_j \eta_j} = \infty$.*

The conditions $\sum_k \eta_k = \infty$ and $\sum_k \eta_k^2 < \infty$ are known as the Robbins-Monro conditions (Robbins and Monro, 1951) and are classical in the SGD literature (see Bertsekas and Tsitsiklis (2000) for example). The additional conditions, $\sum_k \frac{\eta_k}{\sum_j \eta_j} = \infty$ and $(\eta_k)_k$ is decreasing, allow us to derive convergence rates for the *almost sure* convergence using a new proof technique. However, as we will see in the next remark, the usual choices of step sizes which verify the Robbins-Monro conditions verify Condition 1 as well.

Remark 7 *Let $\eta_k = \frac{\eta}{k^\xi}$ with $\xi, \eta > 0$. Condition 1 is verified for all $\xi \in (\frac{1}{2}, 1]$ when $\sigma^2 \neq 0$, and for all $\xi \in [0, 1]$ when $\sigma^2 = 0$.*

See Appendix A for a proof of this remark. Indeed, all the formal proofs of our results are deferred to the appendix.

Theorem 8 *Let Assumption 1 hold. Consider the iterates of SGD. Choose step sizes $(\eta_k)_k$ which verify Condition 1, where $\forall k \in \mathbb{N}$, $0 < \eta_k \leq 1/(4\mathcal{L})$. Define for all $k \in \mathbb{N}$*

$$w_k = \frac{2\eta_k}{\sum_{j=0}^k \eta_j} \quad \text{and} \quad \begin{cases} \bar{x}_0 = x_0 \\ \bar{x}_{k+1} = w_k x_k + (1 - w_k) \bar{x}_k. \end{cases} \quad (7)$$

Then, we have a.s. that $f(\bar{x}_k) - f_ = o\left(\frac{1}{\sum_{t=0}^{k-1} \eta_t}\right)$.*

Proof We present the main elements of the proof which help in understanding the difference between the classical non-asymptotic analysis of SGD in expectation and our analysis. We present the complete proof in Section C of the appendix.

In the convex setting, the bulk of the convergence proofs of SGD is in using convexity and smoothness of f to establish that, if $\eta_k \leq \frac{1}{4\mathcal{L}}$, we have

$$\mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + \eta_k (f(x_k) - f_*) \leq \|x_k - x_*\|^2 + 2\eta_k^2 \sigma^2. \quad (8)$$

Classic non-asymptotic convergence analysis for SGD. Taking the expectation, using telescopic cancellation and Jensen’s inequality, it is possible to establish that

$$\mathbb{E} [f(\tilde{x}_k) - f_*] \leq \frac{\|x_0 - x_*\|^2}{\sum_{t=0}^{k-1} \eta_t} + \frac{2\sigma^2 \sum_{t=0}^{k-1} \eta_t^2}{\sum_{t=0}^{k-1} \eta_t}, \text{ where } \tilde{x}_k = \sum_{t=0}^{k-1} \frac{\eta_t}{\sum_{j=0}^{k-1} \eta_j} x_t.$$

\tilde{x}_k can then be computed on the fly using:

$$\tilde{x}_{k+1} = \tilde{w}_k x_k + (1 - \tilde{w}_k) \tilde{x}_k, \quad \text{where } \tilde{w}_k = \frac{\eta_k}{\sum_{j=0}^k \eta_j}. \quad (9)$$

This sequence of weights $(\tilde{w}_k)_k$ (which can be computed on the fly as $\tilde{w}_{k+1} = \frac{\eta_{k+1} \tilde{w}_k}{\eta_k + \eta_{k+1} \tilde{w}_k}$) is the one which allows to derive the tightest upper bound on the objective gap $f(x) - f_*$ in expectation. But it does not lend itself to tight *almost sure* asymptotic convergence, as we will show next.

Naive asymptotic analysis. Applying Lemma 6 to (8) gives that $\sum_k \eta_k (f(x_k) - f_*) < \infty$. Unfortunately, this only gives that $\lim_k \eta_k (f(x_k) - f_*) = 0$.

Asymptotic analysis using the iterates defined in (9). What if we had used the sequence of iterates defined in (9)? Let $\tilde{\delta}_k = f(\tilde{x}_k) - f_*$. Using Jensen’s inequality, we have

$$f(x_k) - f_* \geq \frac{1}{\tilde{w}_k} \tilde{\delta}_{k+1} - \left(\frac{1}{\tilde{w}_k} - 1 \right) \tilde{\delta}_k.$$

Using this bound in (8) gives, after replacing \tilde{w}_k by its expression (9) and multiplying by η_k , that

$$\mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + \sum_{j=0}^k \eta_j \tilde{\delta}_{k+1} \leq \|x_k - x_*\|^2 + \sum_{j=0}^{k-1} \eta_j \tilde{\delta}_k + 2\eta_k^2 \sigma^2.$$

Applying Lemma 6 gives that $\left(\sum_{j=0}^{k-1} \eta_j \tilde{\delta}_k \right)_k$ converges *almost surely*. This, together with the fact that $\sum_j \eta_j = \infty$, implies that $\tilde{\delta}_k = O\left(\frac{1}{\sum_{j=0}^{k-1} \eta_j}\right)$. This improves over the naive asymptotic analysis, but we can actually do better with a different approach.

Improved analysis. Now consider the alternative averaging of iterates \bar{x}_k given in (7). Let $\delta_k \stackrel{\text{def}}{=} f(\bar{x}_k) - f_*$. As we have done in the last paragraph, we can use Jensen’s inequality to lower-bound $f(x_k) - f_*$ in (8) (detailed derivations are given in Appendix C), and we obtain:

$$\mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + \frac{1}{2} \sum_{j=0}^k \eta_j \delta_{k+1} + \frac{\eta_k}{2} \delta_k \leq \|x_k - x_*\|^2 + \frac{1}{2} \sum_{j=0}^{k-1} \eta_j \delta_k + 2\eta_k^2 \sigma^2.$$

By Lemma 6, $\left(\sum_{j=0}^{k-1} \eta_j \delta_k \right)_k$ converges *almost surely*, and $\sum_k \eta_k \delta_k < \infty$, which implies that $\lim_k \eta_k \delta_k = 0$. But since $\sum_k \frac{\eta_k}{\sum_{j=0}^{k-1} \eta_j} = \infty$, we have the desired result: $\lim_k \sum_{j=0}^{k-1} \eta_j \delta_k = 0$.

Note that in the first iteration, $w_0 = 2$ and $\bar{x}_1 = x_0$, and we don’t use Jensen’s inequality. ■

With suitable choices of stepsizes, we can extract *almost sure* convergence rates for SGD, as we see in the next corollary. These choices and all the rates we derive are also summarized in Table 1. To the best of our knowledge, these are the first rates for the *almost sure* convergence of SGD in the convex setting.

Corollary 9 (Corollary of Theorem 8) *Let Assumption 1 hold. Let $0 < \eta \leq 1/4\mathcal{L}$ and $\epsilon > 0$.*

- if $\sigma^2 \neq 0$. Let $\eta_k = \frac{\eta}{k^{1/2+\epsilon}}$. Then

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k^{1/2-\epsilon}}\right).$$

- If $\sigma^2 = 0$. Let $\eta_k = \eta$. Then

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k}\right).$$

Although the *almost sure* convergence of SGD with favourable convergence rates only requires the step sizes to verify Condition 1, there are other popular methods to set the step sizes, such as Line-Search (Nocedal and Wright, 2006) or Polyak Stepsizes (Polyak, 1987), which do not require knowing the smoothness constant \mathcal{L} . A natural question is whether the result we have derived in Theorem 8 extends to these methods. We answer this question positively in the next section.

2.2. Convergence of Adaptive step size methods

We first present two adaptive step size selection methods and then present their convergence analysis.

Armijo Line-Search Step Size (ALS). We say that α is an Armijo line-search step-size at $x \in \mathbb{R}^d$ for the function g if, given constants $c, \alpha_{\max} > 0$, α is the *largest step size* in $(0, \alpha_{\max}]$ such that

$$g(x - \alpha \nabla g(x)) \leq g(x) - c\alpha \|\nabla g(x)\|_2^2, \quad (10)$$

which we denote by

$$\alpha \sim \mathbf{ALS}_{c, \alpha_{\max}}(g, x)$$

In practice, we use backtracking to find this α , where we start with a value α_{\max} and decrease it by a factor $\beta \in (0, 1)$ until (10) is verified.

Instead of using a pre-determined step size in SGD, we can choose at each iteration $\eta_k \sim \mathbf{ALS}_{c, \alpha_{\max}}(f_{v_k}, x_k)$ or $\eta_k \sim \mathbf{PS}_{c, \alpha_{\max}}(f_{v_k}, x_k)$. SGD with ALS or PS is known to converge sublinearly to a neighborhood of the minimum and to the minimum exactly if $\sigma^2 = 0$ (Vaswani et al., 2019b, 2020; Loizou et al., 2020). However, it is still not known whether these methods converge to the minimum when $\sigma^2 \neq 0$.

Let $(\eta_k^{\max})_k$ and (γ_k) be two strictly positive decreasing sequences. Consider the following modified SGD methods: at each iteration k , sample $v_k \sim \mathcal{D}$ and update

$$x_{k+1} = x_k - \eta_k \gamma_k \nabla f_{v_k}(x_k), \quad \text{where } \eta_k \sim \mathbf{ALS}_{c, \eta_k^{\max}}(f_{v_k}, x_k), \quad (\text{SGD-ALS})$$

$$x_{k+1} = x_k - \eta_k \gamma_k \nabla f_{v_k}(x_k), \quad \text{where } \eta_k \sim \mathbf{PS}_{c, \eta_k^{\max}}(f_{v_k}, x_k). \quad (\text{SGD-PS})$$

Assumption 2 *For all $v \sim \mathcal{D}$, f_v is lower bounded by $f_v^* > -\infty$ almost surely, and we define $\bar{\sigma}^2 \stackrel{\text{def}}{=} f_* - \mathbb{E}_v[f_v^*]$.*

Similar to our analysis of [SGD](#), we can derive *almost sure* convergence rates to the minimum for an average of the iterates. Remarkably, the analysis of the two methods [SGD-ALS](#) and [SGD-PS](#) can be unified.

Theorem 10 *Let Assumptions 1 and 2 hold. Consider the iterates of [SGD-ALS](#) and [SGD-PS](#). Choose $(\eta_k^{\max})_k$ and $(\gamma_k)_k$ such that $(\eta_k^{\max} \gamma_k)_k$ is decreasing, $\eta_k^{\max} \rightarrow 0$, $\sum_k \eta_k^{\max} \gamma_k = \infty$, $\sum_k \eta_k^{\max} \gamma_k^2 \sigma^2 < \infty$ and $\sum_k \frac{\eta_k^{\max} \gamma_k}{\sum_{j=0}^{k-1} \eta_j^{\max} \gamma_j} = \infty$, $c \geq \frac{1}{2}$ and $\gamma_k \leq c$. Define for all $k \in \mathbb{N}$*

$$w_k = \frac{2\eta_k^{\max} \gamma_k}{\sum_{j=0}^k \eta_j^{\max} \gamma_j} \quad \text{and} \quad \begin{cases} \bar{x}_0 = x_0 \\ \bar{x}_{k+1} = w_k x_k + (1 - w_k) \bar{x}_k. \end{cases} \quad (11)$$

Then, we have almost surely that $f(\bar{x}_k) - f_* = o\left(\frac{1}{\sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t}\right)$.

We also present upper bounds on the suboptimality for [SGD-ALS](#) and [SGD-PS](#) in expectation, from which we can derive convergence rates.

Theorem 11 *Let Assumptions 1 and 2 hold. Let (η_k^{\max}) and (γ_k) be strictly positive, decreasing sequences with $\gamma_k \leq c$, for all $k \in \mathbb{N}$ and $c \geq \frac{1}{2}$. Then the iterates of [SGD-ALS](#) and [SGD-PS](#) satisfy*

$$\mathbb{E}[f(\bar{x}_k) - f_*] \leq \frac{2ca_0 \|x_0 - x_*\|^2 + 4c \sum_{t=0}^{k-1} \gamma_t \eta_t^{\max} \left(\frac{\eta_t^{\max} \mathcal{L}}{2(1-c)} - 1\right)_+ \bar{\sigma}^2 + 2 \sum_{t=0}^{k-1} \gamma_t^2 \eta_t^{\max} \bar{\sigma}^2}{\sum_{t=0}^{k-1} \gamma_t \eta_t^{\max}}, \quad (12)$$

where $\bar{x}_k = \sum_{t=0}^{k-1} \frac{\eta_t^{\max} \gamma_t}{\sum_{j=0}^{k-1} \eta_j^{\max} \gamma_j} x_t$ and $a_0 = \max\left\{\frac{\eta_0^{\max} \mathcal{L}}{2(1-c)}, 1\right\}$.

We now give precise convergence rates derived from the two previous theorems, in the overparametrized as well as the non-overparametrized cases.

Corollary 12 (Corollary of Theorems 10 and 11) *Let $\epsilon, \eta, \gamma > 0$, with $\gamma \leq c$. If $\eta_k^{\max} = \eta k^{-\frac{4\epsilon}{3}}$ and $\gamma_k = \gamma k^{-\frac{1}{2} + \frac{\epsilon}{3}}$ then*

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k^{\frac{1}{2}-\epsilon}}\right) \text{ a.s.} \quad \text{and} \quad \mathbb{E}[f(\bar{x}_k) - f_*] = O\left(\frac{1}{k^{\frac{1}{2}-\epsilon}}\right). \quad (13)$$

If $\bar{\sigma}^2 = 0$. Then, setting $\eta_k^{\max} = \eta > 0$, $c = \frac{2}{3}$ and $\gamma_k = 1$, then for all $x_* \in \mathcal{X}_*$,

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k}\right) \text{ a.s.} \quad \text{and} \quad \mathbb{E}[f(\bar{x}_k) - f_*] \leq \frac{2 \max\left\{\frac{3\eta L_{\max}}{2}, 1\right\} \|x_0 - x_*\|^2}{\eta k}.$$

Notice from (12) and (13) that our analysis highlights a tradeoff between the asymptotic and the nonasymptotic convergence in expectation of [SGD-ALS](#) and [SGD-PS](#). Indeed, (13) predicts that the slower the convergence of $(\eta_k^{\max})_k$ towards 0 (as $\epsilon \rightarrow 0$), the better is the resulting asymptotic convergence rate. However, according to (12), if $(\eta_k^{\max})_k$ vanishes slowly, the second term on the right hand side of (13) vanishes slowly as well, which makes the bound in (12) looser.

Notice also that to be able to derive convergence rates in the non-overparametrized case from the previous theorem, we not only decrease the maximum step sizes, but also scale the adaptive step size η_k by multiplying it by a decreasing sequence γ_k .

3. Almost sure convergence rates for Stochastic Heavy Ball

The rates we derived for [SGD](#), [SGD-ALS](#) and [SGD-PS](#) in the previous section all hold at some weighted average of the iterates. Yet, in practice, it is the last iterate of SGD which is used. In contrast, we show that these rates hold for *the last iterate* of SHB, which is due to the online averaging inherent to SHB that we highlight in [Proposition 5](#). We present the first *almost sure* convergence rates for SHB, and also show that the deterministic HB converges at a $o(1/k)$ rate, which is asymptotically faster than the previously established $O(1/k)$ ([Ghadimi et al., 2015](#)).

We now present *almost sure* convergence rates for SHB

Theorem 13 *Let $x_{-1} = x_0$ and consider the iterates of [SHB-IMA](#). Let [Assumption 1](#) hold. Let η_k be a sequence of stepsizes which verifies [Condition 1](#) and $\forall k \in \mathbb{N}, 0 < \eta_k \leq 1/8\mathcal{L}$. If*

$$\lambda_0 = 0 \quad \text{and} \quad \lambda_k = \frac{\sum_{t=0}^{k-1} \eta_t}{4\eta_k} \quad \text{for all } k \in \mathbb{N}^*, \quad (14)$$

then we have almost surely that $x_k \xrightarrow[k \rightarrow +\infty]{} x_$ for some $x_* \in \mathcal{X}_*$, and $f(x_k) - f_* = o\left(\frac{1}{\sum_{t=0}^{k-1} \eta_t}\right)$.*

Note that when specialized to full gradients sampling, *i.e.* when we use the deterministic HB method, our results hold without the need for *almost sure* statements.

To the best of our knowledge, [Theorem 13](#) is the first result showing that the iterates of SHB converge to a minimizer assuming only smoothness and convexity. Note that this result is not directly comparable to [Gadat et al. \(2018\)](#), who study the more general nonconvex setting but use assumptions beyond smoothness.

In the general stochastic setting, [Theorem 13](#) shows that SHB enjoys the same *almost sure* convergence rates as SGD with averaging (See [Table 1](#)). However, an added benefit of SHB is that these rates hold for the last iterate, which conforms to what is done in practice.

Corollary 14 *Assume $\sigma^2 = 0$ and let $\eta_k = \eta < 1/4\mathcal{L}$ for all $k \in \mathbb{N}$. By [Theorem 13](#) we have*

$$\lim_k k (f(x_k) - f_*) = 0, \quad \text{almost surely.}$$

This corollary has fundamental implications in the deterministic and the stochastic case. In the stochastic case, it shows that when $\sigma^2 = 0$, [SHB-IMA](#) with a fixed step size converges at a $o(1/k)$ rate at the last iterate. In the deterministic case, $\sigma^2 = 0$ always holds, as at each iteration we use the true gradient $\nabla f(x_k)$, and we have $\nabla f(x_*) = 0$ for all $x_* \in \mathcal{X}_*$. Thus [Corollary 14](#) shows that the HB method enjoys the same $o(1/k)$ asymptotic convergence rate as gradient descent ([Lee and Wright, 2019](#)).

It seems that it is our choice iteration-dependent momentum coefficients given by [\(6\)](#) and [\(14\)](#) that enable this fast ‘small o’ convergence of the objective values for SHB. Recent work by [Attouch and Peypouquet \(2016\)](#) corroborates with this finding, where the authors also showed that a version of (deterministic) Nesterov’s Accelerated Gradient algorithm with carefully chosen iteration dependent momentum coefficients converges at a $o(1/k^2)$ rate, rather than the previously known $O(1/k^2)$.

4. Non-convex *almost sure* convergence rates for SGD

We now move on to the non-convex case, where we use the following assumption from [Khaled and Richtárik \(2020\)](#).

Assumption 3 *There exist constants $A, B, C \geq 0$ s.t. for all $x \in \mathbb{R}^d$,*

$$\mathbb{E}_v \left[\|\nabla f_v(x)\|^2 \right] \leq A(f(x) - f_*) + B\|\nabla f(x)\|^2 + C. \quad (\text{ABC})$$

This assumption is called *Expected Smoothness* in [Khaled and Richtárik \(2020\)](#). It includes the bounded gradients assumption, with $A = B = 0$ and $C = G > 0$, and the bounded gradient variance assumption, with $A = 0, B = 1$ and $C = \sigma^2$, as special cases. See ([Khaled and Richtárik, 2020](#), Th. 1) for a thorough investigation of the other assumptions used in the literature which are implied by (ABC). A major benefit of this assumption is that when f is a finite-sum (Remark 4) and the f_i functions are lower-bounded, (ABC) always holds ([Khaled and Richtárik, 2020](#), Prop. 3).

Remark 15 ([Khaled and Richtárik \(2020\)](#), Prop. 3) *In the setting of Remark 4, and assuming that for all $i \in [n]$, $f_i \geq f_i^* > -\infty$, Assumption (ABC) holds with:*

$$A = \frac{n-b}{b(n-b)}L_{\max}, \quad B = \frac{n(b-1)}{b(n-1)}, \quad \text{and} \quad C = \frac{2A}{n} \sum_{i=1}^n (f_* - f_i^*).$$

Since a global minimizer of f does not always exist in the nonconvex case, we can now only hope to find a stationary point. Hence, we present asymptotic convergence rates for the squared gradient norm.

Theorem 16 *Consider the iterates of SGD. Assume that (ABC) holds. Choose stepsizes which verify Condition 1 (with C in place of σ^2) such that $\forall k \in \mathbb{N}$, $0 < \eta_k \leq 1/(BL)$. Then, we have a.s that $\min_{t=0, \dots, k-1} \|\nabla f(x_t)\|^2 = o\left(\frac{1}{\sum_{t=0}^{k-1} \eta_t}\right)$.*

From this result, we can derive *almost sure* convergence rates arbitrarily close to $o(1/\sqrt{k})$, which can be improved to $o(1/k)$ in the overparametrized setting (See Table 1). Since these results are similar to Corollary 9, we omit them for brevity and report them in Table 1 and Corollary 18 in Appendix A.

5. Experiments

In our experiments, we aimed to examine whether or not [SHB-IMA](#) with the parameter settings suggested by our theory performed better than [SGD](#) and SGD with three common alternative parameter settings used throughout the machine learning literature: SGD with fixed momentum β of 0.9 and 0.99 as well as no momentum.

For our experiments, we selected a diverse set of multi-class classification problems from the LibSVM repository, 25 problems in total. These datasets range from a few classes to a thousand, and they vary from hundreds of data-points to hundreds of thousands. We normalized each dataset by a constant so that the largest data vector had norm 1. We used a multi-class logistic regression loss with no regularization so we could test the non-strongly convex convergence properties, and we ran for 50 epochs with no batching.

	SHB	SGD	Momentum 0.9	Momentum 0.99	No best method
Best method for	11	0	0	0	14

Table 2: Count of how many problems each method is statistically significantly superior to the rest on

We use SHB to denote the method (SHB) with α_k and β_k set using (6) (or equivalently the method (SHB-IMA)) and we left η , as well as the step sizes of all the methods we compare, as a constant to be determined through grid search. For the gridsearch, we used power-of-2 grid (2^i), we ran 5 random seeds and chose the learning rate that gave the lowest loss on average for each combination of problem and method. We widened the grid search as necessary for each combination to ensure that the chosen learning-rate was not from the endpoints of our grid search. Although it is possible to give a closed-form bound for the Lipschitz smoothness constant for our test problems, the above setting is less conservative and has the advantage of being usable without requiring any knowledge about the problem structure.

We then ran 40 different random seeds to produce Figure 1. To determine which method, if any, was best on each problem, we performed t-tests with Bonferroni correction, and we report how often each method was statistically significantly superior to all of the other three methods in Table 2. The stochastic heavy ball method using our theoretically motivated parameter settings performed better than all other methods on 11 of the 25 problems. On the remaining problems, no other method was statistically significantly better than all of the rest.

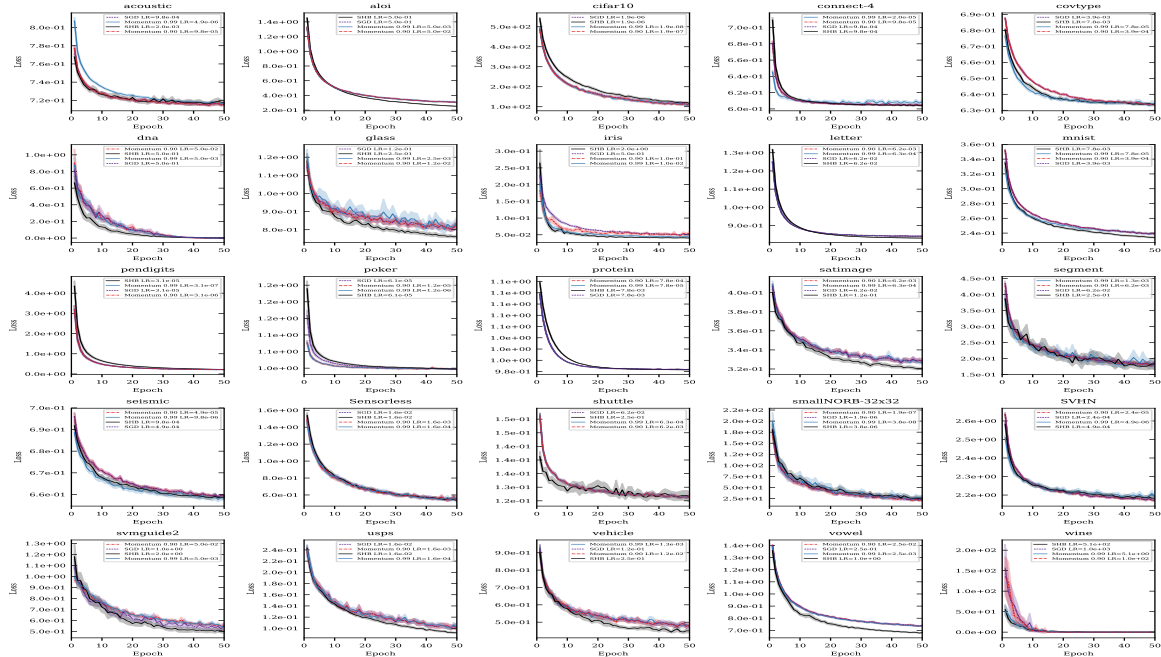


Figure 1: Average training error convergence plots for 25 LibSVM datasets, with using the best learning rate for each method and problem combination. Averages are over 40 runs. Error bars show a range of +/- 2SE.

Acknowledgements

We thank the anonymous reviewers who helped us improve this work.

The work of Othmane Sebbouh was supported in part by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR19- P3IA-0001 (PRAIRIE 3IA Institute). Othmane Sebbouh also acknowledges the support of a "Chaire d'excellence de l'IDEX Paris Saclay".

References

- Hédy Attouch and Juan Peypouquet. The rate of convergence of nesterov’s accelerated forward-backward method is actually faster than $1/k^2$. *SIAM Journal on Optimization*, 26(3):1824–1834, 2016.
- Francis Bach and Eric Moulines. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Neural Information Processing Systems (NIPS)*, 2011.
- Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer Publishing Company, Incorporated, 1st edition, 2011.
- Dimitri P Bertsekas and John N Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.
- Leon Bottou. Stochastic learning. In *Advanced Lectures on Machine Learning*, volume 3176, pages 146–168, 2003.
- Bugra Can, Mert Gürbüzbalaban, and Lingjiong Zhu. Accelerated linear convergence of stochastic momentum methods in wasserstein distances. In *Proceedings of the 36th International Conference on Machine Learning*, pages 891–901, 2019.
- Aaron Defazio. On the curved geometry of accelerated optimization. *Advances in Neural Information Processing Systems 33 (NIPS 2019)*, 2019.
- Sébastien Gadat, Fabien Panloup, and Sofiane Saadane. Stochastic heavy ball. *Electronic Journal of Statistics*, 12:461–529, 2018.
- Nidham Gazagnadou, Robert Mansel Gower, and Joseph Salmon. Optimal mini-batch and step sizes for saga. *The International Conference on Machine Learning*, 2019.
- Euhanna Ghadimi, Hamid Feysmahdavian, and Mikael Johansson. Global convergence of the heavy-ball method for convex optimization. In *2015 European Control Conference (ECC)*, pages 310–315, 2015.
- Saeed Ghadimi and Guanghui Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM J. Optimization*, 23(4):2341–2368, 2013.
- Antoine Godichon-Baggioni. Lp and almost sure rates of convergence of averaged stochastic gradient algorithms with applications to online robust estimation. *arXiv preprint arXiv:1609.05479*, 2016.
- Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. SGD: General Analysis and Improved Rates. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 5200–5209, 2019.
- Prateek Jain, Dheeraj Nagaraj, and Praneeth Netrapalli. Making the last iterate of sgd information theoretically optimal. In *Conference on Learning Theory, COLT 2019*, volume 99, pages 1752–1755, 2019.
- Ahmed Khaled and Peter Richtárik. Better theory for SGD in the nonconvex world. *arXiv:2002.03329*, 2020.

- Rahul Kidambi, Praneeth Netrapalli, Prateek Jain, and Sham M. Kakade. On the insufficiency of existing momentum schemes for stochastic optimization. In *International Conference on Learning Representations*, 2018.
- Guanghui Lan and Yi Zhou. An optimal randomized incremental gradient method. *Mathematical programming*, pages 1–49, 2017.
- Ching-Pei Lee and Stephen Wright. First-order algorithms converge faster than $o(1/k)$ on convex problems. volume 97, pages 3754–3762, 2019.
- Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS*, pages 983–992, 2019.
- Nicolas Loizou and Peter Richtárik. Momentum and Stochastic Momentum for Stochastic Gradient, Newton, Proximal Point and Subspace Descent Methods. *arXiv:1712.09677*, 2018.
- Nicolas Loizou, Sharan Vaswani, Issam Laradji, and Simon Lacoste-Julien. Stochastic polyak step-size for sgd: An adaptive learning rate for fast convergence. *arXiv preprint arXiv:2002.10542*, 2020.
- Panayotis Mertikopoulos, Nadav Hallak, Ali Kavis, and Volkan Cevher. On the almost sure convergence of stochastic gradient descent in non-convex problems. *arXiv preprint arXiv:2006.11144*, 2020.
- Arkadi Nemirovski, Anatoli B. Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4): 1574–1609, 2009.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. 2013.
- Lam M. Nguyen, Phuong Ha Nguyen, Marten van Dijk, Peter Richtárik, Katya Scheinberg, and Martin Takáč. SGD and hogwild! convergence without the bounded gradients assumption. In *Proceedings of the 35th International Conference on Machine Learning*, pages 3747–3755, 2018.
- Jorge Nocedal and Stephen J Wright. Sequential quadratic programming. *Numerical optimization*, pages 529–562, 2006.
- Francesco Orabona. Almost sure convergence of sgd on smooth non-convex functions. <https://parameterfree.com/2020/10/05/almost-sure-convergence-of-sgd-on-smooth-non-convex-functions/>, 2020a. Accessed: 2021-01-28.
- Francesco Orabona. Last iterate of sgd converges (even in bounded domains). <https://parameterfree.com/2020/08/07/last-iterate-of-sgd-converges-even-in-unbounded-domains/>, 2020b. Accessed: 2021-01-28.
- Antonio Orvieto, Jonas Köhler, and Aurélien Lucchi. The role of memory in stochastic optimization. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI*, page 128, 2019.

- B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4:1–17, 1964.
- BT Polyak. Introduction to optimization, translations series in mathematics and engineering. *Optimization Software*, 1987.
- H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.
- Herbert Robbins and David Siegmund. A convergence theorem for nonnegative almost supermartingales and some applications. *Optimizing methods in Statistics*, pages 233–257, 1971.
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, page III–1139–III–1147, 2013.
- Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of SGD for over-parameterized models and an accelerated perceptron. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89, pages 1195–1204, 2019a.
- Sharan Vaswani, Aaron Mishkin, Issam Laradji, Mark Schmidt, Gauthier Gidel, and Simon Lacoste-Julien. Painless stochastic gradient: Interpolation, line-search, and convergence rates. In *Advances in Neural Information Processing Systems 32*, pages 3727–3740. 2019b.
- Sharan Vaswani, Frederik Kunstner, Issam Laradji, Si Yi Meng, Mark Schmidt, and Simon Lacoste-Julien. Adaptive gradient methods converge faster with over-parameterization (and you can do a line-search). *arXiv preprint arXiv:2006.06835*, 2020.
- Tianbao Yang, Qihang Lin, and Zhe Li. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. *arXiv:1604.03257*, 2016.
- Zhengyuan Zhou, Panayotis Mertikopoulos, Nicholas Bambos, Stephen P. Boyd, and Peter W. Glynn. Stochastic mirror descent in variationally coherent optimization problems. In *Advances in Neural Information Processing Systems 30*, pages 7040–7049, 2017.

APPENDIX

Contents

1	Introduction	1
1.1	Contributions and Background	2
1.2	Assumptions and general consequences	3
1.3	SGD and an iterate-averaging viewpoint of SHB	5
2	Almost sure convergence rates for SGD and SGD with adaptive stepsizes	5
2.1	SGD: average-iterates <i>almost sure</i> convergence	5
2.2	Convergence of Adaptive step size methods	8
3	Almost sure convergence rates for Stochastic Heavy Ball	10
4	Non-convex <i>almost sure</i> convergence rates for SGD	11
5	Experiments	11
A	Proofs of corollaries on convergence rates and stepsize choices	19
A.1	Proof of remark 7	19
A.2	SGD: Proof of Corollary 9	19
A.3	SGD with adaptive step sizes: Proof of Corollary 12	20
A.4	SHB	21
A.5	SGD, nonconvex	21
B	Proofs for Section 1	22
B.1	Proof of Lemma 2	22
B.2	Proof of Proposition 5	22
C	Proofs for Section 2	23
C.1	Proof of Theorem 8	23
C.2	Proofs of Theorems 10 and 11	25
C.2.1	Proof of Theorem 10	26
C.2.2	Proof of Theorem 11	27
D	Proofs for Section 3	27
D.1	Proof of Theorem 13	27
E	Proofs for Section 4	32
E.1	Proof of Theorem 16	32
F	Extension of our results to the nonsmooth setting	33

G	Convergence rates for SHB in expectation without the bounded gradients and bounded gradient variance assumptions	35
G.1	Proof of Theorem 24	36
G.2	Proof of Corollary 25	37

The appendix is organized as follows:

- In Section A we present the proofs of Remark 7 and the stepsize choices and the corresponding convergence rates derived in the corollaries listed in Table 1.
- In Section B we present proofs for Section 1.
- In Section C we present proofs for Section 2.
- In Section D we present proofs for Section 3.
- In Section E we present proofs for Section 4.
- In Section F, we present our results for the convergence of stochastic subgradient descent under the bounded gradients assumptions.
- In Section G, we present convergence rates for SHB in expectation without the bounded gradients and bounded gradient variance assumptions.

A. Proofs of corollaries on convergence rates and stepsize choices

A.1. Proof of remark 7

Proof Let $\eta_k = \frac{\eta}{k^\xi}$ with $\eta > 0$ and $\xi \in [0, 1)$. Clearly, $(\eta_k)_k$ is decreasing and $\sum_k \eta_k = \infty$. And we have

$$\sum_{t=0}^{k-1} \eta_t \sim \eta k^{1-\xi}.$$

Hence,

$$\frac{\eta_k}{\sum_{t=0}^{k-1} \eta_t} \sim \frac{1}{k}.$$

Hence, $\sum_k \frac{\eta_k}{\sum_{t=0}^{k-1} \eta_t} = \infty$.

- If $\sigma^2 \neq 0$. Let $\xi \in (\frac{1}{2}, 1]$. Then $\sum_k \eta_k^2 < \infty$, and the stepsizes verify Condition 1.
- If $\sigma^2 = 0$. Let $\xi \in [0, 1)$. We have $\sum_k \eta_k^2 \sigma^2 = 0 < \infty$. Hence, the stepsizes verify Condition 1. ■

A.2. SGD: Proof of Corollary 9

Proof

- If $\sigma^2 \neq 0$. Let $\eta_k = \frac{\eta}{k^{1/2+\epsilon}}$. From Remark 7, we have that the stepsizes verify Condition 1. Moreover, $\sum_k \frac{1}{k^{\frac{1}{2}+\epsilon}} \sim k^{-1/2+\epsilon}$. Thus, from Theorem 8:

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k^{1/2-\epsilon}}\right).$$

- If $\sigma^2 = 0$. Let $\eta_k = \eta$. From Remark 7, the stepsizes verify Condition 1. Thus, from Theorem 8:

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k}\right).$$

■

A.3. SGD with adaptive step sizes: Proof of Corollary 12

Proof Let $\eta, \gamma, c > 0$. We first prove the *almost sure* convergence results.

- If $\bar{\sigma}^2 \neq 0$. Let $\eta_k^{\max} = \eta k^{-4\epsilon/3}$ and $\gamma_k = \gamma k^{-\frac{1}{2} + \epsilon/3}$. Clearly, $(\eta_k^{\max} \gamma_k)_k$ is decreasing, $\sum_k \eta_k = \infty$ and $\sum_k \eta_k^{\max} \gamma_k^2 < \infty$. And

$$\sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t \sim \eta \gamma k^{1/2 - \epsilon}.$$

Hence,

$$\frac{\eta_k^{\max} \gamma_k}{\sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t} \sim \frac{1}{k}.$$

Hence, $\sum_k \frac{\eta_k^{\max} \gamma_k}{\sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t} = \infty$. Thus, the stepsizes verify the conditions of Theorem 10, and we have

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k^{1/2 - \epsilon}}\right).$$

- If $\bar{\sigma}^2 = 0$. Let $\eta_k^{\max} = \eta$ and $\gamma_k = 1$. Clearly $(\eta_k^{\max} \gamma_k)_k$ is decreasing since it is constant, $\sum_k \eta_k^{\max} \gamma_k = \infty$, $\sum_k \eta_k^{\max} \gamma_k^2 \bar{\sigma}^2 = 0 < \infty$, and $\sum_k \frac{\eta_k^{\max} \gamma_k}{\sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t} = \sum_k 1 = \infty$. Thus, the stepsizes verify the conditions of Theorem 10, and we have

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{k}\right).$$

We now prove the convergence rates in expectation. Remember that from Theorem 11, we have that if (η_k^{\max}) and (γ_k) are strictly positive, decreasing sequences with $\gamma_k \leq c$ for all $k \in \mathbb{N}$ and $c \geq \frac{1}{2}$, then the iterates of **SGD-ALS** and **SGD-PS** satisfy

$$\mathbb{E} [f(\bar{x}_k) - f_*] \leq \frac{2ca_0 \|x_0 - x_*\|^2 + 4c \sum_{t=0}^{k-1} \gamma_t \eta_t^{\max} \left(\frac{\eta_t^{\max} \mathcal{L}}{2(1-c)} - 1 \right)_+ \bar{\sigma}^2 + 2 \sum_{t=0}^{k-1} \gamma_t^2 \eta_t^{\max} \bar{\sigma}^2}{\sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t}, \quad (15)$$

where $\bar{x}_k = \sum_{t=0}^{k-1} \frac{\eta_t^{\max} \gamma_t}{\sum_{j=0}^{k-1} \eta_j^{\max} \gamma_j} x_t$ and $a_0 = \max \left\{ \frac{\eta_0^{\max} \mathcal{L}}{2(1-c)}, 1 \right\}$.

- If $\bar{\sigma}^2 \neq 0$. Let $\eta_k^{\max} = \eta k^{-4\epsilon/3}$ and $\gamma_k = \gamma k^{-\frac{1}{2} + \epsilon/3}$. Since $\eta_t^{\max} \rightarrow 0$, there exists $k_0 \in \mathbb{N}$ such that for all $t \geq k_0$, $\gamma_t \eta_t^{\max} \left(\frac{\eta_t^{\max} \mathcal{L}}{2(1-c)} - 1 \right)_+ = 0$. Hence, for all $k \geq k_0$

$$\mathbb{E} [f(\bar{x}_k) - f_*] \leq \frac{2ca_0 \|x_0 - x_*\|^2 + 4c \sum_{t=0}^{k_0-1} \gamma_t \eta_t^{\max} \left(\frac{\eta_t^{\max} \mathcal{L}}{2(1-c)} - 1 \right)_+ \bar{\sigma}^2 + 2 \sum_{t=0}^{k-1} \gamma_t^2 \eta_t^{\max} \bar{\sigma}^2}{\sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t}.$$

Replacing η_k^{\max} and γ_k with their values gives

$$\mathbb{E} [f(\bar{x}_k) - f_*] = O\left(\frac{1}{k^{1/2-\epsilon}}\right).$$

- If $\bar{\sigma}^2 = 0$. Let $\eta_k^{\max} = \eta > 0$ and $\gamma = 1$. We have from (15) that for all $k \in \mathbb{N}$,

$$\mathbb{E} [f(\bar{x}_k) - f_*] \leq \frac{2 \max\left\{\frac{3\eta\mathcal{L}}{2}, 1\right\} \|x_0 - x_*\|^2}{\eta k}.$$

■

A.4. SHB

Corollary 17 (Corollary of Theorem 13) *Let Assumption 1 hold. Let $0 < \eta \leq 1/4\mathcal{L}$ and $\epsilon > 0$.*

- if $\sigma^2 \neq 0$. Let $\eta_k = \frac{\eta}{k^{1/2+\epsilon}}$. Then

$$f(x_k) - f_* = o\left(\frac{1}{k^{1/2-\epsilon}}\right).$$

- If $\sigma^2 = 0$. Let $\eta_k = \eta$. Then

$$f(x_k) - f_* = o\left(\frac{1}{k}\right).$$

Proof The proof is the same as the proof of Corollary 9, using Theorem 13 instead of Theorem 8. ■

A.5. SGD, nonconvex

Corollary 18 (Corollary of Theorem 16) *Let Assumption (ABC) hold. Let $0 < \eta \leq 1/4\mathcal{L}$ and $\epsilon > 0$.*

- If $\sigma^2 = 0$. Let $\eta_k = \frac{\eta}{k^{1/2+\epsilon}}$. Then,

$$\min_{t=0, \dots, k-1} \|\nabla f(x_k)\|^2 = o\left(\frac{1}{k^{1/2-\epsilon}}\right).$$

- If $\sigma^2 \neq 0$. Let $\eta_k = \eta$. Then,

$$\min_{t=0, \dots, k-1} \|\nabla f(x_k)\|^2 = o\left(\frac{1}{k}\right).$$

Proof

- If $C \neq 0$. Let $\eta_k = \frac{\eta}{k^{1/2+\epsilon}}$. From Remark 7, we have that the stepsizes verify Condition 1 with C in place of σ^2 . Moreover, $\sum_k \frac{1}{k^{1/2+\epsilon}} \sim k^{-1/2+\epsilon}$. Thus, from Theorem 16:

$$\min_{t=0,\dots,k-1} \|\nabla f(x_k)\|^2 = o\left(\frac{1}{k^{1/2-\epsilon}}\right).$$

- If $C = 0$. Let $\eta_k = \eta$. From Remark 7, the stepsizes verify Condition 1 with C in place of σ^2 . Thus, from Theorem 16:

$$\min_{t=0,\dots,k-1} \|\nabla f(x_k)\|^2 = o\left(\frac{1}{k}\right).$$

■

B. Proofs for Section 1
B.1. Proof of Lemma 2

Proof Since for all $v \sim \mathcal{D}$, f_v is convex and L_v smooth, we have from (Nesterov, 2013, Equation 2.1.7) that

$$\begin{aligned} \|\nabla f_v(x) - \nabla f_v(x_*)\|^2 &\leq 2L_v (f_v(x) - f_v(x_*) - \langle \nabla f_v(x_*), x - x_* \rangle) \\ &\stackrel{\text{Asm. 1}}{\leq} 2\mathcal{L} (f_v(x) - f_v(x_*) - \langle \nabla f_v(x_*), x - x_* \rangle) \end{aligned}$$

Hence,

$$\mathbb{E}_{v \sim \mathcal{D}} \left[\|\nabla f_v(x) - \nabla f_v(x_*)\|^2 \right] \leq 2\mathcal{L} (f(x) - f(x_*)).$$

Therefore,

$$\begin{aligned} \mathbb{E}_{v \sim \mathcal{D}} \left[\|\nabla f_v(x)\|^2 \right] &\leq 2\mathbb{E}_{v \sim \mathcal{D}} \left[\|\nabla f_v(x) - \nabla f_v(x_*)\|^2 \right] + 2\mathbb{E}_{v \sim \mathcal{D}} \left[\|\nabla f_v(x_*)\|^2 \right] \\ &\leq 4\mathcal{L} (f(x) - f(x_*)) + 2\mathbb{E}_{v \sim \mathcal{D}} \left[\|\nabla f_v(x_*)\|^2 \right]. \end{aligned}$$

■

B.2. Proof of Proposition 5

Proof Consider the iterate-averaging method

$$z_{k+1} = z_k - \eta_k \nabla f_{v_k}(x_k), \tag{16}$$

$$x_{k+1} = \frac{\lambda_{k+1}}{\lambda_{k+1} + 1} x_k + \frac{1}{\lambda_{k+1} + 1} z_{k+1}, \tag{17}$$

and let

$$\alpha_k = \frac{\eta_k}{\lambda_{k+1} + 1} \quad \text{and} \quad \beta_k = \frac{\lambda_k}{\lambda_{k+1} + 1}. \quad (18)$$

Substituting (16) into (17) gives

$$x_{k+1} = \frac{\lambda_{k+1}}{\lambda_{k+1} + 1} x_k + \frac{1}{\lambda_{k+1} + 1} (z_k - \eta_k \nabla f_{v_k}(x_k)). \quad (19)$$

Now using (17) at the previous iteration we have that that

$$z_k = (\lambda_k + 1) \left(x_k - \frac{\lambda_k}{\lambda_k + 1} x_{k-1} \right) = (\lambda_k + 1)x_k - \lambda_k x_{k-1}.$$

Substituting the above into (19) gives

$$\begin{aligned} x_{k+1} &= \frac{\lambda_{k+1}}{\lambda_{k+1} + 1} x_k + \frac{1}{\lambda_{k+1} + 1} ((\lambda_k + 1)x_k - \lambda_k x_{k-1} - \eta_k \nabla f_{v_k}(x_k)) \\ &= x_k - \frac{\eta_k}{\lambda_{k+1} + 1} \nabla f_{v_k}(x_k) + \frac{\lambda_k}{\lambda_{k+1} + 1} (x_k - x_{k-1}). \end{aligned}$$

Consequently, (18) gives the desired expression. ■

C. Proofs for Section 2

C.1. Proof of Theorem 8

Proof Consider the setting of Theorem 8. Expanding the squares we have that

$$\|x_{k+1} - x_*\|^2 = \|x_k - x_*\|^2 - 2\eta_k \langle \nabla f_{v_k}(x_k), x_k - x_* \rangle + \eta_k^2 \|\nabla f_{v_k}(x_k)\|^2.$$

Then taking conditional expectation $\mathbb{E}_k[\cdot] \stackrel{\text{def}}{=} \mathbb{E}[\cdot | x_k]$ gives

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] &= \|x_k - x_*\|^2 - 2\eta_k \langle \nabla f(x_k), x_k - x_* \rangle + \eta_k^2 \mathbb{E}_k \left[\|\nabla f_{v_k}(x_k)\|^2 \right] \\ &\stackrel{(2)+(5)}{\leq} \|x_k - x_*\|^2 - 2\eta_k (1 - 2\eta_k \mathcal{L}) (f(x_k) - f_*) + 2\eta_k^2 \sigma^2. \end{aligned}$$

Since $\eta_k \leq \frac{1}{4\mathcal{L}}$, we have that $1 - 2\eta_k \mathcal{L} \geq \frac{1}{2}$. Hence, rearranging, we have

$$\mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + \eta_k (f(x_k) - f_*) \leq \|x_k - x_*\|^2 + 2\eta_k^2 \sigma^2. \quad (20)$$

From (7) we have that $w_k = \frac{2\eta_k}{\sum_{j=0}^k \eta_j}$. Since $w_0 = \frac{2\eta_0}{\eta_0} = 2$ we have that $\bar{x}_1 = 2x_0 - \bar{x}_0 = 2x_0 - x_0 = x_0$. Hence, it holds that

$$f(\bar{x}_1) - f_* = f(x_0) - f_* = w_0 (f(x_0) - f_*) + (1 - w_0) (f(\bar{x}_0) - f_*). \quad (21)$$

Now for $k \in \mathbb{N}^*$ we have that following equivalence

$$w_k \in [0, 1] \iff 2\eta_k \leq \sum_{j=0}^k \eta_j \iff \eta_k \leq \sum_{j=0}^{k-1} \eta_j.$$

The right hand side of the equivalence holds because $(\eta_k)_k$ is a decreasing sequence. Hence, by Jensen's inequality, we have $\forall k \in \mathbb{N}^*$,

$$f(\bar{x}_{k+1}) - f_* \leq w_k (f(x_k) - f_*) + (1 - w_k) (f(\bar{x}_k) - f_*).$$

Together with (21), this shows that the last inequality holds for all $k \in \mathbb{N}$. Thus,

$$\eta_k (f(x_k) - f_*) \geq \frac{\eta_k}{w_k} (f(\bar{x}_{k+1}) - f_*) - \eta_k \left(\frac{1}{w_k} - 1 \right) (f(\bar{x}_k) - f_*).$$

Replacing this expression in (20) gives:

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] &+ \frac{\eta_k}{w_k} (f(\bar{x}_{k+1}) - f_*) \\ &\leq \|x_k - x_*\|^2 + \eta_k \left(\frac{1}{w_k} - 1 \right) (f(\bar{x}_k) - f_*) + 2\eta_k^2 \sigma^2. \end{aligned}$$

Hence substituting in the definition of w_k from (7) gives

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] &+ \frac{1}{2} \sum_{j=0}^k \eta_j (f(\bar{x}_{k+1}) - f_*) \\ &\leq \|x_k - x_*\|^2 + \frac{1}{2} \left(\sum_{j=0}^{k-1} \eta_j - \eta_k \right) (f(\bar{x}_k) - f_*) + 2\eta_k^2 \sigma^2. \end{aligned}$$

Thus re-arranging

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] &+ \frac{1}{2} \sum_{j=0}^k \eta_j (f(\bar{x}_{k+1}) - f_*) + \frac{\eta_k}{2} (f(\bar{x}_k) - f_*) \\ &\leq \|x_k - x_*\|^2 + \frac{1}{2} \sum_{j=0}^{k-1} \eta_j (f(\bar{x}_k) - f_*) + 2\eta_k^2 \sigma^2, \end{aligned}$$

which, by Lemma 6, has the three following consequences:

$$\begin{aligned} &(\|x_k - x_*\|^2)_k \text{ and } \left(\sum_{j=0}^k \eta_j (f(\bar{x}_{k+1}) - f_*) \right)_k \text{ converge almost surely,} \\ &\text{and } \sum_k \eta_k (f(\bar{x}_k) - f_*) < \infty. \end{aligned}$$

Hence, we have that $\lim_k \frac{\eta_k}{\sum_{j=0}^{k-1} \eta_j} \sum_{j=0}^{k-1} \eta_j (f(\bar{x}_k) - f_*) = \lim_k \eta_k (f(\bar{x}_k) - f_*) = 0$. But since we assumed that $\sum_k \frac{\eta_k}{\sum_{j=0}^{k-1} \eta_j}$ diverges, this implies that $\lim_k \sum_{j=0}^{k-1} \eta_j (f(\bar{x}_{k+1}) - f_*) = 0$, that is

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{\sum_{j=0}^{k-1} \eta_j}\right)$$

■

C.2. Proofs of Theorems 10 and 11

The results of Theorems 10 and 11 can be derived as corollaries of the following theorem.

Theorem 19 *Let (η_k^{\max}) and (γ_k) be strictly positive, decreasing sequences with $\gamma_k \leq \frac{3c}{2}$ for all $k \in \mathbb{N}$ and $c \geq \frac{2}{3}$. Then the iterates of (SGD-ALS) and the iterates of (SGD-PS) satisfy*

$$2a_{k+1} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + \eta_k^{\max} \gamma_k (f(x_k) - f_*) \leq 2a_k \|x_k - x_*\|^2 + 4\gamma_k \eta_k^{\max} \left(\frac{\eta_k^{\max} \mathcal{L}}{2(1-c)} - 1 \right)_+ \bar{\sigma}^2 + \frac{2\gamma_k^2 \eta_k^{\max} \bar{\sigma}^2}{c}, \quad (22)$$

where $\bar{\sigma}^2 \stackrel{\text{def}}{=} f_* - \mathbb{E}_v[f_v^*]$, $a_k = \max\left\{\frac{\eta_k^{\max} \mathcal{L}}{2(1-c)}, 1\right\}$ and where $(b)_+ \stackrel{\text{def}}{=} \max(b, 0)$ for all $b \in \mathbb{R}$.

Before proving Theorem 19, we need the following lemma:

Lemma 20 (Vaswani et al. (2019b), Loizou et al. (2020)) *Let g be an L_g -smooth function, $\alpha_{\max} > 0$ and $c \in [0, 1]$. If $\alpha \sim \mathbf{SALS}_{c, \alpha_{\max}}(g, x)$ or $\mathbf{SP}_{c, \alpha_{\max}}(g, x)$, then*

$$\min\left\{\frac{2(1-c)}{L_g}, \alpha_{\max}\right\} \leq \alpha \leq \alpha_{\max} \quad \text{and} \quad \alpha \|\nabla g(x)\|^2 \leq \frac{g(x) - g^*}{c \|\nabla g(x)\|^2}. \quad (23)$$

Proof [Proof of Theorem 19] Let us now prove Theorem 19.

$$\begin{aligned} \|x_{k+1} - x_*\|^2 &= \|x_k - x_*\|^2 - 2\eta_k \gamma_k \langle \nabla f_{v_k}(x_k), x_k - x_* \rangle + \eta_k^2 \gamma_k^2 \|\nabla f_{v_k}(x_k)\|^2 \\ &\stackrel{(2)}{\leq} \|x_k - x_*\|^2 - 2\eta_k \gamma_k (f_{v_k}(x_k) - f_{v_k}^*) + \eta_k^2 \gamma_k^2 \|\nabla f_{v_k}(x_k)\|^2 \\ &\stackrel{(23)}{\leq} \|x_k - x_*\|^2 - 2\eta_k \gamma_k (f_{v_k}(x_k) - f_{v_k}(x_*)) + \frac{\eta_k \gamma_k^2}{c} (f_{v_k}(x_k) - f_{v_k}^*) \\ &= \|x_k - x_*\|^2 - 2\eta_k \gamma_k \left(1 - \frac{\gamma_k}{2c}\right) (f_{v_k}(x_k) - f_{v_k}^*) + 2\eta_k \gamma_k (f_{v_k}(x_*) - f_{v_k}^*). \end{aligned}$$

Rearranging, we have

$$2\eta_k \gamma_k \left(1 - \frac{\gamma_k}{2c}\right) (f_{v_k}(x_k) - f_{v_k}^*) \leq \|x_k - x_*\|^2 - \|x_{k+1} - x_*\|^2 + 2\eta_k \gamma_k (f_{v_k}(x_*) - f_{v_k}^*).$$

Define $\eta_k^{\min} \stackrel{\text{def}}{=} \min\left\{\frac{2(1-c)}{\mathcal{L}}, \eta_k^{\max}\right\}$. Then,

$$2\eta_k^{\min} \gamma_k \left(1 - \frac{\gamma_k}{2c}\right) (f_{v_k}(x_k) - f_{v_k}^*) \stackrel{(23)}{\leq} \|x_k - x_*\|^2 - \|x_{k+1} - x_*\|^2 + 2\eta_k \gamma_k (f_{v_k}(x_*) - f_{v_k}^*).$$

Hence,

$$\begin{aligned}
 2\eta_k^{\min}\gamma_k\left(1 - \frac{\gamma_k}{2c}\right)(f_{v_k}(x_k) - f_{v_k}(x_*)) &\leq \|x_k - x_*\|^2 - \|x_{k+1} - x_*\|^2 \\
 &+ 2\gamma_k\left(\eta_k - \left(1 - \frac{\gamma_k}{2c}\right)\eta_k^{\min}\right)(f_{v_k}(x_*) - f_{v_k}^*) \\
 &\leq \|x_k - x_*\|^2 - \|x_{k+1} - x_*\|^2 \\
 &+ 2\gamma_k\left(\eta_k^{\max} - \left(1 - \frac{\gamma_k}{2c}\right)\eta_k^{\min}\right)(f_{v_k}(x_*) - f_{v_k}^*).
 \end{aligned} \tag{24}$$

Notice that

$$\frac{\eta_k^{\max}}{\eta_k^{\min}} = \max\left\{\frac{\eta_k^{\max}\mathcal{L}}{2(1-c)}, 1\right\}.$$

Since $(\eta_k^{\max})_k$ is decreasing, $\left(\frac{\eta_k^{\max}}{\eta_k^{\min}}\right)_k$ is decreasing as well. Hence, multiplying both sides of (24) by $\frac{\eta_k^{\max}}{\eta_k^{\min}}$,

$$\begin{aligned}
 2\eta_k^{\max}\gamma_k\left(1 - \frac{\gamma_k}{2c}\right)(f_{v_k}(x_k) - f_{v_k}(x_*)) &\leq \frac{\eta_k^{\max}}{\eta_k^{\min}}\|x_k - x_*\|^2 - \frac{\eta_{k+1}^{\max}}{\eta_{k+1}^{\min}}\|x_{k+1} - x_*\|^2 \\
 &+ 2\gamma_k\eta_k^{\max}\left(\frac{\eta_k^{\max}}{\eta_k^{\min}} - 1 + \frac{\gamma_k}{2c}\right)(f_{v_k}(x_*) - f_{v_k}^*).
 \end{aligned}$$

Hence, taking the expectation,

$$\begin{aligned}
 2\eta_k^{\max}\gamma_k\left(1 - \frac{\gamma_k}{2c}\right)(f(x_k) - f_*) &\leq \frac{\eta_k^{\max}}{\eta_k^{\min}}\|x_k - x_*\|^2 - \frac{\eta_{k+1}^{\max}}{\eta_{k+1}^{\min}}\mathbb{E}_k\left[\|x_{k+1} - x_*\|^2\right] \\
 &+ 2\gamma_k\eta_k^{\max}\left(\frac{\eta_k^{\max}}{\eta_k^{\min}} - 1 + \frac{\gamma_k}{2c}\right)\bar{\sigma}^2.
 \end{aligned}$$

where $\bar{\sigma}^2 \stackrel{\text{def}}{=} f_* - \mathbb{E}_v[f_v^*]$. Using the fact that $1 - \frac{\gamma_k}{2c} \geq \frac{1}{4}$ and rearranging, we have

$$\begin{aligned}
 \frac{2\eta_{k+1}^{\max}}{\eta_{k+1}^{\min}}\mathbb{E}_k\left[\|x_{k+1} - x_*\|^2\right] + \eta_k^{\max}\gamma_k(f(x_k) - f_*) &\leq \frac{2\eta_k^{\max}}{\eta_k^{\min}}\|x_k - x_*\|^2 \\
 &+ 4\gamma_k\eta_k^{\max}\left(\frac{\eta_k^{\max}\mathcal{L}}{2(1-c)} - 1\right)_+ \bar{\sigma}^2 + \frac{2\gamma_k^2\eta_k^{\max}\bar{\sigma}^2}{c},
 \end{aligned}$$

where for all $a \in \mathbb{R}$, $(a)_+ = \max(a, 0)$. ■

C.2.1. PROOF OF THEOREM 10

Proof Using the inequality (22) from Theorem 19, the proof of Theorem 10 proceeds exactly as the proof of Theorem 8, with the conditions on the stepsizes of Theorem 10 instead of Condition 1. See Section 2 and C. ■

C.2.2. PROOF OF THEOREM 11

Proof Taking the expectation in (22), rearranging and summing between $t = 0, \dots, k-1$, we have,

$$\begin{aligned} \sum_{t=0}^{k-1} \eta_t^{\max} \gamma_t \mathbb{E} [f(x_t) - f_*] &\leq 2a_0 \|x_0 - x_*\|^2 - 2a_{k+1} \mathbb{E} [\|x_{k+1} - x_*\|^2] \\ &+ 4 \sum_{t=0}^{k-1} \gamma_t \eta_t^{\max} \left(\frac{\eta_t^{\max} \mathcal{L}}{2(1-c)} - 1 \right)_+ \bar{\sigma}^2 + \frac{2 \sum_{t=0}^{k-1} \gamma_t^2 \eta_t^{\max} \bar{\sigma}^2}{c}. \end{aligned}$$

Dividing by $\sum_{j=0}^{k-1} \eta_t^{\max} \gamma_t$ and using Jensen's inequality gives the desired result. \blacksquare

D. Proofs for Section 3

D.1. Proof of Theorem 13

In the remainder of this section and the forthcoming lemmas we consider the iterates of (SHB) and the setting of Theorem 13, that is

$$\lambda_0 = 0, \lambda_k = \lambda_k = \frac{\sum_{t=0}^{k-1} \eta_t}{4\eta_k}, \alpha_k = \frac{\eta_k}{1 + \lambda_{k+1}} \text{ and } \beta_k = \frac{\lambda_k}{1 + \lambda_{k+1}}, \quad (25)$$

where

$$0 < \eta_k \leq 1/4\mathcal{L}, \sum_k \eta_k^2 \sigma^2 < \infty \text{ and } \sum_k \eta_k = \infty. \quad (26)$$

Note that from (SHB-IMA), we have

$$z_k = x_k + \lambda_k (x_k - x_{k-1}). \quad (27)$$

We also assume that Assumption 1 holds throughout.

To make the proof more readable, we first state and prove the two following lemmas.

Lemma 21 $\sum_k \eta_k (f(x_k) - f_*) < +\infty$ almost surely.

Lemma 22 $\sum_k \lambda_{k+1} \|x_k - x_{k-1}\|^2 < +\infty$, and thus, $\lim_k \lambda_{k+1} \|x_{k+1} - x_k\|^2 = 0$ almost surely.

We first prove Lemma 21.

Proof [Proof of Lemma 21] From (43), we have

$$\begin{aligned} \mathbb{E}_k [\|z_{k+1} - x_*\|^2] &\leq \|z_k - x_*\|^2 - 2\eta_k \left(\frac{1}{2} + \lambda_k \right) (f(x_k) - f_*) \\ &+ 2\eta_k \lambda_k (f(x_{k-1}) - f_*) + 2\eta_k^2 \sigma^2. \end{aligned} \quad (28)$$

Using (14) we have that

$$\begin{aligned}
 2\eta_k \left(\frac{1}{2} + \lambda_k \right) &= 2\eta_k \left(\frac{2\eta_k + \sum_{t=0}^{k-1} \eta_t}{4\eta_k} \right) \\
 &= 2\eta_k \left(\frac{\eta_k}{4\eta_k} + \frac{\sum_{t=0}^{k-1} \eta_t}{4\eta_k} \right) \\
 &= 2\eta_k \left(\frac{1}{4} + \frac{\sum_{t=0}^{k-1} \eta_t \eta_{k+1}}{4\eta_{k+1} \eta_k} \right) \\
 &= \frac{\eta_k}{2} + 2\eta_{k+1} \lambda_{k+1}.
 \end{aligned} \tag{29}$$

Using (29) in (28) gives

$$\begin{aligned}
 \mathbb{E}_k \left[\|z_{k+1} - x_*\|^2 + 2\eta_{k+1} \lambda_{k+1} (f(x_k) - f_*) \right] + \eta_k (f(x_k) - f_*) \\
 \leq \|z_k - x_*\|^2 + 2\eta_k \lambda_k (f(x_{k-1}) - f_*) + 2\eta_k^2 \sigma^2.
 \end{aligned}$$

Hence, applying Lemma 6 with

$$V_k = \|z_k - x_*\|^2 + 2\eta_k \lambda_k (f(x_{k-1}) - f_*), \quad \gamma_k = 0, \quad U_{k+1} = \eta_k (f(x_k) - f_*) \quad \text{and} \quad Z_k = 2\eta_k^2 \sigma^2,$$

we have by (26) that

$$\sum_k \eta_k (f(x_k) - f_*) < +\infty \text{ almost surely.}$$

■

We now turn to prove Lemma 22.

Proof [Proof of Lemma 22] We have,

$$\mathbb{E}_k \left[\|x_{k+1} - x_k\|^2 \right] \stackrel{\text{(SHB)}}{=} \beta_k^2 \|x_k - x_{k-1}\|^2 + \alpha_k^2 \|\nabla f_{v_k}(x_k)\|^2 - 2\alpha_k \beta_k \langle \nabla f(x_k), x_k - x_{k-1} \rangle.$$

Multiplying by $(1 + \lambda_{k+1})^2$ and using (25) we have that $\beta_k = \frac{\lambda_k}{1 + \lambda_{k+1}}$ and $\alpha_k = \frac{\eta_k}{1 + \lambda_{k+1}}$ and thus

$$(1 + \lambda_{k+1})^2 \mathbb{E}_k \left[\|x_{k+1} - x_k\|^2 \right] = \lambda_k^2 \|x_k - x_{k-1}\|^2 + \eta_k^2 \|\nabla f_{v_k}(x_k)\|^2 - 2\eta_k \lambda_k \langle \nabla f(x_k), x_k - x_{k-1} \rangle.$$

Thus using the convexity of f and (5), which follows from Lemma 2, we have

$$\begin{aligned}
 (1 + \lambda_{k+1})^2 \mathbb{E}_k \left[\|x_{k+1} - x_k\|^2 \right] &\leq \lambda_k^2 \|x_k - x_{k-1}\|^2 + 4\eta_k^2 \mathcal{L} (f(x_k) - f_*) + 2\eta_k \lambda_k (f(x_{k-1}) - f(x_k)) + 2\eta_k^2 \sigma^2 \\
 &= \lambda_k^2 \|x_k - x_{k-1}\|^2 - 2\eta_k (\lambda_k - 2\eta_k \mathcal{L}) (f(x_k) - f_*) \\
 &\quad + 2\eta_k \lambda_k (f(x_{k-1}) - f_*) + 2\eta_k^2 \sigma^2.
 \end{aligned}$$

Re-arranging the above gives,

$$\begin{aligned}
 (1 + \lambda_{k+1})^2 \mathbb{E}_k \left[\|x_{k+1} - x_k\|^2 \right] + 2\eta_k (\lambda_k - 2\eta_k \mathcal{L}) (f(x_k) - f_*) \\
 \leq \lambda_k^2 \|x_k - x_{k-1}\|^2 + 2\eta_k \lambda_k (f(x_{k-1}) - f_*) + 2\eta_k^2 \sigma^2.
 \end{aligned} \tag{30}$$

Combining both (30) and (28) we have that

$$\begin{aligned} \mathbb{E}_k \left[\|z_{k+1} - x_*\|^2 \right] &+ 4\eta_k \left(\frac{1}{2} - 2\eta_k \mathcal{L} + \lambda_k \right) (f(x_k) - f_*) + (1 + \lambda_{k+1})^2 \mathbb{E}_k \left[\|x_{k+1} - x_k\|^2 \right] \\ &\leq \|z_k - x_*\|^2 + 4\eta_k \lambda_k (f(x_{k-1}) - f_*) + \lambda_k^2 \|x_k - x_{k-1}\|^2 + 4\eta_k^2 \sigma^2. \end{aligned}$$

Hence, since $\eta_k \leq \frac{1}{8\mathcal{L}}$,

$$\begin{aligned} \mathbb{E}_k \left[\|z_{k+1} - x_*\|^2 \right] &+ 4\eta_k \left(\frac{1}{4} + \lambda_k \right) (f(x_k) - f_*) + (1 + \lambda_{k+1})^2 \mathbb{E}_k \left[\|x_{k+1} - x_k\|^2 \right] \\ &\leq \|z_k - x_*\|^2 + 4\eta_k \lambda_k (f(x_{k-1}) - f_*) + \lambda_k^2 \|x_k - x_{k-1}\|^2 + 4\eta_k^2 \sigma^2. \end{aligned}$$

Using (25), we have $\eta_k \left(\frac{1}{4} + \lambda_k \right) = \eta_{k+1} \lambda_{k+1}$. Hence,

$$\begin{aligned} \|z_{k+1} - x_*\|^2 &+ 4\eta_{k+1} \lambda_{k+1} (f(x_k) - f_*) + (1 + \lambda_{k+1})^2 \|x_{k+1} - x_k\|^2 \\ &\leq \|z_k - x_*\|^2 + 4\eta_k \lambda_k (f(x_{k-1}) - f_*) + \lambda_k^2 \|x_k - x_{k-1}\|^2 + 4\eta_k^2 \sigma^2. \end{aligned}$$

Hence, noting $V_k \stackrel{\text{def}}{=} \|z_k - x_*\|^2 + 4\eta_k \lambda_k (f(x_{k-1}) - f_*) + \lambda_k^2 \|x_k - x_{k-1}\|^2$, we have

$$\mathbb{E}_k [V_{k+1}] + (2\lambda_{k+1} + 1) \|x_k - x_{k-1}\|^2 \leq V_k + 4\eta_k^2 \sigma^2.$$

Hence, since $\sum_k \eta_k^2 \sigma^2 < +\infty$, applying Lemma 6, we have

$$\sum_k \lambda_{k+1} \|x_k - x_{k-1}\|^2 < +\infty \quad \text{almost surely, thus} \quad \lim_k \lambda_{k+1} \|x_{k+1} - x_k\|^2 = 0 \quad \text{almost surely.}$$

■

We can now prove Theorem 13.

Proof [Proof of Theorem 13] This proof aims at proving that, *almost surely*

1. $x_k \xrightarrow[k \rightarrow +\infty]{} x_*$ for some $x_* \in \mathcal{X}_*$.

2. $f(x_k) - f_* = o\left(\frac{1}{\sum_{t=0}^{k-1} \eta_t}\right)$.

In our road to prove the first point, we will prove the second point as a byproduct.

We will now prove that $\lim_k \|z_k - x_*\|^2$ exists *almost surely*.

$$\begin{aligned} \|z_k - x_*\|^2 &\stackrel{(27)}{=} \|x_k - x_* + \lambda_k (x_k - x_{k-1})\|^2 \\ &= \lambda_k^2 \|x_k - x_{k-1}\|^2 + 2\lambda_k \langle x_k - x_*, x_k - x_{k-1} \rangle + \|x_k - x_*\|^2 \\ &= (\lambda_k^2 + \lambda_k) \|x_k - x_{k-1}\|^2 + \lambda_k \left(\|x_k - x_*\|^2 - \|x_{k-1} - x_*\|^2 \right) + \|x_k - x_*\|^2. \end{aligned}$$

Define

$$\delta_k \stackrel{\text{def}}{=} \lambda_k \left(\|x_k - x_*\|^2 - \|x_{k-1} - x_*\|^2 \right) + \|x_k - x_*\|^2, \quad (31)$$

so that

$$\|z_k - x_*\|^2 = (\lambda_k^2 + \lambda_k) \|x_k - x_{k-1}\|^2 + \delta_k.$$

We will first prove that $\lim_k (\lambda_k^2 + \lambda_k) \|x_k - x_{k-1}\|^2$ exists *almost surely*, then that $\lim_k \delta_k$ exists *almost surely*.

First, we have from Lemma 22 that $(\lambda_k \|x_k - x_{k-1}\|^2)_k$ converges to zero *almost surely*. Hence, it remains to show that $\lim_k \lambda_k^2 \|x_k - x_{k-1}\|^2$ exists *almost surely*. From (30), we have that

$$\begin{aligned} & \lambda_{k+1}^2 \mathbb{E}_k \left[\|x_{k+1} - x_k\|^2 \right] + 2\eta_k (\lambda_k - 2\eta_k \mathcal{L}) (f(x_k) - f_*) \\ & \leq \lambda_k^2 \|x_k - x_{k-1}\|^2 + 2\eta_k \lambda_k (f(x_{k-1}) - f_*) + 2\eta_k^2 \sigma^2. \end{aligned}$$

Using (25) and the fact that $\eta_k \leq \frac{1}{8\mathcal{L}}$, we have that $2\eta_{k+1} \lambda_{k+1} = 2\eta_k (\frac{1}{4} + \lambda_k) \leq 2\eta_k (\frac{1}{2} - 2\eta_k \mathcal{L} + \lambda_k)$. Hence, $2\eta_k (\lambda_k - 2\eta_k \mathcal{L}) \geq 2\eta_{k+1} \lambda_{k+1} - \eta_k$. Therefore, denoting

$$d_k \stackrel{\text{def}}{=} \|x_k - x_{k-1}\|^2 \text{ and } \theta_k \stackrel{\text{def}}{=} 2\eta_k (f(x_{k-1}) - f_*),$$

we have

$$\mathbb{E}_k [\lambda_{k+1}^2 d_k + \lambda_{k+1} \theta_{k+1}] \leq \lambda_k^2 d_k + \lambda_k \theta_k + \eta_k (f(x_k) - f_*) + 2\eta_k^2 \sigma^2.$$

From Lemma 21, we have $\sum_k \eta_k (f(x_k) - f_*) < +\infty$. Moreover, $\sum_k \eta_k^2 \sigma^2 < +\infty$. Hence, we have by Lemma 6 that $\lim_k \lambda_k^2 d_k + \lambda_k \theta_k$ exists *almost surely*.

Moreover, by Lemma 22, $\sum_k \lambda_k d_k < +\infty$, and we have $\sum_k \theta_k < +\infty$ *almost surely*. Hence, $\sum_k \lambda_k d_k + \theta_k < +\infty$ *a.s.* Rewriting

$$\lambda_k d_k + \theta_k = \frac{1}{\lambda_k} (\lambda_k^2 d_k + \lambda_k \theta_k),$$

we have, since $\lim_k \lambda_k^2 d_k + \lambda_k \theta_k$ exists *almost surely* and

$$\sum_k \frac{1}{\lambda_k} \stackrel{0 < \eta_k \leq \frac{1}{4\mathcal{L}}}{=} 4 \sum_k \frac{\eta_k}{\sum_{t=0}^{k-1} \eta_t} \stackrel{\text{Condition 1}}{=} \infty$$

that

$$\lim_k \lambda_k^2 d_k + \lambda_k \theta_k = 0,$$

which means that both $\lim_k \lambda_k^2 d_k = 0$ and $\lim_k \lambda_k \theta_k = 0$ *a.s.* Writing out $\lim_k \lambda_k \theta_k = 0$ explicitly, we have

$$\boxed{f(x_k) - f_* = o\left(\frac{1}{\sum_{t=0}^{k-1} \eta_t}\right) \text{ almost surely}}$$

This proves the second point of Theorem 13 and that $\lim_k \lambda_k^2 d_k = 0$ *almost surely*. To prove the first point of Theorem 13 it remains to show that $\lim_k \delta_k$ exists *almost surely*.

Note $u_k = \|x_k - x_*\|^2$. We have

$$u_{k+1} \stackrel{\text{(SHB)}}{=} \|x_k - x_* + \beta_k (x_k - x_{k-1})\|^2 + \alpha_k^2 \|\nabla f_{v_k}(x_k)\|^2 - 2\alpha_k \beta_k \langle \nabla f_{v_k}(x_k), x_k - x_* \rangle - 2\alpha_k \langle \nabla f_{v_k}(x_k), x_k - x_* \rangle. \quad (32)$$

Taking expectation conditioned on x_k , and using the convexity of f and Lemma 2 we have that

$$\begin{aligned} \mathbb{E}_k [u_{k+1}] &\leq \|x_k - x_* + \beta_k (x_k - x_{k-1})\|^2 - 2\alpha_k (1 + \beta_k - 2\alpha_k \mathcal{L}) (f(x_k) - f_*) \\ &\quad + 2\alpha_k \beta_k (f(x_{k-1}) - f_*) + 2\alpha_k^2 \sigma^2. \end{aligned} \quad (33)$$

Furthermore note that

$$\begin{aligned} \|x_k - x_* + \beta_k (x_k - x_{k-1})\|^2 &= u_k + \beta_k^2 d_k + 2\beta_k \langle x_k - x_*, x_k - x_{k-1} \rangle \\ &= u_k + (\beta_k^2 + \beta_k) d_k + \beta_k (u_k - u_{k-1}). \end{aligned} \quad (34)$$

Hence, using the fact that $0 \leq \beta_k \leq 1$ and inserting (34) into (33) gives

$$\begin{aligned} \mathbb{E}_k [u_{k+1}] &\leq u_k + 2d_k + \beta_k (u_k - u_{k-1}) - 2\alpha_k (1 + \beta_k - 2\alpha_k \mathcal{L}) (f(x_k) - f_*) \\ &\quad + 2\alpha_k \beta_k (f(x_{k-1}) - f_*) + 2\alpha_k^2 \sigma^2. \end{aligned}$$

Multiplying the above by $(1 + \lambda_{k+1})$, rearranging and using (6) results in

$$\begin{aligned} (1 + \lambda_{k+1}) \mathbb{E}_k [u_{k+1} - u_k] &\leq 2(1 + \lambda_{k+1}) d_k + \lambda_k (u_k - u_{k-1}) - 2\eta_k (1 + \beta_k - 2\alpha_k \mathcal{L}) (f(x_k) - f_*) \\ &\quad + 2\eta_k \beta_k (f(x_{k-1}) - f_*) + 2 \frac{\eta_k^2 \sigma^2}{1 + \lambda_{k+1}}. \end{aligned} \quad (35)$$

Using the definition of δ_{k+1} given in (31) we have that

$$\delta_{k+1} - \delta_k = (1 + \lambda_{k+1}) (u_{k+1} - u_k) - \lambda_k (u_k - u_{k-1}),$$

which we use to re-write (35) as we have

$$\delta_{k+1} + \frac{\eta_k}{\eta_{k+1}} (1 + \beta_k - 2\alpha_k \mathcal{L}) \theta_{k+1} \leq \delta_k + 2(1 + \lambda_{k+1}) d_k + \beta_k \theta_k + 2 \frac{\eta_k^2 \sigma^2}{1 + \lambda_{k+1}}.$$

Hence, since $\eta_{k+1} \leq \eta_k$,

$$\mathbb{E}_k [\delta_{k+1} + (1 + \beta_k - 2\alpha_k \mathcal{L}) \theta_{k+1}] \leq \delta_k + \beta_k \theta_k + 2(1 + \lambda_{k+1}) d_k + 2 \frac{\eta_k^2}{1 + \lambda_{k+1}} \sigma^2.$$

And since,

$$\begin{aligned} 1 + \beta_k - 2\alpha_k \mathcal{L} &= 1 + \frac{\lambda_k}{1 + \lambda_{k+1}} - \frac{2\eta_k \mathcal{L}}{1 + \lambda_{k+1}} = \frac{1}{1 + \lambda_{k+1}} (1 + \lambda_{k+1} + \lambda_k - 2\eta_k \mathcal{L}) \\ &\geq \frac{\lambda_{k+1}}{1 + \lambda_{k+1}} \geq \frac{\lambda_{k+1}}{1 + \lambda_{k+2}} = \beta_{k+1}, \end{aligned}$$

we have

$$\mathbb{E}_k [\delta_{k+1} + \beta_{k+1}\theta_{k+1}] \leq (\delta_k + \beta_k\theta_k) + 2(1 + \lambda_{k+1})d_k + \frac{2\eta_k^2}{1 + \lambda_{k+1}}\sigma^2.$$

Since by Lemma 22 we have that $\sum_k 2(1 + \lambda_{k+1})d_k < +\infty$ *almost surely*, and $\sum_k \frac{\eta_k^2\sigma^2}{1 + \lambda_{k+1}} < +\infty$, we have by Lemma 6 that $\lim_k \delta_k + \beta_k\theta_k$ exists *almost surely*. And since $\lim_k \beta_k\theta_k = 0$ *almost surely*, we deduce that $\lim_k \delta_k$ exists *almost surely*.

Thus we have now shown that $\lim_k \|z_k - x_*\|^2$ exists *almost surely*. Therefore, since $x_k - x_* = z_k - x_* - \lambda_k(x_k - x_{k-1})$ and

$$\| \|x_k - x_*\| - \|z_k - x_*\| \| \leq \lambda_k \|x_k - x_{k-1}\| \xrightarrow[k \rightarrow +\infty]{} 0 \quad \text{almost surely,}$$

we have that $\lim_k \|x_k - x_*\| - \|z_k - x_*\|$ exists *almost surely*, and so does $\lim_k \|x_k - x_*\|$.

We also have that both $\|z_k - x_*\|$ and $\lambda_k \|x_k - x_{k-1}\|$ are bounded *almost surely*, thus $\|x_k - x_*\|$ is bounded *almost surely*. Hence, $(x_k)_k$ is bounded *almost surely*, thus *almost surely* sequentially compact.

Let $(x_{n_k})_k$ be a subsequence of $(x_n)_n$ which converges to some $x \in \mathbb{R}^d$ *a.s.* Since $f(x_n) \rightarrow_n f_*$ *almost surely* for all $x^* \in \text{argmin } f$, we have $x \in \text{argmin } f$ *a.s.* Finally, applying Lemma 2.39 in Bauschke and Combettes (2011) (restricted to our finite dimensional setting, where weak convergence and strong convergence are equivalent), there exists $x_* \in \text{argmin } f$ such that

$$\boxed{x_k \xrightarrow[k \rightarrow +\infty]{} x_*, \quad \text{almost surely}}$$

This proves the first point of Theorem 13. ■

E. Proofs for Section 4

E.1. Proof of Theorem 16

Proof Consider the setting of Theorem 16. Let $\eta_k \leq \frac{1}{LB}$ for all $k \in \mathbb{N}$. From (Khaled and Richtárik, 2020, Proof of Lemma 2), we have

$$\mathbb{E}_k [f(x_{k+1}) - f_*] + \eta_k \|\nabla f(x_k)\|^2 \leq (1 + \eta_k^2 AL) (f(x_k) - f_*) + \frac{\eta_k^2 LC}{2}.$$

Define for all $k \in \mathbb{N}$,

$$w_k = \frac{2\eta_k}{\sum_{j=0}^k \eta_j}, \quad g_0 = \|\nabla f(x_0)\|^2, \quad g_{k+1} = (1 - w_k)g_k + w_k \|\nabla f(x_k)\|^2.$$

Note that since $(\eta_k)_k$ is decreasing, $w_k \in [0, 1]$. Plugging this back in the previous inequality gives

$$\mathbb{E}_k [f(x_{k+1}) - f_*] + \frac{\sum_{j=0}^k \eta_j}{2} g_{k+1} + \frac{\eta_k}{2} g_k \leq (1 + \eta_k^2 AL) (f(x_k) - f_*) + \frac{\sum_{j=0}^{k-1} \eta_j}{2} g_k + \frac{\eta_k^2 LC}{2}.$$

Since $\sum_k \eta_k^2 < \infty$, we also have that $\prod_{k=0}^{\infty} (1 + \eta_k^2 AL) < \infty$. Thus, by Lemma 6, we have

$$(f(x_k) - f_*)_k \text{ and } \left(\left(\sum_{j=0}^{k-1} \eta_j \right) g_k \right)_k \text{ converge almost surely, and } \sum_k \eta_k g_k < \infty \text{ almost surely}$$

Hence, we have that $\lim_k \frac{\eta_k}{\sum_{j=0}^{k-1} \eta_j} \sum_{j=0}^{k-1} \eta_j g_k = \lim_k \eta_k g_k = 0$. But since we assumed that $\sum_k \frac{\eta_k}{\sum_{j=0}^{k-1} \eta_j}$ diverges, this implies that $\lim_k \sum_{j=0}^{k-1} \eta_j g_k = 0$, that is, we have that,

$$g_k = o\left(\frac{1}{\sum_{j=0}^{k-1} \eta_j}\right) \text{ almost surely}$$

But since for all $k \in \mathbb{N}$, $g_{k+1} = (1 - w_k)g_k + w_k \|\nabla f(x_k)\|^2$, g_k is a weighted average of all past $\|\nabla f(x_j)\|^2$, $j = 0, \dots, k-1$. Hence, there exists a sequence $(\tilde{w}_j)_j$ in $[0, 1]$ which verifies $\sum_{j=0}^{k-1} \tilde{w}_j = 1$ such that $g_k = \sum_{j=0}^{k-1} \tilde{w}_j \|\nabla f(x_j)\|^2$. Thus, $g_k \geq \min_{t=0, \dots, k-1} \|\nabla f(x_t)\|^2 \geq 0$. Hence we have *almost surely*

$$\boxed{\min_{t=0, \dots, k-1} \|\nabla f(x_t)\|^2 = o\left(\frac{1}{\sum_{j=0}^{k-1} \eta_j}\right)}$$

■

F. Extension of our results to the nonsmooth setting

In this section, we will consider the stochastic subgradient descent method under the bounded gradients assumption, as in Nemirovski et al. (2009). Under this assumption, we show that we can derive the same convergence rates as in Theorem 8.

Proposition 23 *Consider the following method: at each iteration k , let g_k be such that $\mathbb{E}_k [g_k] = g(x_k)$ for some $g(x_k) \in \partial f(x_k)$, and update*

$$x_{k+1} = x_k - \eta_k g_k,$$

where we assume that f is convex and that there exists G such that $\forall k \in \mathbb{N}$, $\mathbb{E} [\|g_k\|^2] \leq G$.

Choose step sizes $(\eta_k)_k$ which verify Condition 1 (with G in place of σ^2). Define for all $k \in \mathbb{N}$

$$w_k = \frac{2\eta_k}{\sum_{j=0}^k \eta_j} \quad \text{and} \quad \begin{cases} \bar{x}_0 = x_0 \\ \bar{x}_{k+1} = w_k x_k + (1 - w_k) \bar{x}_k. \end{cases} \quad (36)$$

Then, we have a.s. that $f(\bar{x}_k) - f_* = o\left(\frac{1}{\sum_{t=0}^{k-1} \eta_t}\right)$.

Proof The proof proceeds exactly as in the smooth case, but with replacing the bound (5) by the bound $\forall k \in \mathbb{N}$, $\mathbb{E} \left[\|g_k\|^2 \right] \leq G$. Indeed, expanding the squares we have that

$$\|x_{k+1} - x_*\|^2 = \|x_k - x_*\|^2 - 2\eta_k \langle g_k, x_k - x_* \rangle + \eta_k^2 \|g_k\|^2.$$

Then taking conditional expectation $\mathbb{E}_k[\cdot] \stackrel{\text{def}}{=} \mathbb{E}[\cdot | x_k]$ gives, since $\mathbb{E}_k[g_k] = g(x_k)$ for some $g(x_k) \in \partial f(x_k)$, we have

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] &= \|x_k - x_*\|^2 - 2\eta_k \langle g(x_k), x_k - x_* \rangle + \eta_k^2 \mathbb{E}_k \left[\|g_k\|^2 \right] \\ &\leq \|x_k - x_*\|^2 - 2\eta_k (f(x_k) - f_*) + \eta_k^2 G, \end{aligned}$$

where we used in the last inequality the fact that $g(x_k)$ is a subgradient of f at x_k , and that $\mathbb{E}_k \left[\|g_k\|^2 \right] \leq G$. Hence, rearranging, we have

$$\mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + 2\eta_k (f(x_k) - f_*) \leq \|x_k - x_*\|^2 + \eta_k^2 G. \quad (37)$$

From (36) we have that $w_k = \frac{2\eta_k}{\sum_{j=0}^k \eta_j}$. Since $w_0 = 2\frac{\eta_0}{\eta_0} = 2$ we have that $\bar{x}_1 = 2x_0 - \bar{x}_0 = 2x_0 - x_0 = x_0$. Hence, it holds that

$$f(\bar{x}_1) - f_* = f(x_0) - f_* = w_0 (f(x_0) - f_*) + (1 - w_0) (f(\bar{x}_0) - f_*). \quad (38)$$

Now for $k \in \mathbb{N}^*$ we have that following equivalence

$$w_k \in [0, 1] \quad \Longleftrightarrow \quad 2\eta_k \leq \sum_{j=0}^k \eta_j \quad \Longleftrightarrow \quad \eta_k \leq \sum_{j=0}^{k-1} \eta_j.$$

The right hand side of the equivalence holds because $(\eta_k)_k$ is a decreasing sequence. Hence, by Jensen's inequality, we have $\forall k \in \mathbb{N}^*$,

$$f(\bar{x}_{k+1}) - f_* \leq w_k (f(x_k) - f_*) + (1 - w_k) (f(\bar{x}_k) - f_*).$$

Together with (38), this shows that the last inequality holds for all $k \in \mathbb{N}$. Thus,

$$\eta_k (f(x_k) - f_*) \geq \frac{\eta_k}{w_k} (f(\bar{x}_{k+1}) - f_*) - \eta_k \left(\frac{1}{w_k} - 1 \right) (f(\bar{x}_k) - f_*).$$

Replacing this expression in (37) gives:

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + \frac{\eta_k}{w_k} (f(\bar{x}_{k+1}) - f_*) \\ \leq \|x_k - x_*\|^2 + \eta_k \left(\frac{1}{w_k} - 1 \right) (f(\bar{x}_k) - f_*) + \eta_k^2 G. \end{aligned}$$

Hence substituting in the definition of w_k from (36) gives

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] + \sum_{j=0}^k \eta_j (f(\bar{x}_{k+1}) - f_*) \\ \leq \|x_k - x_*\|^2 + \left(\sum_{j=0}^{k-1} \eta_j - \eta_k \right) (f(\bar{x}_k) - f_*) + \eta_k^2 G. \end{aligned}$$

Thus re-arranging

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] &+ \sum_{j=0}^k \eta_j (f(\bar{x}_{k+1}) - f_*) + \eta_k (f(\bar{x}_k) - f_*) \\ &\leq \|x_k - x_*\|^2 + \sum_{j=0}^{k-1} \eta_j (f(\bar{x}_k) - f_*) + \eta_k^2 G, \end{aligned}$$

which, by Lemma 6, has the three following consequences:

$$\begin{aligned} (\|x_k - x_*\|^2)_k \text{ and } \left(\sum_{j=0}^k \eta_j (f(\bar{x}_{k+1}) - f_*) \right)_k &\text{ converge almost surely,} \\ \text{and } \sum_k \eta_k (f(\bar{x}_k) - f_*) &< \infty. \end{aligned}$$

Hence, we have that $\lim_k \frac{\eta_k}{\sum_{j=0}^{k-1} \eta_j} \sum_{j=0}^{k-1} \eta_j (f(\bar{x}_k) - f_*) = \lim_k \eta_k (f(\bar{x}_k) - f_*) = 0$. But since we assumed that $\sum_k \frac{\eta_k}{\sum_{j=0}^{k-1} \eta_j}$ diverges, this implies that $\lim_k \sum_{j=0}^{k-1} \eta_j (f(\bar{x}_{k+1}) - f_*) = 0$, that is

$$f(\bar{x}_k) - f_* = o\left(\frac{1}{\sum_{j=0}^{k-1} \eta_j}\right)$$

■

G. Convergence rates for SHB in expectation without the bounded gradients and bounded gradient variance assumptions

Our first theorem provides a non-asymptotic upper bound on the suboptimality given any sequence of step sizes. Later we develop special cases of this theorem through different choices of the stepsizes.

Theorem 24 *Let Assumption 1 hold. Let $x_{-1} = x_0$. Consider the iterates of SHB-IMA. Let $(\eta_k)_k$ be such that $0 < \eta_k \leq \frac{1}{4\mathcal{L}}$ for all $k \in \mathbb{N}$. Define $\lambda_0 \stackrel{\text{def}}{=} 0$ and $\lambda_k = \frac{\sum_{t=0}^{k-1} \eta_t}{2\eta_k}$ for $k \geq 1$. Then,*

$$\mathbb{E} [f(x_k) - f_*] \leq \frac{\|x_0 - x_*\|^2}{\sum_{t=0}^k \eta_t} + 2\sigma^2 \frac{\sum_{t=0}^k \eta_t^2}{\sum_{t=0}^k \eta_t}.$$

Note that in Theorem 24 the only free parameters are the η_k 's which in the iterate-moving-average viewpoint (SHB-IMA) play the role of a learning rate. The scaled step sizes α_k and the momentum parameters β_k of the usual formulation (SHB) are given by (6) once we have chosen η_k . We now explore three different settings of the η_k 's in the following corollaries.

Corollary 25 *Consider the setting of Theorem 24. Let $\eta \leq 1/4\mathcal{L}$.*

1. *Let $\eta_k = \eta$. Then, $\mathbb{E} [f(x_k) - f_*] \leq \frac{\|x_0 - x_*\|^2}{\eta(k+1)} + 2\eta\sigma^2$.* (39)

2. Let $\eta_k = \frac{\eta}{\sqrt{k+1}}$. Then, $\mathbb{E} [f(x_k) - f_*] \leq \frac{\|x^0 - x^*\|_2^2 + 4\sigma^2\eta^2(\log(k+1)+1)}{2\eta(\sqrt{k+1}-1)} \sim O\left(\frac{\log(k)}{\sqrt{k}}\right)$. (40)

3. Suppose Algorithm (SHB) is run for T iterations. Let $\eta_k = \frac{\eta}{\sqrt{T+1}}$ for all $k \in \{0, \dots, T\}$. Then,

$$\mathbb{E} [f(x_T) - f_*] \leq \frac{\|x^0 - x^*\|_2^2 + 2\sigma^2\eta^2}{\eta\sqrt{T+1}}. \quad (41)$$

(39) shows how to set the parameters of SHB so that the last iterate converges sublinearly to a neighborhood of the minimum. In particular, for overparametrized models with $\sigma^2 = 0$, the last iterate of SHB converges sublinearly to the minimum. Moreover, when using the full gradient, which corresponds to directly using the gradient $\nabla f(x_k)$ at each iteration, we have $\mathcal{L} = L$ and $\sigma^2 = 0$, which recovers the rate derived in Ghadimi et al. (2015) for the deterministic HB method upto a constant.

The $O\left(\log(k)/\sqrt{k}\right)$ convergence rate in (40) is the same rate that can be derived for the iterates of SGD, as is done by Nemirovski et al. (2009) for a weighted average of the iterates of SGD, or by Orabona (2020b) for the last iterate. The difference with SGD is that it is also possible to drop the $\log(k)$ factor in (40) for the last iterate if we know the stopping time of the algorithm as shown in (41). So far in the litterature, shaving of this log factor has been shown only for convex lipschitz functions over closed bounded sets (Jain et al., 2019).

G.1. Proof of Theorem 24

The proof uses the following Lyapunov function

$$L_k = \mathbb{E} \left[\|z_k - x_*\|^2 \right] + 2\eta_k \lambda_k \mathbb{E} [f(x_{k-1}) - f_*]$$

Proof We have

$$\begin{aligned} \|z_{k+1} - x_*\|^2 &= \|z_k - x_* - \eta_k \nabla f_{v_k}(x_k)\|^2 \\ &\stackrel{\text{(SHB-IMA)}}{=} \|z_k - x_*\|^2 - 2\eta_k \langle \nabla f_{v_k}(x_k), z_k - x_* \rangle + \eta_k^2 \|\nabla f_{v_k}(x_k)\|^2 \\ &\stackrel{\text{(SHB-IMA)}}{=} \|z_k - x_*\|^2 - 2\eta_k \langle \nabla f_{v_k}(x_k), x_k - x_* \rangle - 2\eta_k \lambda_k \langle \nabla f_{v_k}(x_k), x_k - x_{k-1} \rangle + \eta_k^2 \|\nabla f_{v_k}(x_k)\|^2 \end{aligned}$$

Then taking conditional expectation $\mathbb{E}_k[\cdot] \stackrel{\text{def}}{=} \mathbb{E}[\cdot | x_k]$ we have

$$\begin{aligned} \mathbb{E}_k \left[\|z_{k+1} - x_*\|^2 \right] &= \|z_k - x_*\|^2 - 2\eta_k \langle \nabla f(x_k), x_k - x_* \rangle \\ &\quad - 2\eta_k \lambda_k \langle \nabla f(x_k), x_k - x_{k-1} \rangle + \eta_k^2 \mathbb{E}_k \left[\|\nabla f_{v_k}(x_k)\|^2 \right], \\ &\stackrel{(5)+(2)}{\leq} A_k + 4\eta_k^2 \mathcal{L}(f(x_k) - f_*) + 2\eta_k^2 \sigma^2 \\ &\quad - 2\eta_k (f(x_k) - f_*) - 2\eta_k \lambda_k (f(x_k) - f(x_{k-1})) \\ &= \|z_k - x_*\|^2 - 2\eta_k (1 + \lambda_k - 2\eta_k \mathcal{L})(f(x_k) - f_*) \\ &\quad + 2\eta_k \lambda_k (f(x_{k-1}) - f_*) + 2\eta_k^2 \sigma^2. \end{aligned} \quad (42)$$

$$\begin{aligned} &\leq \|z_k - x_*\|^2 - 2\eta_k \left(\frac{1}{2} + \lambda_k \right) (f(x_k) - f_*) \\ &\quad + 2\eta_k \lambda_k (f(x_{k-1}) - f_*) + 2\eta_k^2 \sigma^2, \end{aligned} \quad (43)$$

where we used the fact that $\eta_k \leq \frac{1}{4\mathcal{L}}$ in the last inequality. Since $\lambda_{k+1} = \frac{\sum_{t=0}^k \eta_t}{2\eta_{k+1}}$ we have that

$$\eta_{k+1}\lambda_{k+1} = \eta_k \left(\frac{1}{2} + \lambda_k \right).$$

Using this in (42) then taking expectation and rearranging gives

$$\mathbb{E} \left[\|z_{k+1} - x_*\|^2 \right] + 2\eta_{k+1}\lambda_{k+1}\mathbb{E} [f(x_k) - f_*] \leq \mathbb{E} \left[\|z_k - x_*\|^2 \right] + 2\eta_k\lambda_k\mathbb{E} [f(x_{k-1}) - f_*] + 2\eta_k^2\sigma^2.$$

Summing over $t = 0$ to k and using a telescopic sum, we have

$$\mathbb{E} \left[\|z_{k+1} - x_*\|^2 \right] + \left(\sum_{t=0}^k \eta_t \right) \mathbb{E} [f(x_k) - f_*] \leq \|x_0 - x^*\|^2 + 2\sigma^2 \sum_{t=0}^k \eta_t^2,$$

where we used that $\lambda_0 = 0$. Thus, writing λ_k explicitly, gives

$$\mathbb{E} [f(x_k) - f_*] \leq \frac{\|x_0 - x^*\|^2}{\sum_{t=0}^k \eta_t} + \frac{2\sigma^2 \sum_{t=0}^k \eta_t^2}{\sum_{t=0}^k \eta_t}.$$

■

G.2. Proof of Corollary 25

Proof (39) and (41) can be easily derived from Theorem 24. (40) requires some additional sum computations. Using the integral bound and plugging in our choice of η_k gives

$$\sum_{t=0}^{k-1} \eta_t^2 = \eta^2 \sum_{t=0}^{k-1} \frac{1}{t+1} \leq \eta^2 (\log(k) + 1) \quad \text{and} \quad \sum_{t=0}^{k-1} \eta_t \geq 2\eta (\sqrt{k} - 1),$$

which we use to obtain (40).

■