
MAKING THE LAST ITERATE OF SGD INFORMATION THEORETICALLY OPTIMAL

A PREPRINT

Prateek Jain
Microsoft Research
Bengaluru, India
prajain@microsoft.com

Dheeraj Nagaraj*
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, USA 02139
dheeraj@mit.edu

Praneeth Netrapalli
Microsoft Research
Bengaluru, India
praneeth@microsoft.com

May 30, 2019

ABSTRACT

Stochastic gradient descent (SGD) is one of the most widely used algorithms for large scale optimization problems. While classical theoretical analysis of SGD for convex problems studies (suffix) *averages* of iterates and obtains information theoretically optimal bounds on suboptimality, the *last point* of SGD is, by far, the most preferred choice in practice. The best known results for last point of SGD [1] however, are suboptimal compared to information theoretic lower bounds by a $\log T$ factor, where T is the number of iterations. [2] shows that in fact, this additional $\log T$ factor is tight for standard step size sequences of $\Theta\left(\frac{1}{\sqrt{t}}\right)$ and $\Theta\left(\frac{1}{t}\right)$ for non-strongly convex and strongly convex settings, respectively. Similarly, even for subgradient descent (GD) when applied to non-smooth, convex functions, the best known step-size sequences still lead to $O(\log T)$ -suboptimal convergence rates (on the final iterate). The main contribution of this work is to design new step size sequences that enjoy information theoretically optimal bounds on the suboptimality of *last point* of SGD as well as GD. We achieve this by designing a modification scheme, that converts one sequence of step sizes to another so that the last point of SGD/GD with modified sequence has the same suboptimality guarantees as the average of SGD/GD with original sequence. We also show that our result holds with high-probability. We validate our results through simulations which demonstrate that the new step size sequence indeed improves the final iterate significantly compared to the standard step size sequences.

Keywords Stochastic Gradient Descent · Machine Learning · Convex Optimization

1 Introduction

Stochastic Gradient Descent (SGD) is one of the most popular algorithms for solving large-scale empirical risk minimization (ERM) problems [3, 4, 5]. The algorithm updates the iterates using stochastic gradients obtained by sampling data points uniformly at random. The algorithm has been studied for several decades [6] but there are still significant gaps between *practical implementations* and theoretical analyses. In particular, the standard analyses hold only for some kind of average of iterates, but most practitioners just use the final iterate of SGD. So, [7] asked the natural question of whether the final iterate of SGD, as opposed to average of iterates, is provably good. It was partly

*

answered in [1] which gave sub-optimality bound for the last point of SGD but the obtained sub-optimality rates are $O(\log T)$ worse than the information theoretically optimal rates; T is the number of iterations.

[2] showed that the above result is tight for the standard step-size sequence used by most existing theoretical results. The extra logarithmic factor is not due to the stochastic nature of SGD. In fact, even for *subgradient descent* (GD) when applied to general non-smooth, convex functions, the last point’s convergence rates are sub-optimal by $O(\log T)$ factor.

So, this work addresses the following two fundamental questions:

“Does there exist a step-size sequence for which the last point of SGD when applied to general convex functions as well as to strongly-convex functions has optimal error (sub-optimality) rate?”, and,

“Does there exist a step-size sequence for which the last point of GD when applied to general non-smooth convex functions has optimal error (sub-optimality) rate?”

In this paper, we answer both the questions in the affirmative. That is, we provide novel step size sequences and show that the final iterate of SGD run with these step size sequences has the information theoretically optimal error (suboptimality) rate. In particular, for general non-smooth convex functions, our results ensure an error rate of $O(\frac{1}{\sqrt{T}})$ and for strongly-convex functions, the error rate is $O(\frac{1}{T})$. We also present high-probability versions, i.e., we show

that with probability at least $1 - \delta$, the suboptimality is $O\left(\sqrt{\frac{\log \frac{1}{\delta}}{T}}\right)$ and $O\left(\frac{\log \frac{1}{\delta}}{T}\right)$ respectively (see Theorems 1

and 2). For GD, we show that a similarly modified step-size sequence leads to suboptimality of $O(\frac{1}{T})$ and $O(\frac{1}{\sqrt{T}})$ for non-smooth convex functions, with and with out strong convexity respectively, which is optimal.

In general, SGD takes the iterates near the optimum value but since the objective isn’t smooth near the optimizer x^* , the gradients don’t become small even when the points are close to x^* . Standard step sizes don’t decay appreciably with time to ensure fast enough convergence to x^* . Therefore the iterates x_t , after going close to x^* , start oscillating around it without actually approaching it (See Section 4 for concrete examples). Our new step sizes, given in Section 2.1 ensure that the step sizes decay fast enough after a certain point, making the iterates go closer to the optimum x^* . The exact mode of this decay ensures that the last iterate approaches the optimum at the information theoretic rate.

Our results utilize a general step size modification scheme which ensures that the upper bounds for the average function value with the original step sizes gets transferred to the last iterate when the modified step sizes are used (see Theorems 3 and 4). A key technical contribution of the paper is the proof of Theorem 2 that constructs a sequence of averaging schemes which are ‘good’ with high probability such that the last averaging scheme consists only of the last iterate and hence lets us conclude that the last iterate is ‘good’ with high probability.

Our new step-size sequence requires that the number of iterations or horizon T is known apriori. In contrast, standard step-size sequences do not require T apriori, and hence guarantee any-time results. Information about T apriori helps us in ensuring that we do not drop step-size too early; only after we are close to the optimum, does the step size drop rapidly. In fact, we conjecture that in absence of apriori information about T , *no step-size* sequence can ensure the information theoretically optimal error rates for final iterate of SGD. As a step towards proving this, we show that in the case of strongly convex objectives, any choice of step sizes with infinite horizon (i.e, without the knowledge of total number of iterations) is either suboptimal almost surely or suboptimal in expectation for infinitely many points. We show this in Theorem 5.

Related Work: Averaging was used first in the stochastic approximation setting by [8] to show optimal rates of convergence. Gradient Descent type methods have been shown to achieve information theoretically optimal error rates in the convex and strongly convex settings when averaging of iterates is used ([9],[10],[11], [12], Epoch GD in [13], SGD [14] and [15]). The question of the last iterate was first considered in [1] and it gives a bound of $O(\frac{\log T}{\sqrt{T}})$ and $O(\frac{\log T}{T})$ in expectation for the general case and strongly convex case respectively. [2] show matching high probability bounds and show that for the standard step sizes ($O(\frac{1}{\sqrt{t}}$) in the general case and $O(\frac{1}{t})$ in the strongly convex case), the logarithmic-suboptimal bounds are tight.

Organization: The setting and main results are presented in Section 2. In particular, Section 2.1 describes the general step size modification considered and states key results regarding this modification and the lower bound is presented in Section 2.2. Key technical ideas are developed in Section 3 and the main theorems are proved. We present some experimental results in Section 4 and conclude in Section 5. Skipped proofs of technical lemmas are given in the appendix.

2 Problem Setup and Main Results

Consider the following optimization problem:

$$\min_{x \in \mathcal{W}} F(x), \quad (1)$$

where objective function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex function and $\mathcal{W} \subset \mathbb{R}^d$ is a closed convex set. Let the global minimizer of $F(\cdot)$ be $x^* \in \mathcal{W}$. We start the SGD algorithm at a point $x_1 \in \mathcal{W}$ and iteratively obtain estimates x_t for the minimizer of $F(\cdot)$. We assume that at each time step, we have access to independent, unbiased estimate \hat{g}_t to a subgradient $g_t \in \partial F$. That is, $\mathbb{E}[\hat{g}_t(x)] = g_t(x) \in \partial F(x)$ for every $x \in \mathcal{W}$ and $(\hat{g}_t - g_t)_{t=1}^T$ are independent. We pick step sizes $(\alpha_t)_{t=1}^T \geq 0$. Let $\Pi_{\mathcal{W}}$ be the projection operator to the set \mathcal{W} . The SGD algorithm is given as follows:

Input: total time T and step sizes α_t
Output: x_T
for $t \leftarrow 1$ **to** T **do**
 | $x_{t+1} \leftarrow \Pi_{\mathcal{W}}(x_t - \alpha_t \hat{g}_t(x_t))$.
end

Algorithm 1: Stochastic Gradient Descent

Henceforth, we will retain the assumptions made above. Whenever we use $g_t(x)$, it is implied that $g_t(x) \in \partial F(x)$. Throughout the paper, we assume that F is a Lipschitz continuous convex function.

Assumption 1 (Lipschitz Continuity). $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is G -Lipschitz continuous convex function over closed convex set \mathcal{W} , i.e., $\|g(x)\| \leq G$ for every $x \in \mathcal{W}$ and every $g(x) \in \partial F(x)$. Furthermore, the stochastic gradients \hat{g} satisfy: $\|\hat{g}(x)\| \leq G$ almost surely for every $x \in \mathcal{W}$.

Assumption 2 (Closed and bounded set). Diameter of closed convex set \mathcal{W} is bounded by D , i.e., $\text{diam}(\mathcal{W}) \leq D$.

Assumption 3 (Strong convexity). Let $\lambda > 0$. A convex function F is said to be λ strongly convex over \mathcal{W} iff $F(y) \geq F(x) + \langle \nabla F(x), y - x \rangle + \frac{\lambda}{2} \|y - x\|^2 \forall x, y \in \mathcal{W}$.

Step size sequence for general convex functions: we first define,

$$k := \inf\{i : T \cdot 2^{-i} \leq 1\}, \quad T_i := T - \lceil T \cdot 2^{-i} \rceil, \quad 0 \leq i \leq k, \quad \text{and } T_{k+1} := T. \quad (2)$$

Clearly, $0 = T_0 < T_1 < \dots < T_k = T - 1 < T_{k+1} = T$. We note in particular that $T_1 \approx \frac{T}{2}$. Let $C > 0$ be arbitrary. Then, we choose the step size α_t as follows:

$$\alpha_t = \frac{C \cdot 2^{-i}}{\sqrt{T}} \quad \text{when } T_i < t \leq T_{i+1}, \quad 0 \leq i \leq k. \quad (3)$$

The theorem below provides suboptimality guarantee for the SGD algorithm with the step-size sequence mentioned above.

Theorem 1 (SGD/GD Last Point for General Convex Functions). *Let Assumptions 1 and 2 hold. Given $T \geq 4$, let x_1, \dots, x_T be the iterates of SGD (Algorithm 1) with step size α_t as defined in Equation (3). Then, the following holds for all $T \geq 4$:*

$$\mathbb{E}[F(x_T)] \leq F(x^*) + \frac{4D^2}{C\sqrt{T}} + \frac{11G^2C}{\sqrt{T}}.$$

In particular, if we choose $C = \frac{D}{\delta}$, we have: $\mathbb{E}F(x_T) \leq F(x^) + \frac{15GD}{\sqrt{T}}$. Furthermore, the following holds w.p. $\geq 1 - \delta$ for any $0 < \delta < \frac{1}{e}$:*

$$F(x_T) = F(x^*) + O\left(\frac{D^2}{C\sqrt{T}} + \frac{CG^2}{\sqrt{T}} \log\left(\frac{1}{\delta}\right)\right) \leq F(x^*) + O\left(DG\sqrt{\frac{\log\frac{1}{\delta}}{T}}\right).$$

Finally, under the same assumptions, GD update ($x_{t+1} = \Pi_{\mathcal{W}}(x_t - \alpha_t \nabla F(x_t))$) with the same step-size sequence given in (3) also ensures the following after T iterations:

$$F(x_T) \leq F(x^*) + \frac{4D^2}{C\sqrt{T}} + \frac{11G^2C}{\sqrt{T}}.$$

We will prove this theorem in Section 3 after developing some general ideas.

Remarks: (1) Note that the bounds on sub-optimality (for SGD and GD) are information theoretically optimal up to constants.

(2) Our result on the expected sub-optimality improves upon that of [1] by a multiplicative $\log T$ factor and our result on the high probability sub-optimality improves upon [2] by a multiplicative factor of $\log T \sqrt{\log \frac{1}{\delta}}$. On the other hand, our step-size sequence requires apriori knowledge of T . We conjecture that for any-time algorithm (i.e., without apriori knowledge of T) expected error rate of $\frac{GD \log T}{\sqrt{T}}$ is information theoretically optimal.

(3) The rate obtained above for last point of GD (in the deterministic setting) is also optimal in the gradient oracle model and to the best of our knowledge, is the first such result for last point of GD.

Step size sequence for strongly-convex functions: Let $F(\cdot)$ be λ strongly convex (Assumption 3). Let $k := \inf\{i : T \cdot 2^{-i} \leq 1\}$. We pick α_t as follows:

$$\alpha_t = 2^{-i} \frac{1}{\lambda t}, \quad \forall T_i < t \leq T_{i+1}, 0 \leq i \leq k. \quad (4)$$

We now present our result for last point of SGD with strong-convexity assumption.

Theorem 2 (SGD Last Point for Strongly Convex Functions). *Let F satisfy Assumptions 1 and 3. Then the following holds for the T -th iterate of the SGD algorithm (Algorithm 1) when run with the step size sequence given in Equation (4):*

$$\mathbb{E}[F(x_T)] \leq F(x^*) + \frac{130G^2}{\lambda T}.$$

Furthermore, the following holds for all $0 < \delta \leq 1/e$ with probability at least $1 - \delta$:

$$\mathbb{E}[F(x_T)] = F(x^*) + O\left(\frac{G^2 \log(\frac{1}{\delta})}{\lambda T}\right).$$

Under the same assumptions, GD update ($x_{t+1} = \Pi_{\mathcal{W}}(x_t - \alpha_t \nabla F(x_t))$) with the same step-size sequence given in (4) also ensures the following after T iterations:

$$F(x_T) \leq F(x^*) + \frac{130G^2}{\lambda T}.$$

Here again, we note that the result is information theoretically optimal up to $\log(1/\delta)$ factor.

2.1 General Step Size Modification

Theorems 1 and 2 are consequences of our general results on step size modification that we present below. Consider SGD step size sequence $(\gamma_t)_{t=1}^T$. We obtain modified step size sequence $(\alpha_t)_{t=1}^T$ as follows:

$$\alpha_t := 2^{-i} \gamma_t \quad \forall T_i < t \leq T_{i+1} \text{ and } 0 \leq i \leq k. \quad (5)$$

Under certain mild conditions, we will show that the last iterate of SGD with step size α_t is as good as the average iterate of SGD with step size γ_t . We make these notions precise below:

Assumption 4 (Slowly Decreasing Step Size Sequence). *We call a step size sequence (γ_t) ‘decreasing’ if $\gamma_{t+1} \leq \gamma_t$. We say that step size sequence γ_t has ‘at most polynomial decay’ with decay constant $0 < \beta \leq 1$ if $\gamma_{2t} \geq \beta \gamma_t$ for every $t \geq 1$.*

We have the following general theorem:

Theorem 3. *Let $(\gamma_t)_{t=1}^T$ be a decreasing step size sequence with at most polynomial decay with decay constant $0 < \beta \leq 1$. Let the iterates of SGD with step size γ_t be y_1, \dots, y_T . Let α_t be the modification of γ_t as defined in Equation (5). Let the iterates of SGD with step size α_t be x_1, \dots, x_T . Then, for all $T \geq 4$, we have:*

$$\mathbb{E}[F(x_T)] \leq 5G^2 \gamma_T \left(\frac{1}{\beta^2} + \frac{1}{\beta^4} \right) + \inf_{\lceil \frac{T}{4} \rceil \leq t \leq T_1} \mathbb{E}[F(y_t)].$$

We also give a high probability version of Theorem 3.

Theorem 4. Let $T \geq 4$. Let $q^{(0)}$ be any arbitrary fixed probability distribution over the set $\{\lceil \frac{T}{4} \rceil, \dots, T_1\}$. With probability atleast $1 - \frac{\delta}{2}$, we have:

$$F(x_T) \leq \gamma_T \cdot G^2 \left(120 \log\left(\frac{1}{\delta}\right) + 400\right) \cdot \left(\frac{1}{\beta^2} + \frac{1}{\beta^4}\right) + \sum_{s=\lceil \frac{T}{4} \rceil}^{T_1} q^{(0)}(s) F(y_s).$$

That is, the above theorems show that compared to any weighted average of function values of iterates in the $[\lceil T/4 \rceil, T_1]$ iterations, the error is not significantly larger if β is reasonably large and γ_T is small. Now, using standard analysis, we can ensure small average function value for iterates in $[\lceil T/4 \rceil, T_1]$ iterations. Small value of γ_T and bound on β hold trivially for standard step-size sequences.

See Section 3 for detailed proofs of the above theorems. We first develop general technique and prove key lemmas in the next section, and then present proofs for all the theorems.

2.2 Lower Bounds

The step size modification procedure described above assumed the knowledge of the last iterate T (this is not a setback in practice). We study the case of infinite horizon SGD. In this section we state our bounds on the last iterate of ‘any time’ (infinite horizon) SGD in the case of strongly convex objectives. We will first introduce the notion of suboptimality that we consider. In particular, we look at two kinds of ‘bad performance’ in infinite horizon SGD for non-smooth strongly convex optimization. Consider any infinite step size sequence γ_t .

1. The sequence γ_t is said to be ‘bad in expectation’ if for an objective F satisfying assumptions 1,2 and 3, some choice of subgradient oracle, and SGD iterates $(x_t)_{t \in \mathbb{N}}$ with step size γ_t , there is a fixed subsequence $\{t_k\}_{k \in \mathbb{N}}$ such that $\lim_{k \rightarrow \infty} t_k \mathbb{E}[F(x_{t_k}) - F(x^*)] = \infty$.
2. The sequence γ_t is said to be ‘bad almost surely’ if for an objective F satisfying assumptions 1,2 and 3, some choice of subgradient oracle, and SGD iterates $(x_t)_{t \in \mathbb{N}}$ with step size γ_t , with probability 1 there exists a random infinite sequence of times $\{t_k\}$ such that $\lim_{k \rightarrow \infty} t_k [F(x_{t_k}) - F(x^*)] = \infty$

We give a ‘no free lunch’ theorem: that is we show that infinite horizon step-size sequence for non-smooth strongly convex optimization is either ‘bad in expectation’ or ‘bad almost surely’. More precisely, we will show that if any infinite horizon SGD is good in ‘expectation’ for every t for every strongly convex function, then it is ‘bad almost surely’ for some function F .

Theorem 5. Consider infinite horizon SGD with step size γ_t such that assumptions 1 2 and 3 hold for the objective function. Then, for any choice of $\gamma_t > 0$, the algorithm is either bad in expectation or bad almost surely.

We give the proof in Section B.

3 Technical Ideas and Proofs

Recall the definition of T_i from Section 2. The rough idea behind the proof is as follows: we will find a ‘good point’ in the range $[\lceil \frac{T}{4} \rceil, T_1]$ and then show that this implies that there is a ‘good point’ between $T_1 = T/2$ and $T_2 \approx 3T/4$ and so on, until we conclude that x_T is a good point.

To this end, we first provide a key lemma that bounds the total weighted deviation of SGD iterates from a given iterate x_{t_0} (in terms of function value), i.e., it intuitively shows that once we find an iterate with small function value, the remaining iterates cannot deviate from it significantly. The lemma uses a trick that was first used in [16] and then also in [1].

Lemma 1. Let x_1, \dots, x_T be the output of SGD algorithm (Algorithm 1) with step size sequence α_t defined by (3). Then, given any $1 < t_0 < t_1 \leq T$,

$$\sum_{t=t_0}^{t_1} 2\alpha_t \mathbb{E}[F(x_t) - F(x_{t_0})] \leq \sum_{t=t_0}^{t_1} G^2 \alpha_t^2.$$

Proof. By convexity of \mathcal{W} , we have:

$$\|x_{t+1} - x_{t_0}\| = \|\Pi_{\mathcal{W}}(x_t - \alpha_t \hat{g}_t(x_t)) - x_{t_0}\| \leq \|x_t - \alpha_t \hat{g}_t(x_t) - x_{t_0}\|$$

Taking squares and expanding on both sides,

$$\|x_{t+1} - x_{t_0}\|^2 \leq \|x_t - x_{t_0}\|^2 + \alpha_t^2 \|\hat{g}_t(x_t)\|^2 - 2\alpha_t \langle \hat{g}_t(x_t), x_t - x_{t_0} \rangle$$

Taking expectation on both sides, and realizing that \hat{g}_t is independent of x_t and x_{t_0} , we conclude,

$$\mathbb{E} [\|x_{t+1} - x_{t_0}\|^2] \leq \mathbb{E} [\|x_t - x_{t_0}\|^2] + \alpha_t^2 G^2 - 2\alpha_t \mathbb{E} \langle g_t, x_t - x_{t_0} \rangle$$

Here we have used the fact that $\mathbb{E} [\hat{g}_t(x_t) | x_t, x_{t_0}] = g_t(x_t)$. Using convexity, $\langle g_t, x_t - x_{t_0} \rangle$ is lower bounded by $F(x_t) - F(x_{t_0})$. We conclude that:

$$\mathbb{E} [\|x_{t+1} - x_{t_0}\|^2] \leq \mathbb{E} [\|x_t - x_{t_0}\|^2] + \alpha_t^2 G^2 - 2\alpha_t \mathbb{E} [F(x_t) - F(x_{t_0})]$$

The result now follows by summing the above term from $t = t_0$ to $t = t_1$. \square

We now provide a high probability version of Lemma 1. To this end, we construct an exponential super-martingale that when combined with a Chernoff bound leads to exponential concentration bound. The method used is somewhat similar to the one used in [2], but our technique is specifically for Lemma 1 and is more concise.

For simplicity of exposition, we first define a few key quantities. Let $1 < t_0 < t_1 \leq T$ and $r = t_1 - t_0 + 1$. We define the sequence L_t as follows: for $t_0 \leq t \leq t_1$ as follows:

$$L_{t_1} = \frac{1}{e \cdot r}, \quad t \in [t_0, t_1], \quad L_{t-1} = L_t + L_t^2, \quad t_0 \leq t-1 < t_1. \quad (6)$$

Using Lemma 3, $\frac{1}{e \cdot r} \leq L_t \leq \frac{1}{r}$. Now, for any l such that $t_0 \leq l \leq t_1$, we define the following random variables :

$$A(l, t_1) := \sum_{t=l}^{t_1} L_t [2\alpha_t (F(x_t) - F(x_l)) - \alpha_t^2 G^2], \quad A^*(t_0, t_1) := \sum_{t=t_0}^{t_1} L_t [2\alpha_t (F(x_t) - F(x^*)) - \alpha_t^2 G^2]. \quad (7)$$

We note the difference between $A^*(t_0, t_1)$ and $A(l, t_1)$: $A(l, t_1)$ considers suboptimality with respect to x_l whereas $A^*(t_0, t_1)$ considers the suboptimality with respect to the optimizer x^* .

Lemma 2. *Let A and A^* be as defined by (7). Let $p(t_0), \dots, p(t_1)$ be any probability distribution over $\{t_0, \dots, t_1\}$. We let $(p.A)(t_0, t_1) := \sum_{l=t_0}^{t_1} p(l)A(l, t_1)$. Also, let α_t be a decreasing step size sequence. Then,*

$$\mathbb{P} [(p.A)(t_0, t_1) > \eta] \leq \exp\left(-\frac{\eta}{8\alpha_{t_0}^2 G^2}\right).$$

Additionally, if $\text{diam}(\mathcal{W}) \leq D$ almost surely, we have:

$$\mathbb{P} [A^*(t_0, t_1) > \eta] \leq \exp\left(\frac{2D^2 L_{t_0}}{8\alpha_{t_0}^2 G^2}\right) \exp\left(\frac{-\eta}{8\alpha_{t_0}^2 G^2}\right).$$

Lemma 3. *Let $\Gamma > 0$ be fixed. Let $\lambda_0 = \frac{1}{re\Gamma}$, $\lambda_1 = \lambda_0 + \Gamma\lambda_0^2$, \dots , $\lambda_{i+1} = \lambda_i + \Gamma\lambda_i^2$. Then, for every $i \leq r$, $\lambda_i \leq (1 + \frac{1}{r})^i \lambda_0$*

See Section A for proofs of the above given lemmata. We also require the following technical lemma:

Lemma 4. *Let T_i be as defined in Section 2. Then, for all $0 \leq i \leq k-1$:*

$$4(T_{i+2} - T_{i+1}) \geq T_{i+1} - T_i.$$

Proof. Lemma follows from the fact that $2\lceil a \rceil - 1 \leq \lceil 2a \rceil \leq 2\lceil a \rceil$. \square

3.1 Step Size Modification

Henceforth, we will assume that γ_t is a decreasing step size sequence with at most polynomial decay (decay constant being β). We let α_t be the modification of γ_t as defined in Equation 5. Let,

$$\tau_i := \arg \inf_{T_i < t \leq T_{i+1}} \mathbb{E}[F(x_t)], \quad i \in [k+1], \quad \text{and} \quad \tau_0 := \arg \inf_{\lceil \frac{T}{4} \rceil \leq t \leq T_1} \mathbb{E}[F(x_t)]. \quad (8)$$

Note that $\tau_{k+1} = T$. We note that τ_i are completely deterministic and only used as part of the proof. The ability to compute τ_i is not necessary.

Lemma 5. Let x_t 's be iterates of SGD (Algorithm 1) with modified step size sequence α_t of γ_t defined in (5); γ_t sequence satisfies Assumption 4. Let T_i, k be as defined by (2), and $\tau_i, 0 \leq i \leq k+1$ be as defined in (8). Also, let $T \geq 4$. Then, the following holds for all $i \in [k]$:

$$\mathbb{E}[F(x_{\tau_{i+1}}) - F(x_{\tau_i})] \leq \frac{5G^2\gamma_T}{\beta^2} 2^{-i}, \quad \mathbb{E}[F(x_{\tau_1}) - F(x_{\tau_0})] \leq \frac{5G^2\gamma_T}{\beta^4}.$$

Proof. We first consider $i \geq 1$. If $\mathbb{E}[F(x_{\tau_{i+1}})] \leq \mathbb{E}[F(x_{\tau_i})]$, the proof is done. Else, using Lemma 1 with $t_0 = \tau_i$ and $t_1 = T_{i+2}$, and the fact that α_t is a decreasing sequence, we get:

$$\frac{\sum_{t=\tau_i}^{T_{i+2}} 2\alpha_t \mathbb{E}[F(x_t) - F(x_{\tau_i})]}{T_{i+2} - \tau_i + 1} \leq \frac{\sum_{t=\tau_i}^{T_{i+2}} G^2 \alpha_t^2}{T_{i+2} - \tau_i + 1} \leq G^2 \alpha_{T_{i+1}}^2. \quad (9)$$

By definition of τ_i , $\mathbb{E}[F(x_{\tau_i})] \leq \mathbb{E}[F(x_t)]$ whenever $T_i < t \leq T_{i+1}$. Hence,

$$G^2 2^{-2i} \gamma_{T_{i+1}}^2 = G^2 \alpha_{T_{i+1}}^2 \geq \frac{\sum_{t=\tau_i}^{T_{i+2}} 2\alpha_t \mathbb{E}[F(x_t) - F(x_{\tau_i})]}{T_{i+2} - \tau_i + 1} \geq \frac{\sum_{t=T_{i+1}+1}^{T_{i+2}} 2\alpha_t \mathbb{E}[F(x_t) - F(x_{\tau_i})]}{T_{i+2} - \tau_i + 1}, \quad (10)$$

where the first equality follows from the definition of α_t in (5), first inequality follows from Equation (9), and the final inequality follows from the fact that $\mathbb{E}[F(x_t) - F(x_{\tau_i})] \geq 0$ when $T_i < t \leq T_{i+1}$ (see definition of τ_i in (8)).

Now, by using the above inequality with the assumption $\mathbb{E}[F(x_{\tau_{i+1}})] \geq \mathbb{E}[F(x_{\tau_i})]$, and the fact that $T_{i+2} - T_i \geq T_{i+2} - \tau_i + 1$, we have:

$$\begin{aligned} G^2 2^{-2i} \gamma_{T_{i+1}}^2 &\geq 2\alpha_{T_{i+2}} \frac{T_{i+2} - T_{i+1}}{T_{i+2} - T_i} \mathbb{E}[F(x_{\tau_{i+1}}) - F(x_{\tau_i})] \stackrel{\zeta_1}{\geq} \frac{2\alpha_{T_{i+2}}}{5} \mathbb{E}[F(x_{\tau_{i+1}}) - F(x_{\tau_i})] \\ &= \frac{2^{-i} \gamma_{T_{i+2}}}{5} \mathbb{E}[F(x_{\tau_{i+1}}) - F(x_{\tau_i})] \geq \frac{2^{-i} \beta \gamma_{T_{i+1}}}{5} \mathbb{E}[F(x_{\tau_{i+1}}) - F(x_{\tau_i})], \end{aligned} \quad (11)$$

where ζ_1 follows from Lemma 4. The equality follows from definition of α_t and the last inequality follows from the β -slowly decaying assumption for γ_t (Assumption 4). That is we obtain the result for the case $i \geq 1$. The proof for the case when $i = 0$ follows with minor modifications to the arguments given above. \square

We now present a high probability version of Lemma 5.

Lemma 6. Consider the setting of Lemma 5. Let $0 \leq i \leq k$ and define $t_0 = T_i + 1$ for $1 \leq i$ and $t_0 = \lceil \frac{T}{4} \rceil$ for $i = 0$. Let $q^{(i)}$ be any probability distribution over $\{t_0, \dots, T_{i+1}\}$. Let $p^{i+1}(t) := \frac{L_t \alpha_t}{\sum_{s=T_{i+1}+1}^{T_{i+2}} L_s \alpha_s}$, where $t \in [T_{i+1} + 1, T_{i+2}]$

and the sequence $(L_t)_{T_{i+1}+1}^{T_{i+2}}$ is defined by (6). Then, for any $\delta_i \in (0, 1)$ and $i \in [1, k-1]$, the following holds with probability at least $1 - \delta_i$:

$$\sum_{t=T_{i+1}+1}^{T_{i+2}} p^{(i+1)}(t) F(x_t) \leq \frac{G^2 \gamma_T 2^{-i}}{\beta^2} \left(15 + 120 \log \frac{1}{\delta_i}\right) + \sum_{s=T_{i+1}}^{T_{i+1}} q^{(i)}(s) F(x_s).$$

For $i = 0$, the following holds with probability atleast $1 - \delta_0$:

$$\sum_{t=T_1+1}^{T_2} p^{(1)}(t) F(x_t) \leq \frac{G^2 \gamma_T 2^{-i}}{\beta^4} \left(15 + 120 \log \frac{1}{\delta_i}\right) + \sum_{s=\lceil \frac{T}{4} \rceil}^{T_1} q^{(0)}(s) F(x_s).$$

Proof. We will only show the case $1 \leq i \leq k-1$. The $i = 0$ case follows by a similar proof. For $T_i \leq t \leq T_{i+1}$, we define $\Gamma(t) = \sum_{s=t+1}^{T_{i+2}} \alpha_s L_s$. We let κ be defined as follows over $\{T_i + 1, \dots, T_{i+2}\}$:

$$\begin{aligned} \kappa(T_i + 1) &:= \frac{\Gamma(T_{i+1})}{\Gamma(T_i + 1)} \cdot q^{(i)}(T_i + 1), \quad \kappa(t) := \frac{\Gamma(T_{i+1})}{\Gamma(t)} q^{(i)}(t) + \frac{\alpha_t L_t \cdot \left(\sum_{s=T_i+1}^{t-1} \kappa(s)\right)}{\Gamma(t)}, \quad t \in (T_i + 1, T_{i+1}], \\ \kappa(t) &:= 0, \quad \forall t \geq T_{i+1}. \end{aligned} \quad (12)$$

From Lemma 7, we conclude that κ is a probability distribution over $\{T_i + 1, \dots, T_{i+2}\}$. From Lemma 2, we conclude that with probability atleast $1 - \delta_i$:

$$(\kappa.A)(t_0, t_1) \leq 8\alpha_{t_0}^2 G^2 \log \frac{1}{\delta_i} \quad (13)$$

We will show that when this event happens, the inequality in the statement of the lemma holds. If $\sum_{t=T_{i+1}+1}^{T_{i+2}} p^{(i+1)}(t)F(x_t) \leq \sum_{s=T_{i+1}}^{T_{i+1}} q^{(i)}(s)F(x_s)$, then the statement of the lemma holds trivially. Now assume $\sum_{t=T_{i+1}+1}^{T_{i+2}} p^{(i+1)}(t)F(x_t) > \sum_{s=T_{i+1}}^{T_{i+1}} q^{(i)}(s)F(x_s)$. We use the fact that κ is supported over $\{T_i + 1, \dots, T_{i+1}\}$ and hence:

$$(\kappa.A)(t_0, t_1) = \sum_{l=T_{i+1}}^{T_{i+1}} \sum_{t=l}^{T_{i+2}} \kappa(l)L_t [2\alpha_t(F(x_t) - F(x_l)) - \alpha_t^2 G^2]$$

We exchange summation and collect the coefficients of the term $F(x_t)$ to conclude:

$$(\kappa.A)(t_0, t_1) = \sum_{t=T_{i+1}+1}^{T_{i+2}} L_t (\alpha_t F(x_t) - \alpha_t^2 G^2) - \sum_{s=T_{i+1}}^{T_{i+1}} (\alpha_s^2 G^2 \sigma_s L_s + 2F(x_s) (\sigma_s \alpha_s L_s - \kappa(s) \Gamma(s-1))),$$

where $\sigma(s) := \sum_{t=T_{i+1}}^s \kappa(s)$ (empty sum being 0 by definition). By definition of $\sigma(s) = \kappa(s) + \sigma(s-1)$, $\Gamma(s-1) = \alpha_s L_s + \Gamma(s)$ and $\kappa(s) = \frac{\Gamma(T_{i+1})}{\Gamma(s)} q^{(i)}(s) + \frac{\alpha_s L_s}{\Gamma(s)} \sigma(s-1)$. Therefore, we conclude:

$$\begin{aligned} (\kappa.A)(t_0, t_1) &= \sum_{T_{i+1}+1}^{T_{i+2}} 2\alpha_t L_t F(x_t) - \sum_{T_{i+1}+1}^{T_{i+2}} \alpha_t^2 G^2 L_t - \sum_{s=T_{i+1}}^{T_{i+1}} \alpha_s^2 G^2 L_s \sigma_s \\ &\quad - \left(\sum_{T_{i+1}+1}^{T_{i+2}} 2\alpha_t L_t \right) \left(\sum_{s=T_{i+1}}^{T_{i+1}} q^{(i)}(s) F(x_s) \right). \end{aligned} \quad (14)$$

We recall that $p^{(i+1)}(t) \cdot \left(\sum_{s=T_{i+1}+1}^{T_{i+2}} \alpha_s L_s \right) = \alpha_t L_t$ whenever $T_{i+1} < t \leq T_{i+2}$. The rest of the proof is similar to Equation (11) in Lemma 5. We use the fact that α_t is the modification of γ_t , γ_t has at most polynomial decay, $\frac{1}{e^{(T_{i+2}-T_i)}} \leq L_t \leq \frac{1}{T_{i+2}-T_i}$ and Lemma 4 in Equation (14) to conclude the result. \square

Lemma 7. *Let κ be as defined in (12). Then, κ is a probability distribution over $\{T_i + 1, \dots, T_{i+2}\}$.*

The proof of this lemma is given in Section A

3.2 Proof of Theorem 3

Proof. Recall the definition of τ_i in (8). Clearly, $\tau_{k+1} = T$. Summing the bounds in Lemma 5 we conclude:

$$\begin{aligned} \mathbb{E}[F(x_T)] &= \mathbb{E}[F(x_{\tau_{k+1}})] = \mathbb{E}[F(x_{\tau_0})] + \sum_{i=0}^k \mathbb{E}[F(x_{\tau_{i+1}}) - F(x_{\tau_i})] \\ &\leq \mathbb{E}[F(x_{\tau_0})] + \frac{5G^2 \gamma_T}{\beta^4} + \sum_{i=1}^k \frac{5G^2 \gamma_T}{\beta^2} 2^{-i} \leq 5G^2 \gamma_T \left(\frac{1}{\beta^2} + \frac{1}{\beta^4} \right) + \inf_{\lceil \frac{T}{4} \rceil \leq t \leq T_1} \mathbb{E}[F(x_t)]. \end{aligned}$$

We conclude the result by noting that $x_t = y_t$ for all $t \leq T_1$. \square

3.3 Proof of Theorem 4

Proof. This proof is similar to the proof of Theorem 3, but instead of Lemma 5 we use Lemma 6. In Lemma 6, we pick $q^{(i)} = p^{(i)}$ for $1 \leq i \leq k-1$ and we let $q^{(0)}$ be arbitrary. We let $\delta_i = \frac{\delta}{2^{i+2}}$. By union bound, the inequalities in the statement of Lemma 6 hold for all $0 \leq i \leq k-1$ simultaneously with probability at least $1 - \sum_{i=0}^{k-1} \delta_i \geq 1 - \frac{\delta}{2}$. Summing all these inequalities, we conclude:

$$\sum_{t=T_k+1}^{T_{k+1}} p^{(k)}(t)F(x_t) \leq \gamma_T G^2 [120 \log(\frac{1}{\delta}) + 400] \left[\frac{1}{\beta^2} + \frac{1}{\beta^4} \right] + \sum_{s=\lceil \frac{T}{4} \rceil}^{T_1} q^{(0)}(s)F(x_s).$$

We note that the distribution $p^{(k)}$ has unit mass over the point $T_{k+1} = T$ and that $x_t = y_t$ when $t \leq T_1$ to conclude the result. \square

3.4 Proof of Theorem 1

Proof. We note that the step size defined in Equation (3) is the modification of the standard step size $\gamma_t = \frac{C}{\sqrt{t}}$. Let y_t be the output of SGD under the assumptions of the theorem when step size γ_t is used. Using the fact that infimum is smaller than any weighted average, we have:

$$\begin{aligned} \inf_{\lceil \frac{T}{4} \rceil \leq t \leq T_1} \mathbb{E}[F(y_t) - F(x^*)] &\leq \frac{1}{T_1 - \lceil \frac{T}{4} \rceil + 1} \sum_{t=\lceil \frac{T}{4} \rceil}^{T_1} \mathbb{E}[F(y_t) - F(x^*)] \\ &\leq \frac{2}{T_1} \sum_{t=1}^{T_1} \mathbb{E}[F(y_t) - F(x^*)] \stackrel{\zeta_1}{\leq} \frac{2}{\sqrt{T_1}} \left[\frac{D^2 \sqrt{T_1}}{C} + CG^2 \sqrt{T_1} \right] \leq \frac{4}{\sqrt{T}} \left[\frac{D^2}{C} + CG^2 \right], \end{aligned}$$

where the second line follows from $T_1 \leq 2(T_1 - \lceil \frac{T}{4} \rceil + 1)$. ζ_1 follows from the standard analysis [6]. Final inequality follows from the fact that $\frac{T}{4} \leq T_1 \leq \frac{T}{2}$. We note that γ_t satisfies the conditions for Theorem 3 with $\beta = 1$. We invoke Theorem 3 to conclude the bound on expectation. The above proof in expectation also works for GD. We take $\hat{g}_t = \nabla F$ and SGD is the same as GD. Here each x_t and y_t is a deterministic point mass. Therefore, the expectation bound for the last iterate of SGD holds for the last iterate of GD.

We will now prove the high probability bound. Let $t_0 = \lceil \frac{T}{4} \rceil$, $t_1 = T_1$, and $\alpha_t = \gamma_t = \frac{C}{\sqrt{t}}$ for $t \in [t_0, t_1]$. Then using Lemma 2, the following holds with probability atleast $1 - \delta$:

$$A^*(t_0, t_1) \leq 2D^2 L_{t_0} + 8\alpha_{t_0}^2 G^2 \log \frac{1}{\delta}.$$

Using $\frac{1}{e(t_1 - t_0 + 1)} \leq L_t \leq \frac{1}{t_1 - t_0 + 1}$ and proceeding similarly as above, we have w.p. $\geq 1 - \frac{\delta}{2}$,

$$\inf_{\lceil \frac{T}{4} \rceil \leq t \leq T_1} \mathbb{E}[F(y_t) - F(x^*)] \leq \frac{6D^2}{C\sqrt{T}} + \frac{5CG^2 \log \frac{2}{\delta}}{\sqrt{T}}.$$

Theorem now follows by using Theorem 4 with $\beta = 1$, $q^{(0)}(t) = \frac{1}{T_1 - \lceil \frac{T}{4} \rceil + 1}$, and union bound. \square

3.5 Proof of Theorem 2

Proof. We note that the step size defined in Equation (4) is the modification of the standard step size $\gamma_t = \frac{1}{\lambda t}$ used for strongly convex functions (see [14]). Let y_1, \dots, y_T be the output of SGD when step size γ_t is used. From Theorem 5 in [14], we conclude that:

$$\inf_{\lceil T/4 \rceil \leq t \leq T_1} \mathbb{E}[F(y_t) - F(x^*)] \leq \frac{1}{T_1 - \lceil \frac{T}{4} \rceil + 1} \sum_{t=\lceil T/4 \rceil}^{T_1} \mathbb{E}[F(y_t) - F(x^*)] \leq 30 \frac{G^2}{\lambda T}.$$

The expectation bound follows from using above equation with Theorem 3 and noting that γ_t satisfies the required conditions with $\beta = 2$. We get high probability bounds by invoking high probability bounds for suffix averaging from [2], i.e., w.p. at least $1 - \frac{\delta}{2}$,

$$\inf_{\lceil \frac{T}{4} \rceil \leq t \leq T_1} F(y_t) - F(x^*) \leq O\left(\frac{G^2 \log(\frac{1}{\delta})}{\lambda T}\right).$$

The result now follows by using Theorem 4 with $\beta = 2$ and $q^{(0)}(t) = \frac{1}{T_1 - \lceil \frac{T}{4} \rceil + 1}$. \square

4 Experiments

We now empirically compare SGD last point with our step-size sequence (Our Method) with the standard steps size sequence (Standard) as well as the averaged iterates of SGD (Averaged). We apply these methods on two non-smooth problems: a) Lasso regression, b) linear SVM training.

Lasso Regression: We consider gradient descent for $F(x) = \frac{1}{n} \sum_{i=1}^n \|\langle a_i, x \rangle - b_i\|^2 + \lambda \|x\|_1$ for $x \in \mathbb{R}^d$. Here $a_i \sim \mathcal{N}(0, I_d)$ and $b_i = \langle a_i, x^* \rangle + z_i$ for some s sparse vector x^* and $z_i \sim \mathcal{N}(0, \sigma^2)$. a_i and z_i are all independent.

Figure 1: (a) $F(x)$ vs number of iterations for the Lasso regression (Section 4). Here $d = 100$, $s = \|x^*\|_0 = 60$, $n = 80$, $C = 4$, $\sigma = 0.1$, $\lambda = 0.2$ and $T = 2^{19}$. Green line indicates the running average of the iterates from 1 to t . (b) SVM loss (15) vs number of iterations. We pick $d = 30$, $\sigma = 5$, $\eta = 1$, $n = 500$, $\lambda = 0.1$, $T = 2^{17}$. (c) Average over 100 independent SGD runs for the SVM loss. Green line is the loss of the average of the last $\frac{T}{4}$ iterates. In all the cases, SGD last point with our step-size sequence produces smaller objective value than the standard step size as well as the averaged iterates of SGD.

We use the step sizes of $\gamma_t = \frac{C}{\sqrt{T}}$ and let α_t be the modification of γ_t as given in Section 2 for total T iterations. Since the objective is not smooth, the gradient doesn't vanish near the optimum. Therefore, when the standard step size was picked, the iterate x_t kept oscillating around the infimum but never really reaches it. In contrast, our method decreased the step size after sometime which allows better convergence to the optimum (see Figure 1(a)).

Training SVMs: We consider training SVMs which is a typical example where non-smooth SGD is heavily used [4]. For our experiments, we generate data as follows. Let $a_i \sim \mathcal{N}(0, \sigma^2 I_d)$ and the label $b_i = \text{sgn}(a_i(1) + z_i)$ where $z_i \sim \mathcal{N}(0, \eta^2)$. We generate $n = 500$ points in $d = 30$ dimensions. The SVM training problem is now:

$$F(x) = \frac{1}{n} \sum_{i=1}^n \max(0, 1 - b_i \langle a_i, x \rangle) + \frac{\lambda}{2} \|x\|^2, \tag{15}$$

where $\lambda = 0.1$. Since the objective is λ strongly convex, we consider step sizes of $\gamma_t := \frac{1}{\lambda t}$ for the standard method and the modified step sizes given in Equation (4) for our method. Figure 1 (b) plots loss during a typical run of SGD and Figure 1 (c) for the loss averaged over 100 independent runs of SGD for the same problem with the same initial point. The last point of SGD with modified step size sequence (Our Method) in blue consistently outperformed the standard SGD (Standard) in red. The green line denotes the loss of the average of the last $\frac{T}{4}$ iterates.

5 Conclusions and Discussion

We studied the fundamental question of sub-optimality of the last point of SGD/GD for general non-smooth convex functions as well as for strongly-convex functions. We proposed a novel step-size sequence that leads to information theoretically optimal rates in both the above mentioned settings. Our result proves a more general result for any “modified step-size” of a decaying standard step-size, and uses a novel technique of tracking best iterate in each time-interval and ensuring that the later iterates do not significantly deviate from the best iterate in the previous time interval. We also provide a high-probability bound using a super-martingale technique from [2]. Simulations show that our step-size indeed leads to better last point than the standard step-size sequences.

Our approach fundamentally exploits an assumption that we apriori know the total number of iterations T . Hence, our result does not provide an any-time algorithm. In contrast, existing any-time results have an extra $\log T$ multiplicative factor in the sub-optimality. We conjecture that this gap is fundamental and every *any-time* algorithm would suffer from the extra $\log T$ factor. We give lower bounds for the strongly convex case to show that for any choice of step sizes, the algorithm is either sub-optimal in expectation or almost surely so infinitely often.

References

- [1] Ohad Shamir and Tong Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning*, pages 71–79, 2013.
- [2] Nicholas JA Harvey, Christopher Liaw, Yaniv Plan, and Sikander Randhawa. Tight analyses for non-smooth stochastic gradient descent. *arXiv preprint arXiv:1812.05217*, 2018.
- [3] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- [4] Shai Shalev-Shwartz, Yoram Singer, Nathan Srebro, and Andrew Cotter. Pegasos: Primal estimated sub-gradient solver for svm. *Mathematical programming*, 127(1):3–30, 2011.
- [5] Takuya Akiba, Shuji Suzuki, and Keisuke Fukuda. Extremely large minibatch sgd: Training resnet-50 on imagenet in 15 minutes. *arXiv preprint arXiv:1711.04325*, 2017.
- [6] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- [7] Ohad Shamir. Open problem: Is averaging needed for strongly convex stochastic gradient descent? In *Conference on Learning Theory*, pages 47–1, 2012.

- [8] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [9] Arkadii Semenovich Nemirovsky and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- [10] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 928–936, 2003.
- [11] Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- [12] Sham M Kakade and Ambuj Tewari. On the generalization ability of online strongly convex programming algorithms. In *Advances in Neural Information Processing Systems*, pages 801–808, 2009.
- [13] Elad Hazan and Satyen Kale. Beyond the regret minimization barrier: optimal algorithms for stochastic strongly-convex optimization. *The Journal of Machine Learning Research*, 15(1):2489–2512, 2014.
- [14] Alexander Rakhlin, Ohad Shamir, Karthik Sridharan, et al. Making gradient descent optimal for strongly convex stochastic optimization. In *ICML*, volume 12, pages 1571–1578. Citeseer, 2012.
- [15] Simon Lacoste-Julien, Mark Schmidt, and Francis Bach. A simpler approach to obtaining an $o(1/t)$ convergence rate for the projected stochastic subgradient method. *arXiv preprint arXiv:1212.2002*, 2012.
- [16] Tong Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the twenty-first international conference on Machine learning*, page 116. ACM, 2004.

A Proofs of Technical Lemmas

A.1 Proof of Lemma 2

Proof of Lemma 2. We fix l such that $t_0 \leq l \leq t_1$. In this proof, we will freely use the fact that $\alpha_{t_0} \geq \alpha_t$ whenever $t_0 \leq t$. Let $l \leq t \leq t_1$. Define $\Delta_t = \langle \hat{g}_t(x_t) - g_t(x_t), x_t - x_l \rangle$. We note that x_t are random variables and are functions of $\hat{g}_1, \dots, \hat{g}_{t-1}$ only. We define the sigma-field $\mathcal{F}_t := \sigma(\hat{g}_1, \dots, \hat{g}_t)$.

We use the following notation for the sake of convenience: $D_t := \|x_t - x_l\|^2$. Clearly, D_t is \mathcal{F}_{t-1} measurable and Δ_t is \mathcal{F}_t measurable. It is clear from the definition of Δ_t that $\mathbb{E}[\Delta_t | \mathcal{F}_{t-1}] = 0$ and $|\Delta_t| \leq 2G\|x_t - x_l\| = 2G\sqrt{D_t}$.

By Hoeffding’s lemma, we conclude that for any $\mu \in \mathbb{R}$, we conclude:

$$\mathbb{E}[\exp(\mu\Delta_t) | \mathcal{F}_{t-1}] \leq \exp(2G^2 D_t \mu^2) \quad (16)$$

Let $\lambda = \frac{1}{8\alpha_{t_0}^2 G^2}$. For $t_0 \leq t \leq t_1$, consider

$$M_t := \exp\left(\sum_{s=l}^t -2\lambda L_s \alpha_s \Delta_s + \lambda(L_s - L_{s-1})D_s\right).$$

Clearly, M_t is \mathcal{F}_t measurable. $M_l = 1$ since $D_l = \Delta_l = 0$ almost surely. We will show that M_t is a super martingale:

$$\begin{aligned} \mathbb{E}[M_t | \mathcal{F}_{t-1}] &= M_{t-1} \mathbb{E}[\exp(-2\lambda L_t \alpha_t \Delta_t + \lambda(L_t - L_{t-1})D_t) | \mathcal{F}_{t-1}] \\ &= M_{t-1} \exp(-\lambda L_t^2 D_t) \mathbb{E}[\exp(-2\lambda L_t \alpha_t \Delta_t) | \mathcal{F}_{t-1}] \\ &\leq M_{t-1} \exp(-\lambda L_t^2 D_t + 8\lambda^2 L_t^2 \alpha_t^2 G^2 D_t) \\ &\leq M_{t-1} \end{aligned}$$

Therefore,

$$\mathbb{E}[M_{t_1}] \leq 1 \quad (17)$$

From the proof of Lemma 1, for $l \leq t \leq t_1$ we have:

$$\begin{aligned} \|x_{t+1} - x_l\|^2 &\leq \|x_t - x_l\|^2 + \alpha_t^2 \|\hat{g}_t(x_t)\|^2 - 2\alpha_t \langle \hat{g}_t(x_t), x_t - x_l \rangle \\ &\leq \|x_t - x_l\|^2 + \alpha_t^2 G^2 - 2\alpha_t \langle g_t(x_t), x_t - x_l \rangle - 2\alpha_t \langle \hat{g}_t(x_t) - g_t(x_t), x_t - x_l \rangle \\ &\leq \|x_t - x_l\|^2 + \alpha_t^2 G^2 + 2\alpha_t (F(x_l) - F(x_t)) - 2\alpha_t \Delta_t \end{aligned} \quad (18)$$

In the third step, we have used the convexity of $F(\cdot)$. Reordering Equation (18) and using the notation defined above:

$$2\alpha_t(F(x_t) - F(x_l)) - \alpha_t^2 G^2 \leq D_t - D_{t+1} - 2\alpha_t \Delta_t.$$

Multiplying the equation above by L_t and adding from $t = l$ to $t = t_1$, noting the fact that $D_l = 0$ and $D_{t_1+1} \geq 0$, we conclude:

$$\begin{aligned} \sum_{t=l}^{t_1} L_t [2\alpha_t(F(x_t) - F(x_l)) - \alpha_t^2 G^2] &\leq \sum_{t=l}^{t_1} L_t (D_t - D_{t+1} - 2\alpha_t \Delta_t) \\ &\leq \sum_{t=l}^{t_1} -2L_t \alpha_t \Delta_t + (L_t - L_{t-1}) D_t \end{aligned} \quad (19)$$

We recall the random variable

$$A(l, t_1) := \sum_{t=l}^{t_1} L_t [2\alpha_t(F(x_t) - F(x_l)) - \alpha_t^2 G^2]$$

From equations (17) and (19), we conclude that for every l such that $t_0 \leq l \leq t_1$:

$$\mathbb{E}[\exp(\lambda A(l, t_1))] \leq \mathbb{E}[M_t] \leq 1.$$

By convexity of the exponential function, we have:

$$\mathbb{E}[\exp(\lambda(p.A)(t_0, t_1))] \leq \mathbb{E} \sum_{l=t_0}^{t_1} p(l) \mathbb{E}[\exp(\lambda A(l, t_1))] \leq 1$$

By Chernoff Bound, we conclude:

$$\mathbb{P}[(p.A)(t_0, t_1) > \eta] \leq \exp\left(-\frac{\eta}{8\alpha_{t_0}^2 G^2}\right)$$

The case for $A^*(t_0, t_1)$ proceeds similarly but this time we use x^* in place of x_{t_0} . We define $D_t^* := \|x_t - x^*\|^2$, $\Delta_t^* := \langle \hat{g}_t(x_t) - g_t(x_t), x_t - x^* \rangle$ and

$$M_t^* := \exp\left(\sum_{i=t_0}^t -2\lambda L_i \alpha_i \Delta_i^* + \lambda(L_i - L_{i-1}) D_i^*\right)$$

We note that for $t_0 < t \leq t_1$, $\mathbb{E}[M_t^* | \mathcal{F}_{t-1}] \leq M_{t-1}^*$ and $D_{t_0}^* \leq D^2$. Therefore,

$$\mathbb{E} M_{t_1}^* \leq \mathbb{E} M_1^* \leq \exp\left(\frac{L_{t_0}^2 D^2 + L_{t_0} D^2}{8\alpha_{t_0}^2 G^2}\right) \leq \exp\left(\frac{2L_{t_0} D^2}{8\alpha_{t_0}^2 G^2}\right)$$

Here we have used the fact that $L_{t_0} \leq \frac{1}{t_0 - t_1 + 1} \leq 1$. Noting that $\exp(\lambda A^*(t_0, t_1)) \leq M_{t_1}^*$ we use Chernoff bound to conclude the result. \square

A.2 Proof of Lemma 3

Proof of Lemma 3. We prove this by induction. The assertion is true for $i = 0$. Suppose it is true for $i = k \leq r - 1$. Then,

$$\begin{aligned} \lambda_{k+1} &= \lambda_k (1 + \Gamma \lambda_k) \\ &\leq \left(1 + \frac{1}{r}\right)^k \lambda_0 \left(1 + \frac{\left(1 + \frac{1}{r}\right)^k}{re}\right) \\ &\leq \left(1 + \frac{1}{r}\right)^k \lambda_0 \left(1 + \frac{\left(1 + \frac{1}{r}\right)^k}{re}\right) \\ &\leq \left(1 + \frac{1}{r}\right)^{k+1} \lambda_0 \end{aligned}$$

The we have proved the assertion through induction. \square

A.3 Proof of Lemma 7

Proof of Lemma 7. We take the definitions of the terms used from the proof of Lemma 6. It is clear from the definition that $\kappa(t) \geq 0$. Since $\kappa(t) = 0$ for $t > T_{i+1}$, it is sufficient to show that $\sum_{s=T_{i+1}}^{T_{i+1}} \kappa(s) = 1$.

We define $\sigma(t) = \sum_{s=T_{i+1}}^t \kappa(s)$ (an empty sum denotes 0). By definition of κ , for $T_i + 1 \leq t \leq T_{i+1}$

$$\begin{aligned} \sigma(t) &= \frac{\Gamma(T_{i+1})}{\Gamma(t)} q^{(i)}(t) + \left(1 + \frac{\alpha_t L_t}{\Gamma(t)}\right) \sigma(t-1) \\ &= \frac{\Gamma(T_{i+1})}{\Gamma(t)} q^{(i)}(t) + \frac{\Gamma(t-1)}{\Gamma(t)} \sigma(t-1) \end{aligned}$$

Continuing the above recursion, we conclude:

$$\sigma(t) = \frac{\Gamma(T_{i+1})}{\Gamma(t)} \sum_{s=T_{i+1}}^t q^{(i)}(s)$$

Since q is a probability distribution over $\{T_i + 1, \dots, T_i\}$, we conclude

$$\sigma(T_{i+1}) = \sum_{s=T_{i+1}}^{T_{i+1}} q^{(i)}(s) = 1$$

□

B Proofs of Lower Bounds

We will prove theorem 5 for $G = 5$ and $\mu = 1$ for the sake of convenience. We can handle the general case by considering the transformation $F_0(x) = \frac{25\mu}{G^2} F(\frac{G}{2\mu}x)$. We scale the domain as $D_0 := \frac{5\mu}{G}D$. If F is μ strongly convex and G Lipschitz, then F_0 is 1 strongly convex and 5 Lipschitz. We take the subgradient oracle for F_0 to be $\hat{g}_t^0(x) := \frac{5}{G} \hat{g}_t(\frac{Gx}{5\mu})$. It is easy to check that if SGD for $F(\cdot)$ with step sizes α_t , the iterates are x_t , then starting from $x_0^0 := \frac{5\mu}{G}x_0$ and using step sizes $\alpha_t^0 := \mu\alpha_t$ and the subgradient oracle defined above, the iterates for F_0 is $x_t^0 = \frac{5\mu}{G}x_t$. Therefore, $F_0(x_t^0) = \frac{25\mu}{G^2}F(x_t)$ and the proof below goes through seamlessly. This is similar to the rescaling used for the lowerbounds in [2].

Without loss of generality, we will restrict our attention to strictly positive step size sequences: $\gamma_t > 0$. We further restrict the possible values of γ_t in the following lemma:

Lemma 8. *If the step size sequence γ_t is such that there is an infinite sequence of times t_k such that $\lim_{k \rightarrow \infty} t_k \gamma_{t_k} = \infty$, then SGD is bad in expectation. Therefore, we can restrict our consideration to step size sequences of the form $\gamma_t = O(\frac{1}{t})$.*

Proof. Consider the function $F : [-1, 1] \rightarrow \mathbb{R}$ defined by $F(x) = |x| + \frac{x^2}{2}$. F has a global optimum at $x = 0$ and it is 1 strongly convex. Let ϵ_t be a sequence of i.i.d. rademacher random variables (i.e, uniform over $\{-1, 1\}$). We let the subgradient oracle to return $\hat{g}_t(x) = \text{sgn}(x) + x + 3\epsilon_t$. Clearly,

$$|x_{t+1}| = \min(|x_t - \gamma_t(\text{sgn}(x_t) + x_t + 3\epsilon_t)|, 1).$$

ϵ_t is independent of x_t and conditioned on the value of x_t , with probability atleast $\frac{1}{2}$, ϵ_t has the opposite sign as x_t . When this happens, $(\text{sgn}(x_t) + x_t + 3\epsilon_t)$ has the opposite sign of x_t and $|(\text{sgn}(x_t) + x_t + 3\epsilon_t)| \geq 1$. Therefore under this event, $|x_t - \gamma_t(\text{sgn}(x_t) + x_t + 3\epsilon_t)| \geq |x_t| + \gamma_t \geq \gamma_t$.

Therefore, we conclude:

$$\mathbb{E}|x_{t+1}| \geq \frac{1}{2} \min(1, \gamma_t).$$

Considering the fact that $(t_k + 1)\mathbb{E}(F(x_{t_k+1}) - F(0)) \geq t_k \mathbb{E}|x_{t_k+1}| \geq \frac{1}{2} \min(t_k, \gamma_{t_k} t_k) \rightarrow \infty$, we conclude that SGD with this step size is bad in expectation.

□

Henceforth, we will restrict our attention without loss of generality to step size sequences such that $\gamma_t = O(\frac{1}{t})$. We will first consider the function $F_1(x) = \frac{1}{2}x^2$ over the set $[-1, 1]$. Let the infinite horizon learning rate be $(\gamma_t)_{t \in \mathbb{N}}$ at each time instant, the subgradient oracle returns $x + \epsilon_t$ where ϵ_t is a sequence of i.i.d. uniform random variable over $\{-1, 1\}$ (that is rademacher random variables). Let the iterates of SGD be z_t and $z_1 = 1$.

Lemma 9. *Let T_0 be the smallest time such that $\gamma_t < 1$ for all $t \geq T_0$. Then, for every $t \geq T_0$*

1.

$$\mathbb{E}|z_{t+1}|^2 = (1 - \gamma_t)^2 \mathbb{E}z_t^2 + \gamma_t^2$$

2.

$$\mathbb{E}\|z_t\|^2 \geq \frac{1}{t - T_0 + 1}$$

Proof. Suppose $\gamma_t \geq 1$. Then:

$$|z_t(1 - \gamma_t) + \gamma_t \epsilon_t| \geq \gamma_t |\epsilon_t| - |z_t(1 - \gamma_t)| \geq \gamma_t - |1 - \gamma_t| = 1$$

Therefore, when $\gamma_t \geq 1$, $|z_{t+1}| = 1$. Therefore, $z_{T_0} = 1$ almost surely. When $t \geq T_0$,

$$\begin{aligned} |z_t - \gamma_t(z_t + \epsilon_t)| &\leq |(1 - \gamma_t)z_t| + \gamma_t |\epsilon_t| \\ &= (1 - \gamma_t)|z_t| + \gamma_t \leq 1 \end{aligned}$$

Therefore, when $t \geq T_0$, the iteration of SGD won't leave the set $[-1, 1]$ almost surely, so there is no need for the projection step to obtain the next iterate. That is, for $t \geq T_0$, $z_{t+1} = z_t(1 - \gamma_t) + \epsilon_t \gamma_t$. Squaring and taking expectations, we conclude:

$$\begin{aligned} \mathbb{E}|z_{t+1}|^2 &= (1 - \gamma_t)^2 \mathbb{E}z_t^2 + \gamma_t^2 \\ &\geq \inf_{\gamma \in \mathbb{R}} (1 - \gamma)^2 \mathbb{E}z_t^2 + \gamma^2 \\ &= \frac{\mathbb{E}z_t^2}{1 + \mathbb{E}z_t^2} \end{aligned}$$

Clearly, $\mathbb{E}|z_{T_0}|^2 = 1 = \frac{1}{T_0 - T_0 + 1}$. Using induction in the equation above, we conclude: $\mathbb{E}|z_t|^2 \geq \frac{1}{t - T_0 + 1}$ for every $t \geq T_0$. □

We divide \mathbb{N} into time intervals of the form $\{2^k, 2^k + 1, \dots, 2^{k+1} - 1\} := I_k$. We have the following lemma:

Lemma 10. *If $\gamma_t \leq \frac{C}{t}$ for some constant $C \geq 0$ and there exist positive infinite sequences c_k and d_k such that $\lim_{k \rightarrow \infty} c_k = \infty$, $\lim_{k \rightarrow \infty} d_k = 0$ and every k , either one of the two conditions below hold:*

1. $\sum_{t \in I_k} \gamma_t^2 \geq c_k 2^{-k} (\sum_{t \in I_k} \gamma_t)^2$
2. $\sum_{t \in I_k} \gamma_t \leq d_k$

Then, SGD with step size γ_t is bad in expectation.

Proof. We consider the optimization problem considered in Lemma 9 i.e, optimizing $F(x) = x^2$. Let $T_k = 2^k$. We assume the contrary - that is, $\mathbb{E}|z_{T_k}|^2 \leq \frac{L}{T_k}$ for every k , for some $L > 0$. As shown in the second inequality of Lemma 9, irrespective of the choice of $\gamma_1, \dots, \gamma_{T_k-1}$,

$$\mathbb{E}|z_{T_k}^2| \geq \frac{1}{T_k - 1}.$$

From the first equality in Lemma 9, we conclude that for $t \in I_k$, $\mathbb{E}|z_{t+1}|^2 = (1 - \gamma_t)^2 \mathbb{E}z_t^2 + \gamma_t^2$. Since $\gamma_t \leq \frac{C}{t}$, we can take k large enough so that $\gamma_t \leq \frac{1}{2}$ for every $t \in I_k$. Using the fact that $(1 - \gamma_t) \geq \exp -\frac{\gamma_t}{1 - \gamma_t} \geq \exp -2\gamma_t$. Therefore,

$$\mathbb{E}|z_{t+1}|^2 \geq e^{-4\gamma_t} \mathbb{E}|z_t|^2 + \gamma_t^2$$

Unravelling the recursion above, we conclude:

$$\mathbb{E}|z_{T_{k+1}}|^2 \geq e^{-4 \sum_{t \in I_k} \gamma_t} \left[\mathbb{E}|z_{T_k}|^2 + \sum_{t \in I_k} \gamma_t^2 \right] \quad (20)$$

We define $S_k := \sum_{t \in I_k} \gamma_t$.

1. Suppose for a particular k , the first item in the statement of the lemma holds

By assumption, $\sum_{t \in I_k} \gamma_t^2 \geq \frac{c_k}{T_k} S_k^2$. Using this in Equation (20), we conclude:

$$\mathbb{E}|z_{T_{k+1}}|^2 \geq e^{-4S_k} \left[\mathbb{E}|z_{T_k}|^2 + \frac{c_k}{T_k} S_k^2 \right]$$

Now, since $\gamma_t \leq \frac{C}{t}$, we have $S_k \leq C$. Therefore,

$$\begin{aligned} \mathbb{E}|z_{T_{k+1}}|^2 &\geq e^{-4S_k} \mathbb{E}|z_{T_k}|^2 + \frac{c_k}{T_k} S_k^2 e^{-4C} \\ &= \mathbb{E}|z_{T_k}|^2 \left[e^{-4S_k} + \frac{c_k S_k^2 e^{-4C}}{T_k \mathbb{E}|z_{T_k}|^2} \right] \\ &\geq \mathbb{E}|z_{T_k}|^2 \left[e^{-4S_k} + \frac{c_k S_k^2 e^{-4C}}{L} \right] \\ &\geq \mathbb{E}|z_{T_k}|^2 \inf_{x \geq 0} \left[e^{-4x} + \frac{x^2 c_k e^{-4C}}{L} \right] \end{aligned} \quad (21)$$

In the third step, we have used the fact that $T_k \mathbb{E}|z_{T_k}|^2 \leq L$. We now consider the function $h : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ given by $h(x) = e^{-2x} + \kappa x^2$ for some $\kappa > 0$. Clearly, h is convex, bounded below and tends to infinity as $x \rightarrow \infty$. Therefore, it has a unique minimizer t^* - the unique point such that $h'(t^*) = 0$. That is, t^* is the unique point which satisfies: $\kappa t^* = 2e^{-4t^*} \leq 2$. Therefore, $t^* \leq \frac{2}{\kappa}$. Therefore, $h(t^*) \geq e^{-4t^*} \geq e^{-8/\kappa} \geq 1 - \frac{8}{\kappa}$. In Equation (23), we take $\kappa = \frac{c_k e^{-4C}}{L}$ we conclude:

$$\mathbb{E}|z_{T_{k+1}}|^2 \geq \mathbb{E}|z_{T_k}|^2 \left(1 - \frac{C'}{c_k} \right)$$

Where C' is a constant depending only on L and C .

2. Suppose for a particular k , the second item in the statement of the lemma holds: Then, by Equation (20), we have:

$$\mathbb{E}|z_{T_{k+1}}|^2 \geq e^{-4d_k} \mathbb{E}|z_{T_k}|^2 \geq (1 - 4d_k) \mathbb{E}|z_{T_k}|^2 \quad (22)$$

From Equations (21) and (22), we conclude that there exists an absolute constant \bar{C} depending only on C and L such that:

$$\mathbb{E}|z_{T_{k+1}}|^2 \geq \left(1 - \bar{C} \max \left(d_k, \frac{1}{c_k} \right) \right) \mathbb{E}|z_{T_k}|^2 \quad (23)$$

Since $\max \left(d_k, \frac{1}{c_k} \right) \rightarrow 0$, we can choose k large enough so that $\sup_{s > k} \bar{C} \max \left(d_k, \frac{1}{c_k} \right) \leq 1 - e^{-\epsilon}$ for arbitrary $\epsilon > 0$.

From Equation (23), it follows that for arbitrary $K \in \mathbb{N}$,

$$\mathbb{E}|z_{T_{k+K}}|^2 \geq e^{-\epsilon K} \mathbb{E}|z_{T_k}|^2$$

By Lemma 9, $\mathbb{E}|z_{T_k}|^2 \geq \frac{1}{T_k-1} \geq 2^{-k}$. By our assumption, $\mathbb{E}|z_{T_{k+K}}|^2 \leq L2^{-k-K}$. Therefore, we conclude: $L2^{-k-K} \geq e^{-\epsilon K} 2^{-k}$ for every $K \in \mathbb{N}$. This cannot hold for any finite L when we take $\epsilon < \log 2$. This contradicts our assumption. Therefore, SGD with step size γ_t is bad in expectation. \square

We will show that if conditions for γ_t in Lemma 9 or those in Lemma 10 don't hold, then SGD is bad almost surely. We recall the definition of the interval $I_k = \{2^k + 1, \dots, 2^{k+1}\}$. We prove the following lemma to inspect how frequently long, contiguous segments of ϵ_t are all equal to 1 for $t \in I_k$. We take $\tau_k := 2^{\lfloor \log_2(k/2) \rfloor}$. We note that $\frac{k}{4} \leq \tau_k \leq \frac{k}{2}$. We can divide I_k into $|I_k|/\tau_k$ contiguous, disjoint intervals, each of size τ_k . We call these intervals $J_k(i)$ for $i \in \{1, \dots, |I_k|/\tau_k\}$. We let A_k to be the event that for some $i \in \{1, \dots, |I_k|/\tau_k\}$, $\epsilon_t = 1$ for all $t \in J_k(i)$. In particular, the even A_k implies that there is a contiguous τ_k length sequence of ϵ_t of all 1s in I_k .

Lemma 11.

$\mathbb{P}(A_k^c) \leq Ck2^{-k/2}$ for some absolute constant C .

Proof. We subdivide the interval I_k into disjoint subintervals of length τ_k . There are $\frac{2^k}{\tau_k}$ such intervals. The event A_k holds if over one such subinterval, the random signs are all 1. The probability of a given subinterval having all signs equal to 1 is $p_{\tau_k} := \frac{1}{2^{\tau_k}}$. Therefore, we conclude:

$$\mathbb{P}(A_k) = 1 - (1 - p_{\tau_k})^{\lfloor \frac{2^k}{\tau_k} \rfloor} \geq 1 - e^{-p_{\tau_k} \lfloor \frac{2^k}{\tau_k} \rfloor} \geq 1 - \frac{1}{e p_{\tau_k} \lfloor \frac{2^k}{\tau_k} \rfloor}.$$

Here we have used the inequality $xe^{-x} \leq \frac{1}{e}$ for $x > 0$.

Therefore, we conclude that: $\mathbb{P}(A_k^c) \leq Ck2^{-k/2}$ for some absolute constant C . \square

We now consider the same function which was considered in Lemma 8 i.e, $F : [-1, 1] \rightarrow \mathbb{R}$ defined by $F(x) = |x| + \frac{x^2}{2}$. F has a global optimum at $x = 0$ and it is 1 strongly convex. Let ϵ_t be a sequence of i.i.d. rademacher random variables (i.e, uniform over $\{-1, 1\}$). We let the subgradient oracle to return $\hat{g}_t(x) = \text{sgn}(x) + x + 3\epsilon_t$. Let the iterates of SGD for F with step sizes γ_t be y_t .

Lemma 12. Suppose $\gamma_t \leq \frac{C}{t}$, there exists an infinite sequence $(k_r)_{r \in \mathbb{N}}$ and fixed constants $c_0, d_0 > 0$ such that both the conditions hold:

1.

$$\sum_{t \in I_{k_r}} \gamma_t^2 \leq c_0 2^{-k_r} \left(\sum_{t \in I_{k_r}} \gamma_t \right)^2$$

2.

$$\sum_{t \in I_{k_r}} \gamma_t \geq d_0$$

We note that these conditions are the negations of the conditions for γ_t in Lemma 8 and Lemma 10. Then SGD with step size γ_t is bad almost surely.

Proof. We will show that there exists a sequence of independent events B_{k_r} for $r \in \mathbb{N}$ such that $\mathbb{P}(B_{k_r}) \geq p_0 > 0$ uniformly and whenever B_{k_r} holds,

$$\max_{t \in I_{k_r}} \frac{t}{\log t} [F(y_t) - F(0)] \geq \delta_0$$

For some constant $\delta_0 > 0$. We note that p_0 and δ_0 depend only on C, d_0 and c_0 . We consider a random times $T_{\max}, T_{\min} \in I_k$ as follows:

1. If the event A_k^c holds, pick a uniformly random element i_0 from $\{1, \dots, |I_k|/\tau_k\}$ independent of everything else. Set $T_{\max} := \max J_k(i_0)$ and $T_{\min} := \min J_k(i_0)$

2. If the event A_k holds, pick a uniformly random element i_0 from $\{i : \text{for all } t \in J_k(i), \epsilon_t = 1\}$, independent of everything else. Set $T_{\max} := \max J_k(i_0)$ and $T_{\min} := \min J_k(i_0)$

We note that by symmetry, i_0 is uniformly distributed over the set $\{1, \dots, |I_k|/\tau_k\}$. We will show that, when the event A_k holds, then one of the following is true:

1. $y(T_{\max}) = -1$.
2. $y(T_{\min}) - y(T_{\max}) \geq \sum_{t \in J_k(i_0)} \gamma_t$

Suppose the event A_k holds. Then for $T_{\min} \leq t < T_{\max}$, $y_{t+1} = \max(y_t - \gamma_t(y_t + \text{sgn}(y_t) + 3), -1)$. Since under the event A_k , $\epsilon_t = 1$ for every $t \in J_k(i_0)$, we conclude that $\gamma_t(y_t + \text{sgn}(y_t) + 3) \geq \gamma_t$. That is SGD drifts in the negative direction irrespective of the value of the iterate. It is therefore clear that if for some $T_{\min} \leq t \leq T_{\max}$, y_t hits -1 , then $y(T_{\max}) = -1$. Now suppose that for $y_t > -1$ for every t in this range. Then, $y_{t+1} \leq y_t - \gamma_t$. But unraveling this recursion, it follows that $y(T_{\min}) - y(T_{\max}) \geq \sum_{t \in J_k(i_0)} \gamma_t$. Therefore, it follows that when the event A_k holds:

$$\begin{aligned} \max_{t \in I_k} F(y_t) &\geq \max(F(y(T_{\max})), F(y(T_{\min}))) \\ &\geq \max(|y(T_{\max})|, |y(T_{\min})|) \\ &\geq \min(1, \frac{1}{2} \sum_{t \in J_k(i_0)} \gamma_t) \end{aligned} \tag{24}$$

It is clear that since $\gamma_t = C/t$, for k large enough, $\frac{1}{2} \sum_{t \in J_k(i_0)} \gamma_t \leq 1$. Therefore, we conclude that for k large enough, when the event A_k holds,

$$\max_{t \in I_k} F(y_t) \geq \frac{1}{2} \sum_{t \in J_k(i_0)} \gamma_t.$$

Fix $0 < \beta < 1$. We now consider E_k to be the event $\{\sum_{t \in J_k(i_0)} \gamma_t \geq \frac{\beta \tau}{|I_k|} \sum_{t \in I_k} \gamma_t\}$.

By symmetry, i_0 is uniformly distributed over $\{1, \dots, |I_k|/\tau_k\}$. Therefore,

$$\mathbb{E} \sum_{t \in J_k(i_0)} \gamma_t = \frac{\tau_k}{|I_k|} \sum_{t \in I_k} \gamma_t$$

and

$$\begin{aligned} \mathbb{E} \left(\sum_{t \in J_k(i_0)} \gamma_t \right)^2 &= \frac{\tau_k}{|I_k|} \sum_{i=1}^{|I_k|/\tau_k} \sum_{t, s \in J_k(i_0)} \gamma_t \gamma_s \\ &\leq \frac{\tau_k}{|I_k|} \sum_{i=1}^{|I_k|/\tau_k} \sum_{t, s \in J_k(i_0)} \frac{\gamma_t^2 + \gamma_s^2}{2} \\ &= \frac{\tau_k^2}{|I_k|} \sum_{t \in I_k} \gamma_t^2 \end{aligned}$$

Now, when k is part of the infinite sequence (k_r) , by assumption we have:

$$\mathbb{E} \left(\sum_{t \in J_k(i_0)} \gamma_t \right)^2 \leq c_0 \frac{\tau_k^2}{|I_k|^2} \left(\sum_{t \in I_k} \gamma_t \right)^2$$

Therefore, by Payley-Zigmund inequality, whenever k is part of the infinite sequence (k_r) , for every $\beta < 1$,

$$\mathbb{P} \left[\sum_{t \in J_k(i_0)} \gamma_t \geq \beta \frac{\tau}{|I_k|} \sum_{t \in I_k} \gamma_t \right] \geq (1 - \beta)^2 \frac{(\mathbb{E} \sum_{t \in J_k(i_0)} \gamma_t)^2}{\mathbb{E} (\sum_{t \in J_k(i_0)} \gamma_t)^2} \geq \frac{(1 - \beta)^2}{c_0}$$

Recalling the definition of E_k , we conclude, $\mathbb{P}(E_{k_r}) \geq \frac{(1-\beta)^2}{c_0}$.

We will now define the event $B_k := E_k \cap A_k$. The events B_k are all independent by definition. When the event B_k holds, clearly, from equation (24), we conclude:

$$\max_{t \in I_{k_r}} F(y_t) \geq \frac{1}{2} \sum_{t \in J_{k_r}(i_0)} \gamma_t \geq \frac{\beta \tau_{k_r}}{2|I_{k_r}|} \sum_{t \in I_{k_r}} \gamma_t \geq \frac{\beta \tau_{k_r} d_0}{2|I_{k_r}|}$$

The second inequality follows from the definition of E_k . Using the fact that any $t \in I_{k_r}$ is such that $t \leq 2|I_{k_r}|$ and $\tau_{k_r} = \Theta(k_r)$, we conclude that for some $\delta_0 > 0$, fixed, the following holds whenever the event B_{k_r} holds.

$$\max_{t \in I_{k_r}} \frac{t}{\log t} [F(y_t) - F(0)] \geq \delta_0 \quad (25)$$

$$\begin{aligned} \mathbb{P}(B_{k_r}) &\geq \mathbb{P}(E_{k_r}) - \mathbb{P}(A_{k_r}) \\ &\geq \frac{(1 - \beta)^2}{c_0} - O(k_r 2^{-k_r/2}) \end{aligned}$$

It is clear that we can find a $p_0 > 0$ such that for all k_r large enough, $\mathbb{P}(B_{k_r}) > p_0$.

Since B_{k_r} are independent sets, it follows that infinitely many of them are true with probability 1. From equation (25), we conclude that SGD with step sizes γ_t is bad almost surely. \square

Proof of Theorem 5. We will conclude this from Lemmas 8, 10 and 12. Therefore, it is sufficient to show that any strictly positive infinite sequence γ_t is such that atleast one of the following condition holds

1. There is an infinite sequence of times t_k such that $\lim_{k \rightarrow \infty} t_k \gamma_{t_k} = \infty$. In this case, by Lemma 8, we conclude that it is bad in expectation.
2. There exists a C such that $\gamma_t \leq \frac{C}{t}$ and there exist infinite sequences $c_k \rightarrow \infty$ and $d_k \rightarrow 0$ such that for every k , either $\sum_{t \in I_k} \gamma_t^2 \geq c_k 2^{-k} (\sum_{t \in I_k} \gamma_t)^2$ or $\sum_{t \in I_k} \gamma_t \leq d_k$. In this case, by Lemma 10, we conclude that it is bad in expectation.
3. There exists a C such that $\gamma_t \leq \frac{C}{t}$ and there exist fixed positive constants c_0 and d_0 such that for some infinite sub-sequence (k_r) , $\sum_{t \in I_{k_r}} \gamma_t^2 \leq c_0 2^{-k_r} (\sum_{t \in I_{k_r}} \gamma_t)^2$ and $\sum_{t \in I_{k_r}} \gamma_t \geq d_0$. In this case, by Lemma 12, we conclude that the algorithm is bad almost surely.

It is therefore sufficient to show that if conditions 1 and 2 don't hold then condition 3 holds. The negation of condition 1 is that $\gamma_t \leq \frac{C}{t}$ for some $C > 0$. Now, we denote by

$$\eta_k := 2^k \frac{\sum_{t \in I_k} \gamma_t^2}{(\sum_{t \in I_k} \gamma_t)^2}$$

and

$$\lambda_k := \sum_{t \in I_k} \gamma_t$$

. Therefore, $\eta_k \geq c_k$ or $\lambda_k \leq d_k$ for some $c_k \rightarrow \infty$ and $d_k \rightarrow 0$ is equivalent to $\eta_k + \frac{1}{\lambda_k} \rightarrow \infty$ which is equivalent to the statement that for every subsequence k_r , $\eta_{k_r} + \frac{1}{\lambda_{k_r}} \rightarrow \infty$. Therefore the negation of condition 2 is equivalent to atleast one of the following conditions being true

1. There exists infinite sequence (t_k) such that $t_k \gamma_{t_k} \rightarrow \infty$

2. There exists an infinite subsequence k_r such that $\eta_{k_r} + \frac{1}{\lambda_{k_r}} \leq M$ for some $M > 0$. That is, $\eta_{k_r} \leq M := c_0$ and $\lambda_{k_r} \geq \frac{1}{M} := d_0$

Therefore we conclude that when neither of the conditions 1 and 2 hold, then condition 3 holds. This proves our result. \square