
A High Probability Analysis of Adaptive SGD with Momentum

Xiaoyu Li¹ Francesco Orabona²

Abstract

Stochastic Gradient Descent (SGD) and its variants are the most used algorithms in machine learning applications. In particular, SGD with adaptive learning rates and momentum is the industry standard to train deep networks. Despite the enormous success of these methods, our theoretical understanding of these variants in the non-convex setting is not complete, with most of the results only proving convergence in expectation and with strong assumptions on the stochastic gradients. In this paper, we present a high probability analysis for adaptive and momentum algorithms, under weak assumptions on the function, stochastic gradients, and learning rates. We use it to prove for the first time the convergence of the gradients to zero in high probability in the smooth nonconvex setting for Delayed AdaGrad with momentum.

1. Introduction

Despite the incredible popularity of stochastic gradient methods in practical machine learning applications, our theoretical understanding of these methods is still not complete. In particular, *adaptive learning rates* methods like AdaGrad (Duchi et al., 2011) have been mainly studied in the convex domain, with few analyses in the non-convex domain (Li & Orabona, 2019; Ward et al., 2019). However, even in these latter analyses, the assumptions used are very strong and/or the results limited.

In particular, there are two main problems with the previous analyses of Stochastic Gradient Descent (SGD) and its variants in the nonconvex setting. First, the classic analysis of convergence for SGD in the nonconvex setting uses an analysis in expectation. However, expectation bounds do not rule out extremely bad outcomes. As pointed out by Harvey

et al. (2019a), it is a misconception that for the algorithms who have expectation bounds it is enough to pick the best of several independent runs to have a high probability guarantee: It can actually be a computationally inefficient procedure. Moreover, in practical applications like deep learning, it is often the case that only one run of the algorithm is used since that the training process may take long time. Hence, it is essential to get high probability bounds which guarantee the performance of the algorithm on single runs.

Another very common assumption used in most of the previous papers is the one of bounded stochastic gradients. This is a rather strong assumption and it is false even in the deterministic optimization of a convex quadratic function, e.g., $f(x) = x^2$.

In this work, we overcome both these problems. We prove *high probability convergence rates only assuming that the noises on the gradients are well-behaved*, i.e., subgaussian. In this way, we allow for unbounded gradients and unbounded noise. The weak assumptions, the nonconvex analysis, and the adaptive learning rates make our results particularly challenging to obtain. Indeed, high probability bounds for bounded stochastic gradients are almost trivial to obtain but of limited applicability. Overall, we believe this paper is the first one to prove such guarantees.

Contributions. In this short paper, we present a high probability analysis of SGD with momentum and adaptive learning rates, with weak assumptions on function and stochastic gradients. So, first in Theorem 1 we prove high probability bounds for the gradients of classic momentum SGD step size $O(\frac{1}{\sqrt{t}})$ in the nonconvex setting. Then, in Theorem 2 we prove *for the first time* high probability convergence rates for the gradients of AdaGrad with momentum in the nonconvex setting. In particular, we also show that the high probability bounds are adaptive to the level of noise.

2. Related Work

Stochastic momentum methods. Sutskever et al. (2013) discussed the importance of classic momentum methods in deep learning, which is nowadays widely used in the training of neural networks. On the convergence of stochastic momentum methods, Yang et al. (2016) studied a unified momentum method and provided expectation bounds in

¹Division of System Engineering, Boston University, MA, USA ²Electrical & Computer Engineering, Boston University, MA, USA. Correspondence to: Xiaoyu Li <xiaoyuli@bu.edu>.

the rate of $O(\frac{1}{\sqrt{T}})$ in both convex and nonconvex setting. However, the results hold only for Lipschitz functions. [Gadat et al. \(2018\)](#) provided an in-depth description of the stochastic heavy-ball method. Moreover, for the non-convex functions, they showed some almost sure convergence results. [Loizou & Richtárik \(2017\)](#) provided a general analysis for the momentum variants of several classes of stochastic optimization algorithms and proved the linear rate convergence for quadratic and smooth functions. To the best of our knowledge, there are no high probability bounds for nonconvex stochastic momentum methods without using strong assumptions.

Nonconvex convergence of adaptive methods. In recent years, a variety of adaptive SGD algorithms have been developed to automatically tune the step size by using the past stochastic gradients. The first adaptive algorithm was AdaGrad [Duchi et al. \(2011\)](#), designed to adapt to sparse gradients. [Li & Orabona \(2019\)](#) and [Ward et al. \(2019\)](#) showed the convergence of variants of AdaGrad with a rate of $O(\ln T/\sqrt{T})$ in the non-convex case. Moreover, [Li & Orabona \(2019\)](#) showed that AdaGrad with non-coordinate-wise learning rates is adaptive to the level of noise. [Zou et al. \(2019\)](#) studied AdaGrad with a unified momentum and [Chen et al. \(2019\)](#) considered a large family of Adam-like algorithms ([Kingma & Ba, 2015](#)) including AdaGrad with momentum. Yet, all of these works prove on bounds in expectation and most of them use the very strong assumption of bounded stochastic gradients.

High probability bounds. The results on high probability bounds are relatively rare compared to those in expectation, which are easier to obtain. [Kakade & Tewari \(2009\)](#) used Freeman’s inequality to prove high probability bounds for an algorithm solving the SVM objective function. For classic SGD, [Harvey et al. \(2019b\)](#) and [Harvey et al. \(2019a\)](#) used a generalized Freedman’s inequality to prove bounds in non-smooth and strongly convex case, while [Jain et al. \(2019\)](#) proved the optimal bound for the last iterate of SGD with high probability. As far as we know, there are currently no high probability bounds for adaptive methods in the nonconvex setting.

3. Problem Set-Up

Notation. We denote vectors and matrices by bold letters. The coordinate j of a vector \mathbf{x} is denoted by x_j and as $\nabla f(\mathbf{x})_j$ for the gradient $\nabla f(\mathbf{x})$. To keep the notation concise, all standard operations $\mathbf{x}\mathbf{y}$, \mathbf{x}/\mathbf{y} , \mathbf{x}^2 , $1/\mathbf{x}$, $\mathbf{x}^{1/2}$ and $\max(\mathbf{x}, \mathbf{y})$ on the vectors \mathbf{x} , \mathbf{y} are supposed to be element-wise. We denote by $\mathbb{E}[\cdot]$ the expectation with respect to the underlying probability space and by $\mathbb{E}_t[\cdot]$ the conditional expectation with respect to the past randomness. We use L_2 norms.

Assumptions. In this paper we focus on the optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}),$$

where f is bounded from below and we denote its infimum by f^* . We *do not* assume the function to be convex. We consider stochastic optimization algorithms that have access to a noisy estimate of the gradient of f . This covers the ubiquitous SGD ([Robbins & Monro, 1951](#)), as well modern variants as AdaGrad. We are interested in studying the convergence of the gradients to zero, because without additional assumptions it is the only thing we can study in the nonconvex setting.

We make the following assumption on the objective function $f(\mathbf{x})$:

- (A) f is M -smooth, that is, f is differentiable and its gradient is M -Lipschitz, i.e. $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq M\|\mathbf{x} - \mathbf{y}\|$, $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d$.

Note that (A), for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, implies ([Nesterov, 2004](#), Lemma 1.2.3)

$$|f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle| \leq \frac{M}{2} \|\mathbf{y} - \mathbf{x}\|^2.$$

It is easy to see that this assumption is necessary to have the convergence of the gradients to zero. Indeed, without smoothness the norm of the gradients does not go to zero even in the convex case, e.g., consider the function $f(x) = |x - 1|$.

We assume that we have access to a stochastic first-order black-box oracle, that returns a noisy estimate of the gradient of f at any point $\mathbf{x} \in \mathbb{R}^d$. That is, we will use the following assumption:

- (B1) We receive a vector $g(\mathbf{x}, \xi)$ such that $\mathbb{E}_\xi[g(\mathbf{x}, \xi)] = \nabla f(\mathbf{x})$ for any $\mathbf{x} \in \mathbb{R}^d$.

We will also make the following assumption on the variance of the noise.

- (B2) (*Sub-Gaussian Noise*) The stochastic gradient satisfies $\mathbb{E}_\xi [\exp(\|\nabla f(\mathbf{x}) - g(\mathbf{x}, \xi)\|^2/\sigma^2)] \leq \exp(1)$, $\forall \mathbf{x}$.

The condition (B2) has been used by [Nemirovski et al. \(2009\)](#) and [Harvey et al. \(2019a\)](#) to prove high probability convergence guarantees. Intuitively, it implies that the tails of the noise distribution are dominated by tails of a Gaussian distribution. Note that, by Jensen’s inequality, this condition implies a bounded variance on the stochastic gradients.

4. A General Analysis for Algorithms with Momentum

In this section, we will consider a generic stochastic optimization algorithm with Polyak’s momentum (Polyak, 1964; Qian, 1999; Sutskever et al., 2013), also known as the Heavy-ball algorithm or classic momentum, see Algorithm 1.

Algorithm 1 Algorithms with Momentum

- 1: **Input:** $\mathbf{m}_0 = 0, \{\eta_t\}_{t=1}^T, 0 < \mu \leq 1, \mathbf{x}_1 \in \mathbb{R}^d$
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: Get stochastic gradient $\mathbf{g}_t = g(\mathbf{x}_t, \xi_t)$
 - 4: $\mathbf{m}_t = \mu \mathbf{m}_{t-1} + \eta_t \mathbf{g}_t$
 - 5: $\mathbf{x}_{t+1} = \mathbf{x}_t - \mathbf{m}_t$
 - 6: **end for**
-

Two forms of momentum, but not equivalent. First, we want to point out that there two forms of Heavyball algorithms are possible. The first one is in Algorithm 1, while the second one is

$$\begin{aligned} \mathbf{m}_t &= \mu \mathbf{m}_{t-1} + \mathbf{g}_t, \\ \mathbf{x}_{t+1} &= \mathbf{x}_t - \eta_t \mathbf{m}_t. \end{aligned} \quad (1)$$

This second is used in many practical implementation, see, for example, PyTorch (Paszke et al., 2019). It would seem that there is no reason to prefer one over the other. However, here we argue that the classic form of momentum is the right one if we want to use adaptive learning rates. To see why, let’s unroll the updates in both cases. Using the update in Algorithm 1, we have

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \mathbf{g}_t - \mu \eta_{t-1} \mathbf{g}_{t-1} - \mu^2 \eta_{t-1} \mathbf{g}_{t-2} \dots,$$

while using the update in (1), we have

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \mathbf{g}_t - \mu \eta_t \mathbf{g}_{t-1} - \mu^2 \eta_t \mathbf{g}_{t-2} \dots$$

In words, in the first case the update is composed by a sum of weighted gradients, each one multiplied by a learning rate we decided in the past. On the other hand, in the update (1) the update is composed by a sum of weighted gradients, each one multiplied by the *current* learning rate. From the analysis point of view, the second update destroys the independence between the past and the future, introducing a dependency that breaks our analysis, unless we introduce very strict conditions on the gradients. On the other hand, the update in Algorithm 1 allows us to carry out the analysis because each learning rate was chosen only with the knowledge of the past. Note that this is known problem in adaptive algorithms: the lack of independence between past and present is exactly the reason why Adam fails to converge on simple 1d convex problems, see for example the discussion in Savarese et al. (2019).

It is interesting to note that usually people argue that these two types updates for momentum are usually considered equivalent. This seems indeed true only if the learning rates are not adaptive.

Assumptions on learning rates. Note that in the pseudocode we do not specify the learning rates $\eta_t \in \mathbb{R}^d$. In fact, our analysis covers the case of generic learning rates and adaptive ones too. We only need the following assumptions on the stepsizes η_t :

- (C1) η_t is non-increasing, i.e., $\eta_{t+1} \leq \eta_t, \forall t$.
- (C2) η_t is independent with ξ_t .

The first assumption is very common (e.g., Duchi et al., 2011; Reddi et al., 2018; Li & Orabona, 2019; Chen et al., 2019; Zhou et al., 2018). Indeed, AdaGrad has the non-increasing step sizes by the definition. Also, Reddi et al. (2018) have claimed that the main issue of the divergences of Adam and RMSProp lies in the positive definiteness of $1/\eta_t - 1/\eta_{t-1}$.

The need of the second assumption is technical and shared by similar analysis (Li & Orabona, 2019; Savarese et al., 2019). Indeed, Li & Orabona (2019) showed that delayed step sizes can avoid the possible deviation brought by the step sizes that include the current noise.

High probability guarantee. Adaptive learning rates and in general learning rates that are decided using previous gradients become stochastic variables. This makes the high probability analysis more complex. Hence, we use a new concentration inequality for martingales in which the variance is treated as a random variable, rather than a deterministic quantity. We use this concentration in the proof of Lemma 2. Our proof, in the Appendix, merges ideas from the related results in Beygelzimer et al. (2011, Theorem 1) and Lan et al. (2012, Lemma 2). A similar result has also been shown by Jin et al. (2019, Lemma 6).

Lemma 1. *Assume that Z_1, Z_2, \dots, Z_T is a martingale difference sequence with respect to $\xi_1, \xi_2, \dots, \xi_T$ and $\mathbb{E}_t [\exp(Z_t^2/\sigma_t^2)] \leq \exp(1)$ for all $1 \leq t \leq T$, where σ_t is a sequence of random variables with respect to $\xi_1, \xi_2, \dots, \xi_{t-1}$. Then, for any fixed $\lambda > 0$ and $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have*

$$\sum_{t=1}^T Z_t \leq \frac{3}{4} \lambda \sum_{t=1}^T \sigma_t^2 + \frac{1}{\lambda} \ln \frac{1}{\delta}.$$

Main result. We can now present our general lemma, that allows to analyze SGD with momentum with adaptive learning rates. We will then instantiate it for particular examples.

Lemma 2. Assume (A, B1, B2, C1, C2). Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the iterates of Algorithm 1 satisfy

$$\begin{aligned} \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t) \rangle &\leq \frac{3\|\boldsymbol{\eta}_1\| \sigma^2 (1-\mu^T)^2}{(1-\mu)^2} \ln \frac{1}{\delta} \\ &+ 2(f(\mathbf{x}_1) - f^*) + \frac{M(3-\mu)}{1-\mu} \sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2. \end{aligned}$$

Lemma 2 accomplishes the task of upper bounding the inner product $\sum_{t=1}^T \langle \boldsymbol{\eta}_t \nabla f(\mathbf{x}_t), \mathbf{m}_t \rangle$. Then, it is easy to lower bound the l.h.s by $\sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t) \rangle$ using the assumption (C1), followed by the upper bound of $\sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2$ based on the setting of $\boldsymbol{\eta}_t$.

4.1. SGD with Momentum with $\frac{1}{\sqrt{t}}$ Learning Rates

As a warm-up, we now use Lemma 2 to prove a high probability convergence guarantee for the simple case of deterministic learning rates of $\eta_{t,i} = \frac{c}{\sqrt{t}}$.

Theorem 1. Let T the number of iterations of Algorithm 1. Assume (A, B1, B2). Set step size $\boldsymbol{\eta}_t$ as $\eta_{t,i} = \frac{c}{\sqrt{t}}, i = 1, \dots, d$, where $c \leq \frac{1-\mu^T}{4M(3-2\mu)}$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the iterates of Algorithm 1 satisfy

$$\begin{aligned} \min_{1 \leq t \leq T} \|\nabla f(\mathbf{x}_t)\|^2 &\leq \frac{4(f(\mathbf{x}_1) - f^*)}{c\sqrt{T}} + \frac{6(1-\mu^T)^2 \sigma^2}{(1-\mu)^2 \sqrt{T}} \\ &+ \frac{4(3-\mu)cM\sigma^2 \ln \frac{2Te}{\delta} \ln T}{(1-\mu)\sqrt{T}}. \end{aligned}$$

The proof is in Appendix.

4.2. AdaGrad with Momentum

Now, we are going to prove the convergence rate of a variant AdaGrad in which we use momentum and learning rates that do not contain the current gradient. That is, the step sizes are defined as $\boldsymbol{\eta}_t = (\eta_{t,j})_{j=1, \dots, d}$

$$\eta_{t,j} = \frac{\alpha}{\sqrt{\beta + \sum_{i=1}^{t-1} g_{i,j}^2}}, \quad j = 1, \dots, d, \quad (2)$$

where $\alpha, \beta > 0$. Removing the current gradient from the learning rate was proposed in Li & Orabona (2019); Savarese et al. (2019). Following the naming style in (Savarese et al., 2019), we denote this variant by Delayed AdaGrad.

Obviously, (2) satisfies (C1) and (C2). Hence, we are able to employ Lemma 2 to analyze this variant. Moreover, for Delayed AdaGrad, we upper bound $\sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2$ with the following lemma, whose proof is in the Appendix.

Lemma 3. Assume (A, B1, B2). Let $\boldsymbol{\eta}_t$ set as in (2), where $\alpha, \beta > 0$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have

$$\begin{aligned} \sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2 &\leq \frac{4d\alpha^2 \sigma^2}{\beta} \ln \frac{2Te}{\delta} + \frac{4\alpha}{\sqrt{\beta}} \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t) \rangle \\ &\leq 2\alpha^2 d \ln \left(\sqrt{\beta + \frac{2T\sigma^2 \ln \frac{2Te}{\delta}}{d}} + \sqrt{\frac{2}{d} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \right). \end{aligned}$$

We now present the convergence guarantee for Delayed AdaGrad with momentum.

Theorem 2 (Delayed AdaGrad with Momentum). Assume (A, B1, B2). Let $\boldsymbol{\eta}_t$ set as in (2), where $\alpha, \beta > 0$ and $4\alpha \leq \frac{\sqrt{\beta}(1-\mu)^2}{2M(1+\mu)}$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the iterates of Algorithm 1 satisfy

$$\begin{aligned} \min_{1 \leq t \leq T} \|\nabla f(\mathbf{x}_t)\|^2 &\leq \frac{1}{T} \max \left(\frac{4C(T)^2}{\alpha^2}, \frac{C(T)}{\alpha} \sqrt{2\beta + 4T\sigma^2 \ln \frac{3Te}{\delta}} \right), \end{aligned}$$

$$\text{where } C(T) = O \left(\frac{1}{\alpha} + \frac{d \left(\alpha + \sigma^2 \left(\alpha \ln \frac{T}{\delta} + \frac{\ln \frac{1}{\delta}}{1-\mu} \right) \right)}{1-\mu} \right).$$

Adaptivity to Noise. Observe that when $\sigma = 0$, the convergence rate recovers the rate of Gradient Descent if $O(\frac{1}{T})$ with a constant learning rate. On the other hand, in the noisy case, it matches the rate of SGD $O(\frac{\sigma}{\sqrt{T}})$ with the optimal worst-case learning rate of $O(\frac{1}{\sigma\sqrt{t}})$. In other words, with a unique learning rate, we recover two different optimal convergence rates that requires two different learning rates and the knowledge of σ . This adaptivity of Delayed AdaGrad was already proved in Li & Orabona (2019), but only in expectation and without a momentum term.

Dependency on μ . Observe that the convergence upper bound increases over $\mu \in (0, 1)$ and the optimal upper bound is achieved when taking the momentum parameter $\mu = 0$. In words, the algorithms without momentums have the best theoretical results. This is a known caveat for this kind of analysis and a similar behavior w.r.t. μ is present, e.g., in Zou et al. (2018, Theorem 1) for algorithms with Polyak's momentum.

5. Conclusion and Future Work.

In this work, we present a high probability analysis of adaptive SGD with Polyak's momentum in the nonconvex setting. Without using the common assumption of bounded gradients nor bounded noise, we give the high probability bound

for SGD with Polyak’s momentum with step size $O(\frac{1}{\sqrt{t}})$ and for Delayed AdaGrad with momentum. In particular, to the best of our knowledge, this is the *first* high probability convergence guarantee for adaptive methods.

In the future, we plan to extend our results to more adaptive methods, such as Adam and AMSGrad (Tieleman & Hinton, 2012). Moreover, we will explore other forms of momentum such as exponential moving average and Nesterov’s momentum (Nesterov, 1983).

Acknowledgements

This material is based upon work supported by the National Science Foundation under grant no. 1925930 “Collaborative Research: TRIPODS Institute for Optimization and Learning”.

References

- Beygelzimer, A., Langford, J., Li, L., Reyzin, L., and Schapire, R. Contextual bandit algorithms with supervised learning guarantees. In *Proc. of the International Conference on Artificial Intelligence and Statistics*, pp. 19–26, 2011.
- Chen, X., Liu, S., Sun, R., and Hong, M. On the convergence of a class of Adam-type algorithms for non-convex optimization. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- Duchi, J. C., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- Gadat, S., Panloup, F., and Saadane, S. Stochastic heavy ball. *Electronic Journal of Statistics*, 12(1):461–529, 2018.
- Harvey, N. J., Liaw, C., Plan, Y., and Randhawa, S. Tight analyses for non-smooth stochastic gradient descent. In *Conference on Learning Theory*, pp. 1579–1613, 2019a.
- Harvey, N. J., Liaw, C., and Randhawa, S. Simple and optimal high-probability bounds for strongly-convex stochastic gradient descent. *arXiv preprint arXiv:1909.00843*, 2019b.
- Jain, P., Nagaraj, D., and Netrapalli, P. Making the last iterate of sgd information theoretically optimal. In Beygelzimer, A. and Hsu, D. (eds.), *Proc. of the Conference on Learning Theory (COLT)*, volume 99 of *Proceedings of Machine Learning Research*, pp. 1752–1755, Phoenix, USA, 25–28 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v99/jain19a.html>.
- Jin, C., Netrapalli, P., Ge, R., Kakade, S. M., and Jordan, M. I. A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*, 2019.
- Kakade, S. M. and Tewari, A. On the generalization ability of online strongly convex programming algorithms. In *Advances in Neural Information Processing Systems*, pp. 801–808, 2009.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Lan, G., Nemirovski, A., and Shapiro, A. Validation analysis of mirror descent stochastic approximation method. *Mathematical programming*, 134(2):425–458, 2012.
- Li, X. and Orabona, F. On the convergence of stochastic gradient descent with adaptive stepsizes. In *Proc. of the 22nd International Conference on Artificial Intelligence and Statistics, AISTATS*, 2019. URL <https://arxiv.org/abs/1805.08114>.
- Loizou, N. and Richtárik, P. Momentum and stochastic momentum for stochastic gradient, Newton, proximal point and subspace descent methods. *arXiv preprint arXiv:1712.09677*, 2017.
- Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.
- Nesterov, Y. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In *Doklady AN SSSR (translated as Soviet. Math. Doct.)*, volume 269, pp. 543–547, 1983.
- Nesterov, Y. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer, 2004.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.
- Polyak, B. T. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- Qian, N. On the momentum term in gradient descent learning algorithms. *Neural networks*, 12(1):145–151, 1999.

- Reddi, S. J., Kale, S., and Kumar, S. On the convergence of Adam and beyond. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/pdf?id=ryQu7f-RZ>.
- Robbins, H. and Monro, S. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951. URL <https://projecteuclid.org/euclid.aoms/1177729586>.
- Savarese, P., McAllester, D., Babu, S., and Maire, M. Domain-independent dominance of adaptive methods. *arXiv preprint arXiv:1912.01823*, 2019.
- Sutskever, I., Martens, J., Dahl, G., and Hinton, G. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pp. 1139–1147, 2013.
- Tieleman, T. and Hinton, G. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 2012.
- Ward, R., Wu, X., and Bottou, L. AdaGrad stepizes: Sharp convergence over nonconvex landscapes, from any initialization. In *International Conference on Machine Learning*, pp. 6677–6686, 2019.
- Yang, T., Lin, Q., and Li, Z. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. *arXiv preprint arXiv:1604.03257*, 2016.
- Zhou, D., Tang, Y., Yang, Z., Cao, Y., and Gu, Q. On the convergence of adaptive gradient methods for nonconvex optimization, 2018. arXiv:1808.05671.
- Zou, F., Shen, L., Jie, Z., Sun, J., and Liu, W. Weighted AdaGrad with unified momentum. *arXiv preprint arXiv:1808.03408*, 2018.
- Zou, F., Shen, L., Jie, Z., Zhang, W., and Liu, W. A sufficient condition for convergences of Adam and RMSProp. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 11127–11135, 2019.

A. Appendix

A.1. Details of Section 4

Proof of Lemma 1. Set $\tilde{Z}_t = Z_t/\sigma_t$. By the assumptions of Z_t and σ_t , we have

$$\mathbb{E}_t[\tilde{Z}_t] = \frac{1}{\sigma_t}\mathbb{E}_t[Z_t] = 0 \quad \text{and} \quad \mathbb{E}_t\left[\exp\left(\tilde{Z}_t^2\right)\right] \leq \exp(1).$$

By Jensen's inequality, it follows that for any $c \in [0, 1]$,

$$\mathbb{E}_t\left[\exp\left(c\tilde{Z}_t^2\right)\right] = \mathbb{E}_t\left[\left(\exp\left(\tilde{Z}_t^2\right)\right)^c\right] \leq \left(\mathbb{E}_t\left[\exp\left(\tilde{Z}_t^2\right)\right]\right)^c \leq \exp(c). \quad (3)$$

Also it can be verified that $\exp(x) \leq x + \exp(9x^2/16)$ for all x , hence for $|\kappa| \in [0, 4/3]$ we get

$$\mathbb{E}_t\left[\exp\left(\kappa\tilde{Z}_t\right)\right] \leq \mathbb{E}_t\left[\exp\left(9\kappa^2\tilde{Z}_t^2/16\right)\right] \leq \exp\left(9\kappa^2/16\right) \leq \exp\left(3\kappa^2/4\right), \quad (4)$$

where in the second inequality, we used (3). Besides, $kx \leq 3k^2/8 + 2x^2/3$ holds for any k and x . Hence for $|\kappa| \geq 4/3$, we get

$$\mathbb{E}_t\left[\exp\left(\kappa\tilde{Z}_t\right)\right] \leq \exp\left(3\kappa^2/8\right)\mathbb{E}_t\left[\exp\left(2\tilde{Z}_t^2/3\right)\right] \leq \exp\left(3\kappa^2/8 + 2/3\right) \leq \exp\left(3\kappa^2/4\right), \quad (5)$$

where in the second inequality we used (3). Combining (4) and (5), we get $\forall \kappa$,

$$\mathbb{E}_t\left[\exp\left(\kappa\tilde{Z}_t\right)\right] \leq \exp\left(3\kappa^2/4\right). \quad (6)$$

Note that the above analysis for (6) still hold when κ is a random variable with respect to $\xi_1, \xi_2, \dots, \xi_{t-1}$. So for Z_t , we have $\mathbb{E}_t[\exp(\lambda Z_t)] \leq \exp(3\lambda^2\sigma_t^2/4)$, $\lambda > 0$.

Define the random variables $Y_0 = 1$ and $Y_t = Y_{t-1} \exp(\lambda Z_t - 3\lambda^2\sigma_t^2/4)$, $1 \leq t \leq T$. So, we have $E_t Y_t = Y_{t-1} \exp(-3\lambda^2\sigma_t^2/4) \cdot \mathbb{E}_t[\exp(\lambda Z_t)] \leq Y_{t-1}$. Now, taking full expectation over all variables $\xi_1, \xi_2, \dots, \xi_T$, we have

$$\mathbb{E}Y_T \leq \mathbb{E}Y_{T-1} \leq \dots \leq \mathbb{E}Y_0 = 1.$$

By Markov's inequality, $P(Y_T \geq \frac{1}{\delta}) \leq \delta$, and $Y_T = \exp\left(\lambda \sum_{t=1}^T Z_t - \frac{3}{4}\lambda^2 \sum_{t=1}^T \sigma_t^2\right)$, we have

$$P\left(Y_T \geq \frac{1}{\delta}\right) = P\left(\lambda \sum_{t=1}^T Z_t - \frac{3}{4}\lambda^2 \sum_{t=1}^T \sigma_t^2 \geq \ln \frac{1}{\delta}\right) = P\left(\sum_{t=1}^T Z_t \geq \frac{3}{4}\lambda \sum_{t=1}^T \sigma_t^2 + \frac{1}{\lambda} \ln \frac{1}{\delta}\right) \leq \delta,$$

which completes the proof. \square

To prove Lemma 2, we first need the following technical Lemma.

Lemma 4. $\forall T \geq 1$, it holds

$$\sum_{t=1}^T a_t \sum_{i=1}^t b_i = \sum_{t=1}^T b_t \sum_{i=t}^T a_i \quad \text{and} \quad \sum_{t=1}^T a_t \sum_{i=0}^{t-1} b_i = \sum_{t=0}^{T-1} b_t \sum_{i=t+1}^T a_i.$$

Proof. We prove these equalities by induction. When $T = 1$, they obviously hold. Now, for $k < T$, assume that $\sum_{t=1}^k a_t \sum_{i=1}^t b_i = \sum_{t=1}^k b_t \sum_{i=t}^k a_i$. Then, we have

$$\begin{aligned} \sum_{t=1}^{k+1} a_t \sum_{i=1}^t b_i &= \sum_{t=1}^k a_t \sum_{i=1}^t b_i + a_{k+1} \sum_{i=1}^{k+1} b_i = \sum_{t=1}^k b_t \sum_{i=t}^k a_i + a_{k+1} \sum_{i=1}^k b_i + a_{k+1} b_{k+1} \\ &= \sum_{t=1}^k b_t \sum_{i=t}^{k+1} a_i + a_{k+1} b_{k+1} = \sum_{t=1}^{k+1} b_t \sum_{i=t}^{k+1} a_i. \end{aligned}$$

Hence, by induction, the equality is proved.

Similarly, for second equality assume that for $k < T$ we have $\sum_{t=1}^k a_t \sum_{i=0}^{t-1} b_i = \sum_{t=0}^{k-1} b_t \sum_{i=t+1}^k a_i$. Then, we have

$$\begin{aligned} \sum_{t=1}^{k+1} a_t \sum_{i=0}^{t-1} b_i &= \sum_{t=1}^k a_t \sum_{i=0}^{t-1} b_i + a_{k+1} \sum_{i=0}^k b_i = \sum_{t=0}^{k-1} b_t \sum_{i=t+1}^k a_i + a_{k+1} \sum_{i=0}^{k-1} b_i + a_{k+1} b_k \\ &= \sum_{t=0}^{k-1} b_t \sum_{i=t+1}^{k+1} a_i + a_{k+1} b_k = \sum_{t=0}^k b_t \sum_{i=t+1}^{k+1} a_i. \end{aligned}$$

By induction, we finish the proof. \square

Proof of Lemma 2. By the smoothness of f and the definition of \mathbf{x}_{t+1} , we have

$$f(\mathbf{x}_{t+1}) - f(\mathbf{x}_t) \leq -\langle \nabla f(\mathbf{x}_t), \mathbf{m}_t \rangle + \frac{M}{2} \|\mathbf{m}_t\|^2. \quad (7)$$

We now upper bound $-\langle \nabla f(\mathbf{x}_t), \mathbf{m}_t \rangle$.

$$\begin{aligned} &-\langle \nabla f(\mathbf{x}_t), \mathbf{m}_t \rangle \\ &= -\mu \langle \nabla f(\mathbf{x}_t), \mathbf{m}_{t-1} \rangle - \langle \nabla f(\mathbf{x}_t), \eta_t \mathbf{g}_t \rangle \\ &= -\mu \langle \nabla f(\mathbf{x}_{t-1}), \mathbf{m}_{t-1} \rangle - \mu \langle \nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-1}), \mathbf{m}_{t-1} \rangle - \langle \nabla f(\mathbf{x}_t), \eta_t \mathbf{g}_t \rangle \\ &\leq -\mu \langle \nabla f(\mathbf{x}_{t-1}), \mathbf{m}_{t-1} \rangle + \mu \|\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-1})\| \|\mathbf{m}_{t-1}\| - \langle \nabla f(\mathbf{x}_t), \eta_t \mathbf{g}_t \rangle \\ &\leq -\mu \langle \nabla f(\mathbf{x}_{t-1}), \mathbf{m}_{t-1} \rangle + \mu M \|\mathbf{m}_{t-1}\|^2 - \langle \nabla f(\mathbf{x}_t), \eta_t \mathbf{g}_t \rangle, \end{aligned}$$

where the second inequality is due to the smoothness of f . Hence, iterating the inequality we have

$$\begin{aligned} -\langle \nabla f(\mathbf{x}_t), \eta_t \mathbf{m}_t \rangle &\leq -\mu^2 \langle \nabla f(\mathbf{x}_{t-2}), \mathbf{m}_{t-2} \rangle + \mu^2 M \|\mathbf{m}_{t-2}\|^2 + \mu M \|\mathbf{m}_{t-1}\|^2 \\ &\quad - \mu \langle \nabla f(\mathbf{x}_{t-1}), \eta_{t-1} \mathbf{g}_{t-1} \rangle - \langle \nabla f(\mathbf{x}_t), \eta_t \mathbf{g}_t \rangle \\ &\leq M \sum_{i=1}^{t-1} \mu^{t-i} \|\mathbf{m}_i\|^2 - \sum_{i=1}^t \mu^{t-i} \langle \nabla f(\mathbf{x}_i), \eta_i \mathbf{g}_i \rangle. \end{aligned} \quad \square$$

Thus, denoting by $\epsilon_t = \mathbf{g}_t - \nabla f(\mathbf{x}_t)$ and summing (7) over t from 1 to T , we obtain

$$\begin{aligned} f^* - f(\mathbf{x}_1) &\leq f(\mathbf{x}_{T+1}) - f(\mathbf{x}_1) \\ &\leq M \sum_{t=1}^T \sum_{i=1}^{t-1} \mu^{t-i} \|\mathbf{m}_i\|^2 - \sum_{t=1}^T \sum_{i=1}^t \mu^{t-i} \langle \nabla f(\mathbf{x}_i), \eta_i \mathbf{g}_i \rangle + \sum_{t=1}^T \frac{M}{2} \|\mathbf{m}_t\|^2 \\ &\leq M \sum_{t=1}^T \sum_{i=1}^{t-1} \mu^{t-i} \|\mathbf{m}_i\|^2 - \sum_{t=1}^T \langle \eta_t, \nabla f(\mathbf{x}_t)^2 \rangle - \sum_{t=1}^T \sum_{i=1}^t \mu^{t-i} \langle \nabla f(\mathbf{x}_i), \eta_i \epsilon_i \rangle \\ &\quad + \sum_{t=1}^T \frac{M}{2} \|\mathbf{m}_t\|^2. \end{aligned}$$

By Lemma 4, we have

$$M \sum_{t=1}^T \sum_{i=1}^{t-1} \mu^{t-i} \|\mathbf{m}_i\|^2 \leq \frac{M}{1-\mu} \sum_{t=1}^T \|\mathbf{m}_t\|^2.$$

Also, by Lemma 4, we have

$$\begin{aligned} -\sum_{t=1}^T \sum_{i=1}^t \mu^{t-i} \langle \nabla f(\mathbf{x}_i), \eta_i \epsilon_i \rangle &= -\sum_{t=1}^T \mu^{-t} \langle \nabla f(\mathbf{x}_t), \eta_t \epsilon_t \rangle \sum_{i=t}^T \mu^i \\ &= -\frac{1}{1-\mu} \sum_{t=1}^T \langle \nabla f(\mathbf{x}_t), \eta_t \epsilon_t \rangle (1 - \mu^{T-t+1}) \triangleq S_T. \end{aligned}$$

We then upper bound S_T . Denote by $L_t := -\frac{1-\mu^{T-t+1}}{1-\mu} \langle \nabla f(\mathbf{x}_t), \boldsymbol{\eta}_t \boldsymbol{\epsilon}_t \rangle$, and $N_t := \frac{(1-\mu^{T-t+1})^2}{(1-\mu)^2} \|\boldsymbol{\eta}_t \nabla f(\mathbf{x}_t)\|^2 \sigma^2$. Using the assumptions on the noise, for any $1 \leq t \leq T$, we have

$$\exp\left(\frac{L_t^2}{N_t}\right) \leq \exp\left(\frac{\|\boldsymbol{\eta}_t \nabla f(\mathbf{x}_t)\|^2 \|\boldsymbol{\epsilon}_t\|^2 (1-\mu^{T-t+1})^2 / (1-\mu)^2}{N_t}\right) = \exp\left(\frac{\|\boldsymbol{\epsilon}_t\|^2}{\sigma^2}\right) \leq \exp(1).$$

We can also see that for any t , $\mathbb{E}_t[L_t] = -\frac{1-\mu^{T-t+1}}{1-\mu} \sum_{i=1}^d \boldsymbol{\eta}_{t,i} \nabla f(\mathbf{x}_t)_i \mathbb{E}_t[\boldsymbol{\epsilon}_{t,i}] = 0$. Thus, from Lemma 1, with probability at least $1 - \delta$, any $\lambda > 0$, we have

$$\begin{aligned} S_T = \sum_{t=1}^T L_t &\leq \frac{3}{4} \lambda \sum_{t=1}^T N_t + \frac{1}{\lambda} \ln \frac{1}{\delta} \leq \frac{3\lambda(1-\mu^T)^2}{4(1-\mu)^2} \sum_{t=1}^T \|\boldsymbol{\eta}_t \nabla f(\mathbf{x}_t)\|^2 \sigma^2 + \frac{1}{\lambda} \ln \frac{1}{\delta} \\ &\leq \frac{3\lambda \|\boldsymbol{\eta}_1\| (1-\mu^T)^2}{4(1-\mu)^2} \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle \sigma^2 + \frac{1}{\lambda} \ln \frac{1}{\delta}. \end{aligned}$$

Finally, we upper bound $\sum_{t=1}^T \|\mathbf{m}_t\|^2$. From the convexity of $\|\cdot\|^2$, we have

$$\|\mathbf{m}_t\|^2 = \left\| \mu \mathbf{m}_{t-1} + (1-\mu) \frac{\boldsymbol{\eta}_t \mathbf{g}_t}{1-\mu} \right\|^2 \leq \mu \|\mathbf{m}_{t-1}\|^2 + \frac{1}{1-\mu} \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2.$$

Summing over t from 1 to T , we have

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{m}_t\|^2 &\leq \sum_{t=1}^T \mu \|\mathbf{m}_{t-1}\|^2 + \frac{1}{1-\mu} \sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2 \\ &= \sum_{t=1}^{T-1} \mu \|\mathbf{m}_t\|^2 + \frac{1}{1-\mu} \sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2 \\ &\leq \sum_{t=1}^T \mu \|\mathbf{m}_t\|^2 + \frac{1}{1-\mu} \sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2, \end{aligned}$$

where in the first equality we used $\mathbf{m}_0 = 0$. Reordering the terms, we have that

$$\sum_{t=1}^T \|\mathbf{m}_t\|^2 \leq \frac{1}{(1-\mu)^2} \sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2.$$

Combining things together, and taking $\lambda = \frac{2(1-\mu)^2}{3\|\boldsymbol{\eta}_1\|(1-\mu^T)^2\sigma^2}$, with probability at least $1 - \delta$, we have

$$\begin{aligned} f^* - f(\mathbf{x}_1) &\leq \frac{1}{\lambda} \ln \frac{1}{\delta} + \sum_{t=1}^T \left[\left(\frac{M}{2} + \frac{M}{1-\mu} \right) \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2 - \left(1 - \frac{3\lambda \|\boldsymbol{\eta}_1\| (1-\mu^T)^2 \sigma^2}{4(1-\mu)^2} \right) \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle \right] \\ &= \frac{3\|\boldsymbol{\eta}_1\| (1-\mu^T)^2 \sigma^2}{2(1-\mu)^2} \ln \frac{1}{\delta} + \sum_{t=1}^T \left[\frac{(3-\mu)M}{2(1-\mu)} \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2 - \frac{1}{2} \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle \right]. \end{aligned}$$

Rearranging the terms, we get the stated bound.

A.2. Details of Section 4.1

The proof of this Theorem 1 makes use of the following additional Lemma on the tail of sub-gaussian noise.

Lemma 5. *Assume B2, then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have*

$$\max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2 \leq \sigma^2 \ln \frac{Te}{\delta}.$$

Proof. By Markov's inequality, for any $A > 0$,

$$\begin{aligned}
 \mathbb{P}\left(\max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2 > A\right) &= \mathbb{P}\left(\exp\left(\frac{\max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2}{\sigma^2}\right) > \exp\left(\frac{A}{\sigma^2}\right)\right) \\
 &\leq \exp\left(-\frac{A}{\sigma^2}\right) \mathbb{E}\left[\exp\left(\frac{\max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2}{\sigma^2}\right)\right] \\
 &= \exp\left(-\frac{A}{\sigma^2}\right) \mathbb{E}\left[\max_{1 \leq t \leq T} \exp\left(\frac{\|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2}{\sigma^2}\right)\right] \\
 &\leq \exp\left(-\frac{A}{\sigma^2}\right) \sum_{t=1}^T \mathbb{E}\left[\exp\left(\frac{\|\nabla f(\mathbf{x}_t) - \mathbf{g}_t\|^2}{\sigma^2}\right)\right] \leq \exp\left(-\frac{A}{\sigma^2} + 1\right) T. \quad \square
 \end{aligned}$$

Proof of Theorem 1. With the fact that $\|\mathbf{a} + \mathbf{b}\|^2 \leq 2\|\mathbf{a}\|^2 + 2\|\mathbf{b}\|^2$, we have

$$\begin{aligned}
 \sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2 &= \sum_{t=1}^T \eta_t^2 \|\mathbf{g}_t\|^2 \leq \sum_{t=1}^T 2\eta_t^2 \|\nabla f(\mathbf{x}_t)\|^2 + \sum_{t=1}^T 2\eta_t^2 \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2 \\
 &\leq \sum_{t=1}^T 2\eta_t^2 \|\nabla f(\mathbf{x}_t)\|^2 + \max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2 \sum_{t=1}^T 2\eta_t^2.
 \end{aligned}$$

By Lemma 5, Lemma 2 and the union bound, we have that with probability at least $1 - \delta$,

$$\begin{aligned}
 \frac{\eta_T}{2} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2 &\leq \left(1 - \frac{2M(3-\mu)}{1-\mu} \eta_1\right) \sum_{t=1}^T \eta_t \|\nabla f(\mathbf{x}_t)\|^2 \\
 &\leq 2(f(\mathbf{x}_1) - f^*) + \frac{2(3-\mu)c^2 M \sigma^2 \ln \frac{2Te}{\delta} \ln T}{1-\mu} + \frac{3c(1-\mu^T)^2 \sigma^2}{(1-\mu)^2} \ln \frac{1}{\delta}.
 \end{aligned}$$

Rearranging the terms and lower bounding $\sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2$ by $T \cdot \min_{1 \leq t \leq T} \|\nabla f(\mathbf{x}_t)\|^2$, we have the stated bound. \square

A.3. Details of Section 4.2

For the proof of Lemma 3, we first need the following technical Lemma.

Lemma 6. Let $a_i \geq 0, \dots, T$ and $f : [0, +\infty) \rightarrow [0, +\infty)$ non-increasing function. Then

$$\sum_{t=1}^T a_t f\left(a_0 + \sum_{i=1}^t a_i\right) \leq \int_{a_0}^{\sum_{t=0}^T a_t} f(x) dx.$$

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

$$a_i f(s_i) = \int_{s_{i-1}}^{s_i} f(s_i) dx \leq \int_{s_{i-1}}^{s_i} f(x) dx.$$

Summing over $i = 1, \dots, T$, we have the stated bound. \square

Proof of Lemma 3. First, we separate $\sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2$ into two terms:

$$\sum_{t=1}^T \|\boldsymbol{\eta}_t \mathbf{g}_t\|^2 = \sum_{t=1}^T \|\boldsymbol{\eta}_{t+1} \mathbf{g}_t\|^2 + \sum_{t=1}^T \langle \boldsymbol{\eta}_t^2 - \boldsymbol{\eta}_{t+1}^2, \mathbf{g}_t^2 \rangle.$$

Then, we proceed

$$\begin{aligned}
 \sum_{t=1}^T \langle \boldsymbol{\eta}_t^2 - \boldsymbol{\eta}_{t+1}^2, \mathbf{g}_t^2 \rangle &= \sum_{i=1}^d \sum_{t=1}^T (\eta_{t,i}^2 - \eta_{t+1,i}^2) g_{t,i}^2 \\
 &\leq \sum_{i=1}^d \sum_{t=1}^T 2\eta_{t,i} g_{t,i}^2 (\eta_{t,i} - \eta_{t+1,i}) \\
 &\leq 2 \sum_{i=1}^d \max_{1 \leq t \leq T} \eta_{t,i} g_{t,i}^2 \sum_{t=1}^T (\eta_{t,i} - \eta_{t+1,i}) \\
 &\leq 2 \sum_{i=1}^d \eta_{1,i} \max_{1 \leq t \leq T} \eta_{t,i} g_{t,i}^2 \\
 &\leq 4 \sum_{i=1}^d \eta_{1,i} \max_{1 \leq t \leq T} \eta_{t,i} (g_{t,i}^2 - \nabla f(\mathbf{x}_t)_i^2) + 4 \sum_{i=1}^d \eta_{1,i} \sum_{t=1}^T \eta_{t,i} \nabla f(\mathbf{x}_t)_i^2 \\
 &\leq 4 \sum_{i=1}^d \eta_{1,i}^2 \max_{1 \leq t \leq T} |g_{t,i}^2 - \nabla f(\mathbf{x}_t)_i^2| + 4 \sum_{i=1}^d \eta_{1,i} \sum_{t=1}^T \eta_{t,i} \nabla f(\mathbf{x}_t)_i^2 \\
 &\leq \frac{4d\alpha^2}{\beta} \max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2 + \frac{4\alpha}{\sqrt{\beta}} \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle. \tag{8}
 \end{aligned}$$

Using Lemma 5 on (8), for $\delta \in (0, 1)$, with probability at least $1 - \frac{\delta}{2}$, we have

$$\sum_{t=1}^T \langle \boldsymbol{\eta}_t^2 - \boldsymbol{\eta}_{t+1}^2, \mathbf{g}_t^2 \rangle \leq \frac{4d\alpha^2\sigma^2}{\beta} \ln \frac{2Te}{\delta} + \frac{4\alpha}{\sqrt{\beta}} \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle.$$

We now upper bound $\sum_{t=1}^T \|\boldsymbol{\eta}_{t+1} \mathbf{g}_t\|^2$:

$$\begin{aligned}
 \sum_{t=1}^T \|\boldsymbol{\eta}_{t+1} \mathbf{g}_t\|^2 &= \sum_{i=1}^d \sum_{t=1}^T \frac{\alpha^2 g_{t,i}^2}{\beta + \sum_{j=1}^t g_{j,i}^2} \\
 &\leq \sum_{i=1}^d \alpha^2 \ln \left(\beta + \sum_{t=1}^T g_{t,i}^2 \right) \\
 &\leq \alpha^2 d \ln \left(\beta + \frac{1}{d} \sum_{i=1}^d \sum_{t=1}^T g_{t,i}^2 \right) \\
 &= 2\alpha^2 d \ln \left(\sqrt{\beta + \frac{1}{d} \sum_{t=1}^T \|\mathbf{g}_t\|^2} \right) \\
 &\leq 2\alpha^2 d \ln \left(\sqrt{\beta + \frac{2T}{d} \max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2} + \sqrt{\frac{2}{d} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \right), \tag{9}
 \end{aligned}$$

where in the first inequality we used Lemma 6 and in the second inequality we used Jensen's inequality. Then using Lemma 5 on (9), with probability at least $1 - \frac{\delta}{2}$, we have

$$\sum_{t=1}^T \|\boldsymbol{\eta}_{t+1} \mathbf{g}_t\|^2 \leq 2\alpha^2 d \ln \left(\sqrt{\beta + \frac{2T\sigma^2}{d} \ln \frac{2Te}{\delta}} + \sqrt{\frac{2}{d} \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \right).$$

Putting things together, we have the stated bound. \square

Finally, to prove Theorem 2, we need the two following Lemmas.

Lemma 7 (Lemma 6 in Li & Orabona (2019)). *Let $x \geq 0$, $A, C, D \geq 0$, $B > 0$, and $x^2 \leq (A + Bx)(C + D \ln(A + Bx))$. Then.*

$$x < 32B^3D^2 + 2BC + 8B^2D\sqrt{C} + A/B.$$

Lemma 8 (Lemma 5 in Li & Orabona (2019)). *If $x \geq 0$, and $x \leq C\sqrt{A + Bx}$, then $x \leq \max(2BC^2, C\sqrt{2A})$.*

Proof of Theorem 2. By Lemma 2 and Lemma 3, for $\delta \in (0, 1)$, with probability at least $1 - \frac{2}{3}\delta$, we have

$$\begin{aligned} & \left(1 - \frac{4\alpha M(3-\mu)}{\sqrt{\beta}(1-\mu)}\right) \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle \\ & \leq 2(f(\mathbf{x}_1) - f^*) + \frac{M(3-\mu)}{1-\mu} \left(K + \frac{4d\alpha^2\sigma^2}{\beta} \ln \frac{3Te}{\delta} \right) + \frac{3\|\boldsymbol{\eta}_1\|\sigma^2(1-\mu^T)^2}{(1-\mu)^2} \ln \frac{3}{\delta}. \end{aligned}$$

where K denotes $2\alpha^2d \ln \left(\sqrt{\beta + \frac{2T\sigma^2}{d} \ln \frac{2Te}{\delta}} + \sqrt{\frac{2}{d}} \sqrt{\sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \right)$ for conciseness.

Rearranging the terms, we have

$$\begin{aligned} & \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle \\ & \leq \frac{1}{1 - \frac{4\alpha M(3-\mu)}{\sqrt{\beta}(1-\mu)}} \left[2(f(\mathbf{x}_1) - f^*) + \frac{M(3-\mu)}{1-\mu} \left(K + \frac{4d\alpha^2\sigma^2}{\beta} \ln \frac{3Te}{\delta} \right) + \frac{3\|\boldsymbol{\eta}_1\|\sigma^2(1-\mu^T)^2}{(1-\mu)^2} \ln \frac{3}{\delta} \right] \\ & \leq 4(f(\mathbf{x}_1) - f^*) + \frac{2M(3-\mu)}{1-\mu} \left(K + \frac{4d\alpha^2\sigma^2}{\beta} \ln \frac{3Te}{\delta} \right) + \frac{3\|\boldsymbol{\eta}_1\|\sigma^2(1-\mu^T)^2}{(1-\mu)^2} \ln \frac{3}{\delta} \\ & \triangleq C(T), \end{aligned}$$

where in the second inequality we used $4\alpha \leq \frac{\sqrt{\beta}(1-\mu)}{2M(3-\mu)}$. Also, we have

$$\begin{aligned} \sum_{t=1}^T \langle \boldsymbol{\eta}_t, \nabla f(\mathbf{x}_t)^2 \rangle & \geq \sum_{t=1}^T \langle \boldsymbol{\eta}_T, \nabla f(\mathbf{x}_t)^2 \rangle \\ & = \sum_{i=1}^d \frac{\alpha \sum_{t=1}^T \nabla f(\mathbf{x}_t)_i^2}{\sqrt{\beta + \sum_{t=1}^T g_{t,i}^2}} \\ & \geq \sum_{i=1}^d \frac{\alpha \sum_{t=1}^T \nabla f(\mathbf{x}_t)_i^2}{\sqrt{\beta + 2 \sum_{t=1}^T \nabla f(\mathbf{x}_t)_i^2 + 2 \sum_{t=1}^T (g_{t,i} - \nabla f(\mathbf{x}_t)_i)^2}} \\ & \geq \sum_{i=1}^d \frac{\alpha \sum_{t=1}^T \nabla f(\mathbf{x}_t)_i^2}{\sqrt{\beta + 2 \sum_{i=1}^d \sum_{t=1}^T \nabla f(\mathbf{x}_t)_i^2 + 2 \sum_{i=1}^d \sum_{t=1}^T (g_{t,i} - \nabla f(\mathbf{x}_t)_i)^2}} \\ & \geq \frac{\alpha \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2}{\sqrt{\beta + 2 \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2 + 2T \max_{1 \leq t \leq T} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2}}. \end{aligned}$$

By Lemma 5, with probability at least $1 - \delta$, we have

$$\sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{C(T)}{\alpha} \times \sqrt{\beta + 2 \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2 + 2T\sigma^2 \ln \frac{3Te}{\delta}}. \quad (10)$$

$$\begin{aligned}
 \text{RHS of (10)} &\leq \frac{C(T)}{\alpha} \times \left(\sqrt{\beta + 2T\sigma^2 \ln \frac{3Te}{\delta}} + \sqrt{2 \sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \right) \\
 &\leq \left[C + D \ln \left(A + B \sqrt{\sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \right) \right] \times \left(A + B \sqrt{\sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \right), \quad (11)
 \end{aligned}$$

where $A = \sqrt{\beta + 2T\sigma^2 \ln \frac{3Te}{\delta}}$, $B = \sqrt{2}$, $C = \frac{4(f(\mathbf{x}_1) - f^*)}{\alpha} + \frac{8M(3-\mu)d\alpha\sigma^2}{\beta(1-\mu)} \ln \frac{3Te}{\delta} + \frac{3d(1-\mu^T)^2\sigma^2}{\beta(1-\mu)^2} \ln \frac{3}{\delta}$ and $D = \frac{4\alpha dM(3-\mu)}{1-\mu}$. Using Lemma 7, we have that

$$\sqrt{\sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2} \leq 32B^3D^2 + 2BC + 8B^2D\sqrt{C} + \frac{A}{B}.$$

We use this upper bound in the logarithmic term of (11). Thus, we have (10) again, this time with

$$\begin{aligned}
 C(T) &= C + D \ln(2A + 32B^4D^2 + 2B^2C + 8B^3D\sqrt{C}) \\
 &= O\left(\frac{1}{\alpha} + \frac{d\left(\alpha + \sigma^2\left(\alpha \ln \frac{T}{\delta} + \frac{\ln \frac{1}{\delta}}{1-\mu}\right)\right)}{1-\mu}\right).
 \end{aligned}$$

Solving (11) by Lemma 8 and lower bounding $\sum_{t=1}^T \|\nabla f(\mathbf{x}_t)\|^2$ by $T \min_{1 \leq t \leq T} \|\nabla f(\mathbf{x}_t)\|^2$, we get the stated bound. \square