

# A TALE OF TWO GEOMETRIES: ADAPTIVE OPTIMIZERS AND NON-EUCLIDEAN DESCENT

Shuo Xie<sup>1\*</sup> Tianhao Wang<sup>2\*</sup> Beining Wu<sup>3</sup> Zhiyuan Li<sup>1</sup>

<sup>1</sup>Toyota Technological Institute at Chicago <sup>2</sup>University of California, San Diego

<sup>3</sup>University of Chicago

{shuox, zhiyuanli}@ttic.edu, tianhaowang@ucsd.edu, beiningw@uchicago.edu

## ABSTRACT

Adaptive optimizers can reduce to normalized steepest descent (NSD) when only adapting to the current gradient, suggesting a close connection between the two algorithmic families. A key distinction between their analyses, however, lies in the geometries, e.g., smoothness notions, they rely on. In the convex setting, adaptive optimizers are governed by a stronger adaptive smoothness condition, while NSD relies on the standard notion of smoothness. We extend the theory of adaptive smoothness to the nonconvex setting and show that it precisely characterizes the convergence of adaptive optimizers. Moreover, we establish that adaptive smoothness enables acceleration of adaptive optimizers with Nesterov momentum in the convex setting, a guarantee unattainable under standard smoothness for certain non-Euclidean geometry. We further develop an analogous comparison for stochastic optimization by introducing adaptive gradient variance, which parallels adaptive smoothness and leads to dimension-free convergence guarantees that cannot be achieved under standard gradient variance for certain non-Euclidean geometry.

## 1 INTRODUCTION

Adaptive optimizers such as Adam have been indispensable for training large-scale machine learning models (Bi et al., 2024; Dubey et al., 2024; Yang et al., 2025; Wen et al., 2025). Their dominance in training efficiency, however, has recently been challenged by the surprising effectiveness of simpler Normalized Steepest Descent (NSD)-type methods such as Muon and Lion (Jordan et al., 2024; Chen et al., 2023; Team et al., 2025; Liu et al., 2025; Shah et al., 2025). Behind this competition of two family of optimizers, a broader consensus has begun to emerge: their superior performance is critically related to their ability to exploit non-Euclidean geometry of the loss landscape (Balles et al., 2020; Xie & Li, 2024; Zhang et al., 2024; Pethick et al., 2025).

Recent studies have rigorously characterized how adaptive optimizers exploit non-Euclidean geometry. For example, Maladkar et al. (2024) and Xie et al. (2025a) show that AdaGrad and Adam benefit from exploiting the  $\ell_\infty$ -geometry of loss functions, and a one-sided variant of Shampoo has been shown to leverage the geometry induced by the matrix spectral norm (Xie et al., 2025b; An et al., 2025). Notably, Bernstein & Newhouse (2024) proposed a striking connection between adaptive optimizers and NSD: with exponential moving average (EMA) turned off, certain adaptive optimizers reduce exactly to their NSD counterparts. For example, without EMA, Adam coincides with NSD under the  $\ell_\infty$  norm, and Shampoo coincides with NSD under the matrix spectral norm, which is proposed to be an independent algorithm Muon. Yet, beyond these connections, there is no formal result that systematically characterizes the relationship between the two families of algorithms. This naturally motivates the following question:

**Q1.** *Do adaptive methods (like Adam, Shampoo) and their corresponding non-Euclidean descent (like Lion, Muon) exploit the non-Euclidean geometry of loss landscape in the same way?*

To address this question, we adopt a theoretical perspective and focus on comparing different types of smoothness assumptions that underpin the analysis of these methods. In fact, even under the

---

\*Equal contribution.

same geometry, two distinct notions of smoothness arise. The first is the standard smoothness under a general norm (cf. Definition 2.3), which governs the convergence rate of NSD. The second is called the *adaptive smoothness* (cf. Definition 2.4), introduced by Xie et al. (2025b) and shown to govern the convergence rate of adaptive optimizers in the convex case. Indeed, a main contribution of our work is to show that adaptive smoothness also characterizes the convergence rate of adaptive optimizers in the nonconvex setting. Therefore, while both adaptive optimizers and NSD can exploit non-Euclidean geometry, they rely on fundamentally different smoothness assumptions.

This difference is not merely terminological but quantitative: adaptive smoothness is always no smaller than the standard smoothness under the same geometry. In other words, from the standpoint of technical conditions, the adaptive smoothness represents a stronger assumption, which in turn motivates our second question:

**Q2.** *Does the stronger smoothness assumption in adaptive methods offer optimization benefit?*

We answer this question affirmatively. In particular, we show that by leveraging Nesterov acceleration, adaptive optimizers can attain an accelerated  $O(T^{-2})$  rate under adaptive smoothness in the convex setting. In sharp contrast, it has been shown by Guzmán & Nemirovski (2015) that the convergence rate of any optimizer is no better than  $\Omega(T^{-1})$  under the standard  $\ell_\infty$  smoothness assumption. This establishes a clear separation: adaptive smoothness enables adaptive optimizers to achieve acceleration under non-Euclidean geometry, while the standard smoothness fails. Therefore, the stronger adaptive smoothness assumption indeed translates into concrete optimization benefits, showing its difference from the standard smoothness.

In fact this difference has a direct and interesting analogy in terms of the noise assumption in the stochastic setting. When gradient noise is present, its variability can be measured in two distinct ways: the standard variance considers gradient variation under a fixed norm, whereas the *adaptive variance* (cf. Definition 4.1) measures noise in a more stringent but also more adaptive way that requires uniform control over the geometry prescribed by each preconditioner under consideration. By construction, adaptive variance is always no smaller than standard variance, directly paralleling the relationship between adaptive and standard smoothness. Analogous to adaptive smoothness that enables acceleration under a stronger requirement, adaptive variance can likewise yield benefits despite being larger. We demonstrate this through a careful analysis of NSD under two types of noise assumptions: adaptive variance enables a dimension-free rate, which is not attainable in the worst case under the standard variance condition.

Taken together, our results demonstrate that adaptive smoothness and adaptive variance are different from their standard counterparts as adaptive smoothness enables an acceleration rate and adaptive noise enables a dimension-free rate. These findings reveal an intricate interplay between adaptivity and non-Euclidean geometry, deepening our theoretical understanding of adaptivity in optimization.

Below we summarize our main contributions.

- In Section 3, we show the convergence rate for adaptive optimizers on nonconvex functions (Theorems D.2, D.7 and D.8), which depends on the adaptive smoothness and matches optimal  $\tilde{O}(T^{-1/4})$  rate. It theoretically justifies that adaptive methods and NSD exploit the geometry through different smoothness notions in the nonconvex setting.
- In Section 4.2, we identify the benefit of the adaptive smoothness by showing it enables an acceleration rate  $\tilde{O}(T^{-2})$  of adaptive optimizers equipped with Nesterov momentum (Theorem 4.3) in contrast to the convergence rate  $\Omega(T^{-1})$  the standard  $\ell_\infty$  smoothness.
- In Section 4, we extend the benefit of adaptive geometry to noise assumptions by introducing adaptive noise (Definition 4.1). We show that this stronger notion of noise can provide a new type of convergence rate for NSD with momentum on nonconvex functions which gets rid of dependence on parameter size  $d$  (Theorem 4.5). We complement its superiority by providing a lower bound under the standard noise (Theorem 4.7).
- Our analysis of adaptive optimizers is carried out through a unified framework that covers a broad class of methods, including AdaGrad, AdaGrad-Norm, and one-sided Shampoo. The proof technique developed in this framework may be of independent interest.

## 1.1 NOTATIONS

Let  $\mathcal{M}^d$  be the set of all  $d$ -by- $d$  matrices,  $\mathcal{S}^d \subset \mathcal{M}^d$  be the subset of all symmetric matrices. We use  $\mathcal{S}_+^d$  to denote the set of positive semi-definite matrices. We denote by  $\mathbf{I}_d \in \mathcal{M}^d$  the identity matrix. For matrices  $\mathbf{A}, \mathbf{B}$ , we denote their inner product by  $\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr}(\mathbf{A}^\top \mathbf{B})$ .

For  $\mathbf{H} \in \mathcal{S}_+^d$ ,  $\|\mathbf{x}\|_{\mathbf{H}} := \sqrt{\mathbf{x}^\top \mathbf{H} \mathbf{x}}$  is the (semi-)norm of  $\mathbf{x} \in \mathbb{R}^d$  with respect to  $\mathbf{H}$ . For a convex set  $\mathcal{H} \subseteq \mathcal{S}_+^d$ , we define the induced  $\mathcal{H}$ -norm as

$$\|\mathbf{x}\|_{\mathcal{H}} := \sup_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \|\mathbf{x}\|_{\mathbf{H}}. \quad (1)$$

Throughout the paper, we reserve  $f$  for the loss function and  $\mathbf{x}_0$  for the initialization of an optimization algorithm. For convenience, we denote the initial suboptimality as  $\Delta_0 = f(\mathbf{x}_0) - \min_{\mathbf{x}} f(\mathbf{x})$ .

## 2 FROM ADAM/SIGNGD TO ADAPTIVE SMOOTHNESS

We use the example of Adam and SignGD to motivate the notion of adaptive smoothness in Section 2.1, and then present the formal definition in Section 2.2, along with some related background.

### 2.1 ADAM AND SIGNGD CAN EXPLOIT $\ell_\infty$ GEOMETRY, BUT IN DIFFERENT WAYS

We start by discussing a specific pair of algorithms, Adam and SignGD, to illustrate the problem of interest. It is known that SignGD can be viewed as Normalized Steepest Descent (NSD) under the  $\ell_\infty$  norm and its convergence rate for deterministic nonconvex functions admits the following form (Xie et al., 2025a)

$$\min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_1 \leq O\left(\sqrt{\frac{\Delta_0 L_{\|\cdot\|_\infty}(f)}{T}}\right)$$

where  $L_{\|\cdot\|_\infty}(f)$  is the standard smoothness of  $f$  under the  $\ell_\infty$  norm (see Definition 2.3). Note that SignGD can also be viewed as a special case of Adam with  $\beta_1 = \beta_2 = 0$ . However, the convergence rate of Adam for general  $\beta_1, \beta_2$  instead depends on a different diagonal adaptive smoothness notion, which is defined as  $L_{\text{diag}}(f) = \min_{\mathbf{H} \in \mathcal{D}^d, -\mathbf{H} \preceq \nabla^2 f(\mathbf{x}) \preceq \mathbf{H}} \text{Tr}(\mathbf{H})$  in Maladkar et al. (2024); Xie et al. (2025a). In particular, Adam with  $\beta_1 = 0$  (a.k.a. RMSProp) for deterministic nonconvex functions admits the convergence rate  $\min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_1 = \tilde{O}(\sqrt{\Delta_0 L_{\text{diag}}(f)/T})$  (Xie et al., 2024). Notably, this diagonal adaptive smoothness is always no smaller than  $L_{\|\cdot\|_\infty}(f)$  (Balles et al., 2020). This suggests that though both SignGD and Adam admit convergence guarantees for the  $\ell_1$  norm (the dual norm of  $\|\cdot\|_\infty$ ) of the gradients, they achieve so under different smoothness notions. This distinction motivates the following question:

*How does the diagonal adaptive smoothness  $L_{\text{diag}}(f)$  emerge as an  $\ell_\infty$  geometry?*

To address this question, let us consider the convergence rate of NSD under any norm  $\|\cdot\|_{\mathbf{H}}$  for  $\mathbf{H} \in \mathcal{H} = \mathcal{D}_+^d$  (see Theorem 4.5):

$$\min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_{\mathbf{H},*} = O\left(\sqrt{\frac{\Delta_0 L_{\|\cdot\|_{\mathbf{H}}}(f)}{T}}\right) \quad (2)$$

where  $\|\cdot\|_{\mathbf{H},*}$  is the dual norm of  $\|\cdot\|_{\mathbf{H}}$ . Minimizing both sides of (2) over  $\mathbf{H} \in \mathcal{D}_+^d$  with  $\text{Tr}(\mathbf{H}) \leq 1$  yields

$$\inf_{\substack{\text{diagonal} \\ \text{Tr}(\mathbf{H}) \leq 1}} \min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_{\mathbf{H},*} = O\left(\sqrt{\frac{\Delta_0}{T} \inf_{\substack{\text{diagonal} \\ \text{Tr}(\mathbf{H}) \leq 1}} L_{\|\cdot\|_{\mathbf{H}}}(f)}\right) = O\left(\sqrt{\frac{\Delta_0 L_{\text{diag}}(f)}{T}}\right) \quad (3)$$

where the equality can be checked by the definition of  $L_{\text{diag}}(f)$ . Now the right-hand side matches the aforementioned convergence rate of Adam. The adaptivity of Adam is then demonstrated by its ability to automatically identify and adapt to the best diagonal matrix-induced norm for any given loss function, without the need of knowing  $\mathbf{H}$ .

Importantly, the left-hand side of (3) is closely related to the  $\ell_1$  norm of the gradients because

$$\sup_{\text{diagonal } \mathbf{H} \geq 0, \text{Tr}(\mathbf{H}) \leq 1} \|\cdot\|_{\mathbf{H}} = \|\cdot\|_{\infty}, \quad \inf_{\text{diagonal } \mathbf{H} \geq 0, \text{Tr}(\mathbf{H}) \leq 1} \|\cdot\|_{\mathbf{H},*} = \|\cdot\|_1. \quad (4)$$

We illustrate this fact in Fig. 1. In words, this means that *the  $\ell_{\infty}$  norm is the pointwise supremum of all weighted  $\ell_2$  norms induced by diagonal matrices with unit trace, whereas its dual, the  $\ell_1$  norm, is the pointwise infimum of all the corresponding dual norms*. Also, the unit  $\ell_{\infty}$  ball is the intersection of all unit balls for those  $\ell_2$  norms, and the unit  $\ell_1$  ball is the union of all dual unit balls.

Indeed, the duality between supremum of a class of primal norms and infimum of the corresponding dual norms in (4) is not just a coincidence, but rather a special property induced by the structure of the preconditioner set  $\mathcal{H} = \mathcal{D}_+^d$  for Adam. This property holds more generally for any well-structured preconditioner set and we discuss the corresponding adaptive smoothness in Section 2.2.

## 2.2 ADAPTIVE SMOOTHNESS ASSOCIATED WITH WELL-STRUCTURED PRECONDITIONER

The following definition of well-structured preconditioner sets is proposed by Xie et al. (2025b) to unify the analysis of a broad family of adaptive optimizers with structured preconditioners.

**Definition 2.1** (Well-structured preconditioner set).  $\mathcal{H} \subseteq \mathcal{S}_+^d$  is said to be a well-structured preconditioner set if  $\mathcal{H} = \mathcal{S}_+^d \cap \mathcal{K}$  for some matrix subalgebra<sup>1</sup>  $\mathcal{K} \subseteq \mathcal{M}^d$  with  $\mathbf{I}_d \in \mathcal{K}$ .

As will be discussed in Section 3.1, many commonly used adaptive optimizers, including Adam, AdaGrad, and their variants, can be cast into the framework of a meta-algorithm (Algorithm 1) with well-structured preconditioner sets. A specific case is  $\mathcal{H} = \mathcal{D}_+^d$ , the set of all diagonal PSD matrices, which is the running example in the previous subsection. For any such well-structured preconditioner set  $\mathcal{H}$ , we have the duality between the supremum of the primal norms and the infimum of the corresponding dual norms, formalized in the following lemma.

**Lemma 2.2.** Let  $\mathcal{H} \subseteq \mathcal{S}_+^d$  be any well-structured preconditioner set. Recall that its induced norm is defined as  $\|\cdot\|_{\mathcal{H}} = \sup_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \|\cdot\|_{\mathbf{H}}$ . Then it holds that

$$\|\cdot\|_{\mathcal{H},*} = \inf_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \|\cdot\|_{\mathbf{H},*} = \inf_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \|\cdot\|_{\mathbf{H}^{-1}}.$$

Based on this fact, we can generalize the discussion in Section 2.1 to any well-structured preconditioner set  $\mathcal{H}$ , showing that NSD and adaptive optimizers with preconditioner set  $\mathcal{H}$  can exploit the geometry induced by  $\|\cdot\| = \|\cdot\|_{\mathcal{H}}$  via two different smoothness notions, the former being the standard smoothness under  $\|\cdot\|_{\mathcal{H}}$  and the latter being the adaptive smoothness defined in Definition 2.4.

We proceed to introduce the adaptive smoothness associated with any well-structured preconditioner set  $\mathcal{H}$ . We first review the standard smoothness notion under a general norm  $\|\cdot\|$ .

**Definition 2.3.** For a loss function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  and any norm  $\|\cdot\|$ , we will use  $L_{\|\cdot\|}(f)$  to denote the smoothness of  $f$  with respect to  $\|\cdot\|$ , i.e., the smallest positive constant  $L$  such that  $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_* \leq L \|\mathbf{x} - \mathbf{y}\|$  for any  $\mathbf{x}, \mathbf{y}$ .

When  $\|\cdot\| = \|\cdot\|_{\mathcal{H}}$  for some well-structured preconditioner set  $\mathcal{H}$ ,  $L_{\|\cdot\|_{\mathcal{H}}}(f)$  is then the standard smoothness of  $f$  under the norm induced by  $\mathcal{H}$ . In contrast, the adaptive smoothness associated with  $\mathcal{H}$  is defined as the smallest smoothness of  $f$  under all norms  $\|\cdot\|_{\mathbf{H}}$  induced by  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , as formalized below. This term is introduced as  $\mathcal{H}$ -smoothness in Xie et al. (2025b). We rename it to highlight this notion adapts to the structure of  $\mathcal{H}$ , in contrast to the standard smoothness.

**Definition 2.4** (Adaptive Smoothness, Xie et al. 2025b). The adaptive smoothness of a function  $f$  w.r.t. a well-structured preconditioner set  $\mathcal{H}$  is defined as the smallest smoothness of  $f$  under all  $\|\cdot\|_{\mathbf{H}}$  for  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , that is,

$$\Lambda_{\mathcal{H}}(f) := \min_{\substack{\mathbf{H} \in \mathcal{H} \\ \text{Tr}(\mathbf{H}) \leq 1}} L_{\|\cdot\|_{\mathbf{H}}}(f) = \min_{\substack{\mathbf{H} \in \mathcal{H} \\ \forall \mathbf{x}, -\mathbf{H} \leq \nabla^2 f(\mathbf{x}) \preceq \mathbf{H}}} \text{Tr}(\mathbf{H}). \quad (5)$$

<sup>1</sup>For a set of  $d$ -by- $d$  matrices  $\mathcal{K} \subseteq \mathcal{M}^d$ , we say that  $\mathcal{K}$  is a *subalgebra* if it is closed under scalar multiplication, matrix addition, and matrix multiplication. More concretely, we require that for any  $\alpha \in \mathbb{R}$  and  $\mathbf{A}, \mathbf{B} \in \mathcal{K}$ , it holds that  $\alpha \mathbf{A}, \mathbf{A}\mathbf{B}, \mathbf{A} + \mathbf{B} \in \mathcal{K}$ .

**Algorithm 1** General Adaptive Optimization Algorithm**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , convex cone  $\mathcal{H} \subset \mathcal{S}_+$ , decay factor  $\beta$ **Input:** initialization  $\mathbf{x}_0$ , stochastic loss functions  $\{f_t\}_{t=1}^T : \mathbb{R}^d \rightarrow \mathbb{R}$ 


---

```

 $\mathbf{M}_{-1} \leftarrow \mathbf{0}$ 
for  $t = 0, 1, \dots, T - 1$  :
   $\mathbf{g}_t \leftarrow \nabla f_t(\mathbf{x}_t)$ 
   $\mathbf{M}_t \leftarrow \begin{cases} \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top, & \text{Cumulative variant,} \\ \beta \mathbf{M}_{t-1} + (1 - \beta) \mathbf{g}_t \mathbf{g}_t^\top, & \text{EMA variant,} \\ \beta \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top, & \text{Weighted variant.} \end{cases}$ 
   $\mathbf{V}_t \leftarrow \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$ 
   $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t - \eta \mathbf{V}_t^{-1} \mathbf{g}_t$ 
return  $\mathbf{x}_T$ 

```

---

In the deterministic convex setting, it has been shown by Xie et al. (2025b) that the convergence rate of an adaptive optimizer with any well-structured preconditioner set  $\mathcal{H}$  is of order  $O(\Lambda_{\mathcal{H}}(f) \|\mathcal{X}\|_{\mathcal{H}}^2 / T)$ . In Section 3, we extend such characterization to the nonconvex setting, demonstrating that the adaptive smoothness  $\Lambda_{\mathcal{H}}(f)$  governs the convergence behavior of any adaptive optimizer with well-structured preconditioner set  $\mathcal{H}$ .

**Comparison between two smoothness notions.** For any  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) = 1$ , it always holds  $\|\mathbf{x} - \mathbf{y}\|_{\mathcal{H}} \geq \|\mathbf{x} - \mathbf{y}\|_{\mathbf{H}}$  and  $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathcal{H},*} \leq \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathbf{H},*}$ . Therefore,

$$L_{\|\cdot\|_{\mathbf{H}}}(f) = \sup_{\mathbf{x}, \mathbf{y}} \frac{\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathbf{H},*}}{\|\mathbf{x} - \mathbf{y}\|_{\mathbf{H}}} \geq \sup_{\mathbf{x}, \mathbf{y}} \frac{\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathcal{H},*}}{\|\mathbf{x} - \mathbf{y}\|_{\mathcal{H}}} = L_{\|\cdot\|_{\mathcal{H}}}(f).$$

Minimizing over  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) = 1$  then yields  $\Lambda_{\mathcal{H}}(f) = L_{\|\cdot\|_{\mathbf{H}}}(f) \geq L_{\|\cdot\|_{\mathcal{H}}}(f)$ . In other words, as a condition, the adaptive smoothness is arguably stronger than the standard smoothness. But they can differ by at most a multiplicative factor of  $d$ , as summarized in Proposition 2.5.

**Proposition 2.5.** For any well-structured preconditioner set  $\mathcal{H} \subseteq \mathcal{S}_+^d$  and any loss function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ , it always holds that  $L_{\|\cdot\|_{\mathcal{H}}}(f) \leq \Lambda_{\mathcal{H}}(f) \leq d \cdot L_{\|\cdot\|_{\mathcal{H}}}(f)$ .

### 3 UNIFIED ANALYSIS IN THE NONCONVEX SETTING

In the nonconvex setting, we establish a *unified* analysis that encompasses a broad family of adaptive optimization algorithms. Our result highlights how the convergence behavior of these methods depends critically on the notion of adaptive smoothness.

#### 3.1 ADAPTIVE OPTIMIZERS WITH WELL-STRUCTURED PRECONDITIONER SETS

We adopt the framework in Gupta et al. (2017) and Xie et al. (2025b) to describe adaptive optimizers in a unified way, as displayed in Algorithm 1. This meta-algorithm is flexible in two aspects: the way of aggregating past gradients and the choice of preconditioner set  $\mathcal{H}$ . First, there are three different ways to aggregate the past gradients in Algorithm 1, each of which is presented in a separate algorithm block in Appendix D.1. The cumulative and EMA variants are the most common ways, and they are indeed equivalent to the weighted variant up to hyperparameter transformations. Therefore, it suffices to study the weighted variant, and the results for the other two variants follow as corollaries. Another flexibility of Algorithm 1 comes from the choice of convex cone  $\mathcal{H}$ . More specifically, Algorithm 1 recovers several standard optimizers by specifying  $\mathcal{H}$  as follows:

- $\mathcal{H} = \{\text{all diagonal PSD matrices}\}$  recovers AdaGrad and Adam.
- $\mathcal{H} = \{c \mathbf{I}_d \mid c > 0\}$  recovers AdaGrad-Norm and AdaSGD (Wang & Wiens, 2020).
- $\mathcal{H} = \mathcal{S}_+^d$  recovers full-matrix AdaGrad (Duchi et al., 2011).
- $\mathcal{H} = \mathcal{S}_+^{d_r} \otimes \mathbf{I}_{d_R}$  yields one-sided Shampoo/ASGO recently proposed by (Xie et al., 2025a; An et al., 2025)

In particular, based on the notion of well-structured preconditioner sets defined in Definition 2.1, Xie et al. (2025b) develops a unified convergence analysis for Algorithm 4 in the convex setting, and the convergence rate depends on the adaptive smoothness with respect to  $\mathcal{H}$  defined in Definition 2.4.

**Additional notations.** We define  $P_{\mathcal{H}}(\mathbf{M}) := \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$  for any  $\mathbf{M} \in \mathcal{S}_{++}^d$ . Then in Algorithm 1,  $\mathbf{V}_t = P_{\mathcal{H}}(\mathbf{M}_t)$  and Lemma B.4 will show that  $P_{\mathcal{H}}(\mathbf{M})^2$  is the projection of  $\mathbf{M}$  onto  $\mathcal{H}$ . Specifically, when  $\mathcal{H}$  contains all the PSD matrices,  $\mathbf{V}_t$  is  $\mathbf{M}_t^{\frac{1}{2}}$ .

### 3.2 CONVERGENCE RATE IN THE DETERMINISTIC NONCONVEX SETTING

Here we only present results for the deterministic case to highlight the role of adaptive smoothness, and the complete results for the (stochastic) nonconvex setting and corresponding proofs can be found in Appendix D.2. We first present the convergence guarantee for the weighted variant of Algorithm 1 in the following theorem.

**Theorem 3.1.** For any  $\epsilon \geq 0$ ,  $\beta \in (0, 1]$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 1 with well-structured preconditioner set  $\mathcal{H}$ , where the update of  $\mathbf{M}_t$  follows the weighted version, i.e.,  $\mathbf{M}_t = \beta \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$  for all  $t \in [T]$ . Let  $\Lambda_{\mathcal{H}}(f)$  be the adaptive smoothness of the loss  $f$  according to Definition 2.4. Then when  $f_t \equiv f$ , it holds that

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \leq \frac{\sqrt{\sum_{i=0}^{T-1} \beta^{i/2}}}{T} \xi + \frac{\sqrt{d} \epsilon^{1/4}}{\sqrt{T}} \sqrt{\xi}.$$

where  $\xi = \frac{2\Delta_0}{\eta} + \eta \Lambda_{\mathcal{H}}(f) \|\mathbf{S}_T\|_{\text{op}}$  and  $\mathbf{S}_T = \mathbb{E} \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1}$ .

For general well-structured preconditioner set,  $\|\mathbf{S}_T\|_{\text{op}} = \tilde{O}(\log(d)[(1-\beta)T/\beta + \log(d)])$ . When the preconditioner set only has diagonal matrices,  $\|\mathbf{S}_T\|_{\text{op}} = (1-\beta)T + \tilde{O}(1)$ .

The above result for the weight variant can be converted to guarantees for the cumulative and EMA variants. Specifically, the cumulative variant is equivalent to weighted accumulation with  $\beta = 1$  while the EMA variant with learning rate  $\eta^E$  and stability constant  $\epsilon^E$  produces identical iterates as weighted accumulation with  $\eta^W = \eta^E / \sqrt{1-\beta}$  and  $\epsilon^W = \epsilon^E / (1-\beta)$ . Below we present the result for the cumulative variant in Theorem 3.2, and the result for the EMA variant is in Theorem D.8.

**Theorem 3.2.** For any  $\epsilon \geq 0$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 1 with well-structured preconditioner set  $\mathcal{H}$ , where the update of  $\mathbf{M}_t$  follows the cumulative version, i.e.,  $\mathbf{M}_t = \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$  for all  $t \in [T]$ . Let  $\Lambda_{\mathcal{H}}(f)$  be the adaptive smoothness of the loss  $f$  according to Definition 2.4. Then when  $f_t \equiv f$ , it holds that

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \leq \frac{1}{\sqrt{T}} \left( \xi + \sqrt{d} \epsilon^{1/4} \sqrt{\xi} \right)$$

where  $\xi = \tilde{O}\left(\frac{\Delta_0}{\eta} + \eta \cdot \Lambda_{\mathcal{H}}(f) \log^2 d\right)$ . Moreover, when setting  $\eta = \sqrt{\frac{\Delta_0}{\Lambda_{\mathcal{H}}(f) \log^2 d}}$ , it holds that  $\xi = \tilde{O}(\sqrt{\Delta_0 \cdot \Lambda_{\mathcal{H}}(f)} \log d)$ .

At a high level, Theorem 3.2 shows that, with appropriate hyperparameters, Algorithm 1 with any well-structured preconditioner set  $\mathcal{H}$  achieves a convergence rate of order  $\tilde{O}(\log d \cdot \sqrt{\Delta_0 \Lambda_{\mathcal{H}}(f) / T})$  on deterministic nonconvex functions where  $\tilde{O}(\cdot)$  hides logarithmic factors in problem parameters other than the dimension  $d$ . This result illustrates that the adaptive smoothness  $\Lambda_{\mathcal{H}}(f)$  governs the convergence rate of adaptive optimizers in the nonconvex setting, complementing previous results for the convex setting in Xie et al. (2025b). Moreover, we remark that when  $\mathcal{H}$  contains only diagonal matrices, the  $\log d$  factor disappears, recovering the bounds in Xie et al. (2025a).

It is worth noticing that the convergence guarantees in the above two theorems are concerned with  $\|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*}$  depending on specific  $\mathcal{H}$  rather than  $\|\nabla f(\mathbf{x}_t)\|_2$ . For the specific case of Adam where  $\mathcal{H}$  is the set of all diagonal PSD, this becomes a guarantees in terms of  $\ell_1$  norm of the gradients, as we discussed in Section 2.1. On the other hand, Pethick et al. (2025); Kovalev (2025a) show that NSD achieves  $O\left(\left(\frac{\Delta_0 L_{\mathcal{H}}(f)}{T}\right)^{\frac{1}{2}}\right)$  in the deterministic case. Taken together, these two rates suggest that adaptive optimizers and NSD exploit different smoothness notions in the nonconvex setting.

### 3.3 TECHNICAL CONTRIBUTION: A NOVEL MATRIX INEQUALITY

Previous theoretical results for one-sided Shampoo/ASGO (Algorithm 7) and other well-structured preconditioners primarily focus on convex objectives (Xie et al., 2025b; Kovalev, 2025a). In the non-convex regime, existing convergence analyses apply essentially when the preconditioner set contains only diagonal matrices<sup>2</sup> (Xie et al., 2025a). In contrast, we show the first unified convergence analysis that applies to any general well-structured preconditioner set, well beyond the diagonal cases.

A central difficulty in our analysis is the extension from diagonal preconditioners to a general preconditioner set  $\mathcal{H}$ . In the diagonal case, the proof basically decomposes to entry-wise analyses, and scalar telescoping readily yields the desired bounds. However, for general  $\mathcal{H}$ , *noncommutativity prevents such simplification*, and bounding the second-order terms requires handling delicate matrix inequalities. Our resolution of this challenge yields a key technical contribution, formalized below.

**Lemma 3.3.** *Let  $\epsilon \geq 0$  and  $\beta \in (0, 1]$ . For any  $T \in \mathbb{N}$ , consider any sequence of vectors  $\mathbf{g}_0, \dots, \mathbf{g}_{T-1} \in \mathbb{R}^d$ . Let  $\mathbf{M}_{-1} = 0$ , and recursively define  $\mathbf{M}_t = \beta \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$  for  $t = 0, \dots, T-1$ . For any well-structured preconditioner set  $\mathcal{H}$ , define  $\mathbf{V}_t = \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$  for each  $t \in [T-1]$ . Then for any  $\mathbf{H} \in \mathcal{H} \cap \mathcal{S}_{++}^d$ , it holds*

$$\sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2 \leq \text{Tr}(\mathbf{H}) \|\mathbf{S}_T\|_{\text{op}} \quad \text{where } \mathbf{S}_T = \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1}.$$

Moreover, there exists an absolute constant  $C_1, C_2 > 0$ , independent of  $d, T, \epsilon, \beta$  and  $\mathcal{H}$ , such that

$$\|\mathbf{S}_T\|_{\text{op}} \leq C_1 \left( 1 + \log \left( 1 + \frac{d}{\epsilon} \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 + d^2(1-\beta)T \right) \right) \left( \frac{1-\beta}{\beta} T + \log \|\mathbf{V}_{T-1}^2 / \epsilon\|_{\text{op}} \right) + C_2.$$

In the special case when  $\mathcal{H}$  is commutative, the above bound can be further improved to

$$\|\mathbf{S}_T\|_{\text{op}} \leq (1-\beta)T + \log \|\mathbf{V}_{T-1}^2 / \epsilon\|_{\text{op}}.$$

Specializing  $\mathbf{g}_t$ 's to be the gradients, Lemma 3.3 provides a general upper bound on the sum of second-order terms. It highlights the essential gap between diagonal and general preconditioner sets: noncommutativity introduces an additional  $\log d$  factor, making the dependence strictly worse than in the diagonal case. Nevertheless, this is the first bound that applies to arbitrary well-structured preconditioner sets, and it plays a central role in extending convex analyses to the nonconvex setting.

The proof of Lemma 3.3 can be found in Appendix C. A key step is to establish a novel matrix inequality Lemma C.1 that relates the difference between two positive definite matrices to the difference between their logarithms, which may be of independent interest.

## 4 BENEFIT OF ADAPTIVE GEOMETRY

We have shown in Section 3 that the nonconvex convergence rate of adaptive optimizers relies critically on the adaptive smoothness of the loss, and the bound is worse than that of the corresponding NSD. This naturally raises the concern in Question 2: does the stronger adaptive smoothness lead to stronger results? In this section, we address this question from two complementary angles:

1. Under the adaptive smoothness assumption, adaptive optimizers can achieve faster convergence rates on convex functions via Nesterov acceleration.
2. The distinction between standard smoothness and adaptive smoothness mirrors a parallel separation in the assumptions on gradient noise.

At a high level, these two angles share the same underlying mechanism: *Under non-Euclidean geometry, averaging might not be effective in reducing the norm*, which we will explain below.

<sup>2</sup>This can be generalized to any commutative well-structured preconditioner set.

**Algorithm 2** Accelerated Adaptive Algorithm

**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , convex cone  $\mathcal{H} \subseteq \mathcal{S}_+$   
**Input:** initial  $\mathbf{x}_0$ , constants  $\alpha_0, \dots, \alpha_T \in (0, 1]$   
 $\bar{\mathbf{x}}_0 \leftarrow \mathbf{x}_0, \mathbf{M}_{-1} \leftarrow \mathbf{0}$   
**for**  $t = 0, 1, \dots, T - 1$  :  
 $\mathbf{g}_t \leftarrow \nabla f_t^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t)$  where  $f_t^{\alpha_t, \bar{\mathbf{x}}_t}$  is in (8)  
 $\mathbf{M}_t \leftarrow \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$   
 $\mathbf{V}_t \leftarrow \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$   
 $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t - \eta \mathbf{V}_t^{-1} \mathbf{g}_t$   
 $\bar{\mathbf{x}}_{t+1} \leftarrow \alpha_t \mathbf{x}_{t+1} + (1 - \alpha_t) \bar{\mathbf{x}}_t$   
**return**  $\bar{\mathbf{x}}_T$

**Algorithm 3** NSD with momentum

**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , norm  $\|\cdot\|$ , averaging parameter  $\alpha$   
**Input:** initialization  $\mathbf{x}_0$ , initialization  $\mathbf{m}_0$ , stochastic loss functions  $\{f_t\}_{t=1}^T$   
**for**  $t = 0, 1, \dots, T - 1$  :  
 $\mathbf{g}_t \leftarrow \nabla f_t(\mathbf{x}_t)$   
 $\mathbf{m}_t \leftarrow (1 - \alpha) \mathbf{m}_{t-1} + \alpha \mathbf{g}_t$   
 $\mathbf{u}_t \leftarrow \arg \max_{\|\mathbf{u}\| \leq 1} \langle \mathbf{m}_t, \mathbf{u} \rangle$   
 $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t - \eta \mathbf{u}_t$   
**return**  $\mathbf{x}_T$

## 4.1 ADAPTIVE VARIANCE: AN ANALOGUE OF ADAPTIVE SMOOTHNESS

Our main results in this section are concerned about accelerated adaptive algorithms for convex functions and NSD in the nonconvex setting. Before presenting these results, we introduce a key technical ingredient: *adaptive variance*, a quantity that serves as the analogue of adaptive smoothness for gradient noise.

**Definition 4.1** (Standard and adaptive gradient variance). *For an index set  $\mathcal{T}$ , let  $\{f_t\}_{t \in \mathcal{T}}$  be a set of stochastic loss functions where each  $f_t : \mathbb{R}^d \rightarrow \mathbb{R}$ .*

- For any norm  $\|\cdot\|$  on  $\mathbb{R}^d$ , the gradient variance of  $\{f_t\}_{t \in \mathcal{T}}$  under  $\|\cdot\|$  is defined as

$$\sigma_{\|\cdot\|}(\{f_t\}_{t \in \mathcal{T}})^2 := \sup_{t \in \mathcal{T}, \mathbf{x} \in \mathbb{R}^d} \mathbb{E}[\|\nabla f_t(\mathbf{x}) - \mathbb{E}[\nabla f_t(\mathbf{x})]\|^2] \quad (6)$$

- The adaptive gradient variance of  $\{f_t\}_{t \in \mathcal{T}}$  w.r.t. any well-structured preconditioner set  $\mathcal{H}$  is

$$\sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}})^2 = \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \sup_{t \in \mathcal{T}, \mathbf{x} \in \mathbb{R}^d} \mathbb{E}[\|\nabla f_t(\mathbf{x}) - \mathbb{E}[\nabla f_t(\mathbf{x})]\|_{\mathbf{H}^{-1}}^2]. \quad (7)$$

This adaptive variance is inspired by the noise assumption in Kovalev (2025a), both capturing the overall variation of gradient noise in the geometry induced by  $\mathcal{H}$ . Compared with the traditional noise assumption that only characterizes  $\ell_2$  norm variance, adaptive variance provides a more informative measure. In addition, analogous to the comparison between  $\Lambda_{\mathcal{H}}(f)$  and  $L_{\|\cdot\|_{\mathcal{H}}}(f)$ , the adaptive variance is always no smaller than  $\|\cdot\|_{\mathcal{H},*}$ -variance, as formalized in Proposition B.11.

Here we can compare Definition 4.1 with bounded covariance assumption in Xie et al. (2025b); An et al. (2025) that there exists  $\Sigma \succeq \mathbf{0}$  such that  $\mathbb{E}[\nabla f(\mathbf{x}_t) - \nabla f_t(\mathbf{x}_t)][\nabla f(\mathbf{x}_t) - \nabla f_t(\mathbf{x}_t)]^\top \preceq \Sigma$ . When the covariance matrix is upper bounded by  $\Sigma$ ,  $\sigma_{\mathcal{H}} \leq \text{Tr}(P_{\mathcal{H}}(\Sigma))$  for general  $\mathcal{H}$  as shown in Proposition B.10. On the other hand, Definition 4.1 doesn't require the existence of  $\Sigma$  that can upper bound the covariance matrix everywhere. Therefore, Definition 4.1 is a weaker assumption than the bounded covariance assumption.

## 4.2 ACCELERATION ON CONVEX FUNCTIONS

We follow the framework in Kovalev (2025a) to formulate a unified class of adaptive optimizers with well-structured preconditioner sets with Nesterov acceleration in Algorithm 2. Through the perspective introduced by Kovalev & Borodich (2024), the idea is to interpret each step of Nesterov acceleration as a single step of standard gradient on a modified loss, i.e.,  $f_t^{\alpha_t, \bar{\mathbf{x}}_t}$  in Eq. 8. Here for a constant  $\alpha \in (0, 1]$  and a reference point  $\bar{\mathbf{x}}$ , the corresponding modified loss is defined as

$$f^{\alpha, \bar{\mathbf{x}}}(\mathbf{x}) := \alpha^{-2} f(\alpha \mathbf{x} + (1 - \alpha) \bar{\mathbf{x}}). \quad (8)$$

For stochastic convex functions satisfying Assumption 4.2, we establish the convergence rate of Algorithm 2 in the following Theorem 4.3, whose proof is in Appendix E.

**Assumption 4.2.** *Let  $f$  be a convex loss function. For all  $t \in [T]$  and any  $\mathbf{x}$ ,  $\mathbb{E}[\nabla f_t(\mathbf{x})] = \nabla f(\mathbf{x})$ .*

**Theorem 4.3.** For a well-structured preconditioner set  $\mathcal{H}$ , let  $f$  be a convex loss function whose  $\mathcal{H}$ -smoothness constant is  $\Lambda_{\mathcal{H}}(f) \in (0, \infty)$  according to Definition 2.4. For  $\epsilon > 0$ ,  $T > 0$ , consider Algorithm 2 with  $\alpha_t = 2/(t+2)$  for  $t = 0, 1, \dots, T-1$ . Suppose  $\mathbf{x}^*$  is the global minima and  $\max_{t=0,1,\dots,T-1} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathcal{H}} \leq D$  for some  $D > 0$  and Assumption 4.2 holds with adaptive gradient variance  $\sigma_{\mathcal{H}}(\{f_t\}_{t=1}^T)^2 \leq \sigma_{\mathcal{H}}^2$  for some  $\sigma_{\mathcal{H}} \in [0, \infty)$ . Then it holds that

$$\begin{aligned} \mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] &\leq \frac{2D^2\epsilon}{\eta(T+1)^2} \mathbb{E} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left(\frac{D^2}{2\eta} - \frac{\eta}{2}\right) \mathbb{E} \frac{4}{(T+1)^2} \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \\ &\quad + \frac{2\eta^2}{(T+1)^2} \cdot \Lambda_{\mathcal{H}}(f) \cdot \tilde{O}(\log^2 d) + \frac{\eta}{T^{1/2}} \sigma_{\mathcal{H}} \cdot \tilde{O}(\log d). \end{aligned}$$

Moreover, when choosing learning rate  $\eta = D$ , the convergence rate becomes

$$\mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] = \tilde{O}\left(\frac{\Lambda_{\mathcal{H}}(f) D^2 \log^2 d + d\sqrt{\epsilon}D}{T^2} + \frac{\sigma_{\mathcal{H}} D \log d}{\sqrt{T}}\right).$$

**Remark 4.4.** A drawback of Algorithm 2 is that the optimal choice of learning rate  $\eta$  in Theorem 4.3 depends on the unknown parameter  $D = \max_t \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathcal{H}}$ . To circumvent this issue, we follow the approach in Kovalev (2025a) to introduce a projected variant (see Algorithm 8 and discussion in Appendix E.2), which ensures that all iterates remain inside a  $\|\cdot\|_{\mathcal{H}}$ -ball of radius  $D$ . This removes the requirement for prior knowledge of  $D$ , and we establish a same convergence rate in Theorem E.5.

Our convergence guarantee attains a deterministic component of order  $\tilde{O}(\Lambda_{\mathcal{H}}(f) D^2/T^2)$ , an accelerated rate governed by the adaptive smoothness  $\Lambda_{\mathcal{H}}(f)$ . In comparison, Guzmán & Nemirovski (2015) shows that any first order optimizer can only achieve  $\Omega(\frac{L_{\|\cdot\|_{\infty}}(f)}{T \log T})$  for the specific case of  $\ell_{\infty}$  norm smoothness. Taken together, these results show that the adaptive smoothness is necessary to achieve the acceleration, which can't be replaced by the weaker non-Euclidean smoothness, highlighting its algorithmic benefit and answering Question 2 in the affirmative.

The analysis in Kovalev (2025a) yields similar results when  $\mathcal{H}$  contains only diagonal matrices. However, their Assumption 4, which is critical for their analysis, imposes restrictive conditions on both the loss and the gradient noise for more general  $\mathcal{H}$ . By contrast, our approach avoids such requirements by leveraging Lemma 3.3.

### 4.3 NONCONVEX RESULTS FOR NSD UNDER ADAPTIVE NOISE ASSUMPTION

The ineffectiveness of averaging in the dual space also leads to difficulty in reducing gradient variance via averaging. To illustrate this, consider i.i.d. random vectors  $\mathbf{x}_1, \dots, \mathbf{x}_n$  with  $\mathbb{E}[\mathbf{x}_i] = 0$  and  $\mathbb{E}[\|\mathbf{x}_i\|_2^2] \leq \sigma^2$ , we have  $\mathbb{E}[\|\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i\|_2^2] \leq \frac{\sigma^2}{n}$  while  $\mathbb{E}[\|\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i\|_1^2] \leq d \frac{\sigma^2}{n}$ , and the extra  $d$  factor in the latter bound can be tight. This causes the dimension-dependent convergence rate of NSD in the nonconvex setting, as shown in recent works (Pethick et al., 2025; Kovalev, 2025b).

In particular, the dimension-dependent factor  $\rho = \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_{\mathcal{H},*}}{\|\mathbf{x}\|_2}$  captures the mismatch between  $\|\cdot\|_{\mathcal{H},*}$  and  $\|\cdot\|_2$ . For diagonal  $\mathcal{H}$ , NSD reduces to SignGD with  $\|\cdot\|_{\mathcal{H},*} = \|\cdot\|_1$ , yielding  $\rho = \Theta(\sqrt{d})$  and vacuous bounds when  $T \ll d$ . We avoid this by using the adaptive variance assumption. We prove the stochastic rate of NSD in Theorem 4.5 and the proof is deferred to Appendix F.1.

The concurrent work Kovalev & Borodich (2025) also leveraged the adaptive variance assumption to prove a dimension-free nonconvex convergence rate of NSD. However, they used a smoothness metric similar to adaptive smoothness while our Theorem 4.5 uses the standard smoothness. Our rate is strictly better because of the relationship between standard smoothness and adaptive smoothness.

**Theorem 4.5.** Let  $\mathcal{H}$  be a well-structured preconditioner set. For any  $\epsilon \geq 0$ ,  $\alpha \in (0, 1)$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 3 with  $\|\cdot\| = \|\cdot\|_{\mathcal{H}}$  and  $\mathbf{m}_0 = \nabla f_0(\mathbf{x}_0)$ . Let  $L_{\|\cdot\|_{\mathcal{H}}}(f)$  be the smoothness of the loss  $f$  w.r.t.  $\|\cdot\|_{\mathcal{H}}$  according to Definition 2.3. Suppose Assumption 4.2 holds with adaptive gradient variance  $\sigma_{\mathcal{H}}(\{f_t\}_{t=1}^T)^2 \leq \sigma_{\mathcal{H}}^2$ . Then it holds that

$$\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \leq \frac{\Delta_0}{\eta T} + \frac{2\eta}{\alpha} L_{\|\cdot\|}(f) + \frac{2\sigma_{\mathcal{H}}}{\alpha T} + 2\sigma_{\mathcal{H}}\sqrt{\alpha}.$$

Let  $a_0 = \sqrt{\Delta_0 L_{\|\cdot\|}(f)/\sigma_{\mathcal{H}}}$ . If  $a_0 < 1$ , then

- When  $T < a_0^{-6}$ , we choose  $\alpha = T^{-2/3}$  and  $\eta = \sqrt{\frac{\Delta_0}{L_{\|\cdot\|_{\mathcal{H}}}(f)}} T^{-5/12}$ . The rate is  $O(\sigma_{\mathcal{H}} T^{-1/3})$ .
  - When  $T \geq a_0^{-6}$ , we choose  $\alpha = \frac{a_0}{\sqrt{T}}$ ,  $\eta = \frac{\Delta_0^{3/4} T^{-3/4}}{L_{\|\cdot\|_{\mathcal{H}}}(f)^{1/4} \sigma_{\mathcal{H}}^{1/2}}$ . The rate is  $O\left(\frac{(\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f))^{1/4} \sqrt{\sigma_{\mathcal{H}}}}{T^{1/4}}\right)$ .
- If  $a_0 \geq 1$ , then
- When  $T \leq a_0^2$ , we choose  $\alpha = 1$  and  $\eta = \sqrt{\frac{\Delta_0}{L_{\|\cdot\|_{\mathcal{H}}}(f)}} T^{-1/2}$ . The rate is  $O\left(\sqrt{\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f)} T^{-1/2}\right)$ .
  - When  $T \geq a_0^2$ , we choose  $\alpha = \frac{a_0}{\sqrt{T}}$  and  $\eta = \frac{\Delta_0^{3/4} T^{-3/4}}{L_{\|\cdot\|_{\mathcal{H}}}(f)^{1/4} \sigma_{\mathcal{H}}^{1/2}}$ . The rate is  $O\left(\frac{(\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f))^{1/4} \sqrt{\sigma_{\mathcal{H}}}}{T^{1/4}}\right)$ .

Theorem 4.5 shows that, with appropriate choices of the learning rate  $\eta$  and averaging parameter  $\alpha$ , NSD achieves a dimension-free rate depending only on the standard smoothness  $L_{\|\cdot\|_{\mathcal{H}}}(f)$  and the adaptive variance  $\sigma_{\mathcal{H}}$ , thereby avoiding the unfavorable  $\rho$  factor. Next Theorem 4.6 shows that such a dimension-free upper bound is unattainable under the standard gradient variance assumption.

**Theorem 4.6.** For any  $\epsilon \geq 0$ ,  $\alpha \in (0, 1)$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 3 with any norm  $\|\cdot\|$ . Suppose Assumption 4.2 holds with the gradient variance  $\sigma_{\|\cdot\|_*}(\{f_t\}_{t=1}^T)^2 \leq \sigma_{\|\cdot\|_*}^2$  for some  $\sigma_{\|\cdot\|_*} \in [0, \infty)$ . Then it holds that

$$\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_* \leq \frac{\Delta_0}{\eta T} + \frac{2\eta}{\alpha} L_{\|\cdot\|}(f) + \frac{2}{\alpha T} \sigma_{\|\cdot\|_*} + 2\sigma_{\|\cdot\|_*} \cdot \min\left(1, \alpha^{1/2} \psi(\|\cdot\|_*, \|\cdot\|_2)\right)$$

where  $\psi(\|\cdot\|_*, \|\cdot\|_2) = \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_*}{\|\mathbf{x}\|_2} \cdot \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_2}{\|\mathbf{x}\|_*}$  measures the distortion between the two norms.

Here the norm distortion  $\psi(\|\cdot\|_*, \|\cdot\|_2)$  can grow with the dimension  $d$ . Consequently, Theorem 4.6 gives two kinds of upper bound for the convergence rate of NSD:

- When  $T$  is small (e.g.  $T < d$ ), the constant term  $\sigma_{\|\cdot\|_*}$  dominates.
- For sufficiently large  $T$  and small  $\alpha$  (for which the four terms are balanced), the last term in the upper bound becomes  $2\sigma_{\|\cdot\|_*} \alpha^{1/2} \psi(\|\cdot\|_*, \|\cdot\|_2)$ , which depends on  $d$ .

Moreover, such a dependence on  $d$  is inevitable in the worst case as shown by Theorem 4.7.

**Theorem 4.7.** For any fixed  $\Delta_0, L, \sigma^2, d, T$ , learning rate  $\eta$ , and any averaging parameter  $\alpha$ , there exists a loss function  $f: \mathbb{R}^d \rightarrow \mathbb{R}$ , a sequence of stochastic iid loss functions  $f_0, f_1, \dots, f_{T-1}$  and an initialization  $\mathbf{x}_0$  satisfying the following conditions: (1)  $f(\mathbf{x}_0) - \inf_{\mathbf{x}} f(\mathbf{x}) = \Delta_0$  and  $L_{\|\cdot\|_{\infty}}(f) \leq L$ ; (2) For any  $\mathbf{x} \in \mathbb{R}^d$ , it holds that  $\mathbb{E}[\nabla f_t(\mathbf{x})] = \nabla f(\mathbf{x})$  and  $\mathbb{E}[\|\nabla f_t(\mathbf{x}) - \nabla f(\mathbf{x})\|_1^2] \leq \sigma^2$ .

When running Algorithm 3 with  $\|\cdot\| = \|\cdot\|_{\infty}$ , learning rate  $\eta$ , averaging parameter  $\alpha$  and initialization  $\mathbf{x}_0 = \mathbf{0}$ , it holds that

$$\mathbb{E} \left[ \min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_1 \right] = \min\{e^{-2} 5^{-\frac{1}{4}} (dL\Delta_0\sigma^2)^{\frac{1}{4}} T^{-\frac{1}{2}}, e^{-2} 5^{-\frac{1}{2}} \sigma\}$$

Theorem 4.7 also shows two kinds of lower bound we can achieve on signGD with momentum:

- When  $T$  is not large enough, we can achieve the lower bound  $\Omega(\sigma)$ , which shows the hardness induced by the stochasticity and matches the first upper bound in Theorem 4.6.
- On the other hand, if we want to achieve the error  $\epsilon < e^{-2} 5^{-\frac{1}{2}} \sigma$ , we require the number of steps  $T = \Omega(\epsilon^{-2} (dL\Delta_0\sigma^2)^{\frac{1}{2}})$ , whose dependence on dimension  $d$  is  $\Omega(d^{\frac{1}{2}})$ .

In conclusion, under the standard gradient variance assumption with  $\|\cdot\| = \|\cdot\|_{\infty}$  and  $\|\cdot\|_* = \|\cdot\|_1$ , the  $d$ -dependent rate in Theorem 4.6 is unavoidable. In contrast, Theorem 4.5 attains a dimension-free rate under the adaptive gradient variance assumption, highlighting a fundamental gap.

## 5 CONCLUSION

We extend the unified analysis of adaptive optimizers to nonconvex functions, establishing convergence rate depending on the adaptive smoothness. It strengthens the comparison between smoothnesses that adaptive optimizers and NSD use in the convex settings. We further show the benefit of adaptive smoothness by showing the accelerated rate of adaptive optimizers with Nesterov momentum. The benefit of adaptive geometry is also justified by comparing two kinds of noise.

## REFERENCES

- Zeyuan Allen-Zhu and Lorenzo Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. *arXiv preprint arXiv:1407.1537*, 2014. 15
- Kang An, Yuxing Liu, Rui Pan, Yi Ren, Shiqian Ma, Donald Goldfarb, and Tong Zhang. Asgo: Adaptive structured gradient optimization. *arXiv preprint arXiv:2503.20762*, 2025. 1, 5, 8, 15
- Lukas Balles, Fabian Pedregosa, and Nicolas Le Roux. The geometry of sign gradient descent. *arXiv preprint arXiv:2002.08056*, 2020. 1, 3, 15
- Jeremy Bernstein and Laker Newhouse. Old optimizer, new norm: An anthology. *arXiv preprint arXiv:2409.20325*, 2024. 1
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiusi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024. 1
- Lizhang Chen, Jonathan Li, and Qiang Liu. Muon Optimizes Under Spectral Norm Constraints. *arXiv preprint arXiv:2506.15054*, 2025. 15
- Xiangning Chen, Chen Liang, Da Huang, Esteban Real, Kaiyuan Wang, Yao Liu, Hieu Pham, Xuanyi Dong, Thang Luong, Cho-Jui Hsieh, et al. Symbolic discovery of optimization algorithms. *arXiv preprint arXiv:2302.06675*, 2023. 1
- Xiang Cheng, Fred Roosta, Stefan Palombo, Peter Bartlett, and Michael Mahoney. FLAG n’FLARE: Fast Linearly-Coupled Adaptive Gradient Methods. In *International Conference on Artificial Intelligence and Statistics*, pp. 404–414. PMLR, 2018. 15
- Sinho Chewi, Sébastien Bubeck, and Adil Salim. On the complexity of finding stationary points of smooth functions in one dimension. In *International Conference on Algorithmic Learning Theory*, pp. 358–374. PMLR, 2023. 45
- Ashok Cutkosky. Anytime online-to-batch, optimism and acceleration. In *International conference on machine learning*, pp. 1446–1454. PMLR, 2019. 15
- Ashok Cutkosky and Harsh Mehta. Momentum improves normalized sgd. In *International conference on machine learning*. PMLR, 2020. 15
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv e-prints*, pp. arXiv–2407, 2024. 1
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 2011. 5
- Alina Ene, Huy L Nguyen, and Adrian Vladu. Adaptive gradient methods for constrained convex optimization and variational inequalities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021. 15
- Kevin Frans, Pieter Abbeel, and Sergey Levine. What Really Matters in Matrix-Whitening Optimizers? *arXiv preprint arXiv:2510.25000*, 2025. 15
- Vineet Gupta, Tomer Koren, and Yoram Singer. A unified approach to adaptive regularization in online and stochastic optimization. *arXiv preprint arXiv:1706.06569*, 2017. 5, 38
- Vineet Gupta, Tomer Koren, and Yoram Singer. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning*, 2018. 15
- Cristóbal Guzmán and Arkadi Nemirovski. On lower complexity bounds for large-scale smooth convex optimization. *Journal of Complexity*, 31(1):1–14, 2015. 2, 9
- Wei Jiang, Dingzhi Yu, Sifan Yang, Wenhao Yang, and Lijun Zhang. Improved Analysis for Sign-based Methods with Momentum Updates. *arXiv preprint arXiv:2507.12091*, 2025. 15

- Keller Jordan, Yuchen Jin, Vlado Boza, Jiacheng You, Franz Cesista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks, 2024. 1
- Pooria Joulani, Anant Raj, Andras Gyorgy, and Csaba Szepesvari. A simpler approach to accelerated optimization: iterative averaging meets optimism. In *International conference on machine learning*. PMLR, 2020. 15
- Ali Kavis, Kfir Y Levy, Francis Bach, and Volkan Cevher. Unixgrad: A universal, adaptive algorithm with optimal guarantees for constrained optimization. *Advances in neural information processing systems*, 32, 2019. 15
- Jonathan A Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pp. 217–226. SIAM, 2014. 15
- Dmitry Kovalev. SGD with Adaptive Preconditioning: Unified Analysis and Momentum Acceleration. *arXiv preprint arXiv:2506.23803*, 2025a. 6, 7, 8, 9, 34
- Dmitry Kovalev. Understanding gradient orthogonalization for deep learning via non-euclidean trust-region optimization. *arXiv preprint arXiv:2503.12645*, 2025b. 9, 15, 41
- Dmitry Kovalev and Ekaterina Borodich. On linear convergence in smooth convex-concave bilinearly-coupled saddle-point optimization: Lower bounds and optimal algorithms. *arXiv preprint arXiv:2411.14601*, 2024. 8
- Dmitry Kovalev and Ekaterina Borodich. Non-Euclidean SGD for Structured Optimization: Unified Analysis and Improved Rates. *arXiv preprint arXiv:2511.11466*, 2025. 9
- Tim Tsz-Kit Lau, Qi Long, and Weijie Su. PolarGrad: A Class of Matrix-Gradient Optimizers from a Unifying Preconditioning Perspective. *arXiv preprint arXiv:2505.21799*, 2025. 15
- Huan Li, Yiming Dong, and Zhouchen Lin. On the  $O(\frac{\sqrt{d}}{K^{1/4}})$  Convergence Rate of AdamW Measured by  $\ell_1$  Norm. *arXiv preprint arXiv:2505.11840*, 2025. 32
- Elliott H. Lieb. Convex trace functions and the wigner-yanase-dyson conjecture. *Advances in Mathematics*, 11(3):267–288, 1973. 17
- Jingyuan Liu, Jianlin Su, Xingcheng Yao, Zhejun Jiang, Guokun Lai, Yulun Du, Yidao Qin, Weixin Xu, Enzhe Lu, Junjie Yan, et al. Muon is scalable for LLM training. *arXiv preprint arXiv:2502.16982*, 2025. 1
- Devyani Maladkar, Ruichen Jiang, and Aryan Mokhtari. Convergence Analysis of Adaptive Gradient Methods under Refined Smoothness and Noise Assumptions. *arXiv preprint arXiv:2406.04592*, 2024. 1, 3
- Depen Morwani, Itai Shapira, Nikhil Vyas, Eran Malach, Sham Kakade, and Lucas Janson. A New Perspective on Shampoo’s Preconditioner. *arXiv preprint arXiv:2406.17748*, 2024. 15
- Thomas Pethick, Wanyun Xie, Kimon Antonakopoulos, Zhenyu Zhu, Antonio Silveti-Falls, and Volkan Cevher. Training deep learning models with norm-constrained lmos. *arXiv preprint arXiv:2502.07529*, 2025. 1, 6, 9, 15
- Maria-Eleni Sfyraiki and Jun-Kun Wang. Lions and muons: Optimization via stochastic frank-wolfe. *arXiv preprint arXiv:2506.04192*, 2025. 15
- Ishaan Shah, Anthony M Polloreno, Karl Stratos, Philip Monk, Adarsh Chaluvvaraju, Andrew Hojel, Andrew Ma, Anil Thomas, Ashish Tanwer, Darsh J Shah, et al. Practical efficiency of muon for pretraining. *arXiv preprint arXiv:2505.02222*, 2025. 1
- Chongjie Si, Debing Zhang, and Wei Shen. Adamuon: Adaptive muon optimizer. *arXiv preprint arXiv:2507.11005*, 2025. 15

- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025. 1
- Nikhil Vyas, Depen Morwani, Rosie Zhao, Itai Shapira, David Brandfonbrener, Lucas Janson, and Sham Kakade. Soap: Improving and stabilizing shampoo using adam. *arXiv preprint arXiv:2409.11321*, 2024. 15
- Jiaxuan Wang and Jenna Wiens. AdaSGD: Bridging the gap between SGD and Adam. *arXiv preprint arXiv:2006.16541*, 2020. 5
- Kaiyue Wen, David Hall, Tengyu Ma, and Percy Liang. Fantastic Pretraining Optimizers and Where to Find Them. *arXiv preprint arXiv:2509.02046*, 2025. 1
- Shuo Xie and Zhiyuan Li. Implicit Bias of AdamW:  $\ell_\infty$  Norm Constrained Optimization. *arXiv preprint arXiv:2404.04454*, 2024. 1
- Shuo Xie, Mohamad Amin Mohamadi, and Zhiyuan Li. Adam Exploits  $\ell_\infty$ -geometry of Loss Landscape via Coordinate-wise Adaptivity. *arXiv preprint arXiv:2410.08198*, 2024. 3
- Shuo Xie, Mohamad Amin Mohamadi, and Zhiyuan Li. Adam Exploits  $\ell_\infty$ -geometry of Loss Landscape via Coordinate-wise Adaptivity. In *The Thirteenth International Conference on Learning Representations*, 2025a. 1, 3, 5, 6, 7, 29, 32
- Shuo Xie, Tianhao Wang, Sashank J. Reddi, Sanjiv Kumar, and Zhiyuan Li. Structured Preconditioners in Adaptive Optimization: A Unified Analysis. In *Forty-second International Conference on Machine Learning*, 2025b. 1, 2, 4, 5, 6, 7, 8, 15, 17, 19, 26
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 1
- Chenyang Zhang, Difan Zou, and Yuan Cao. The implicit bias of adam on separable data. *Advances in Neural Information Processing Systems*, 37:23988–24021, 2024. 1

## CONTENTS

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Notations . . . . .	3
<b>2</b>	<b>From Adam/SignGD to Adaptive Smoothness</b>	<b>3</b>
2.1	Adam and SignGD can exploit $\ell_\infty$ geometry, but in different ways . . . . .	3
2.2	Adaptive smoothness associated with well-structured preconditioner . . . . .	4
<b>3</b>	<b>Unified Analysis in the Nonconvex Setting</b>	<b>5</b>
3.1	Adaptive optimizers with well-structured preconditioner sets . . . . .	5
3.2	Convergence rate in the deterministic nonconvex setting . . . . .	6
3.3	Technical Contribution: A Novel Matrix Inequality . . . . .	7
<b>4</b>	<b>Benefit of Adaptive Geometry</b>	<b>7</b>
4.1	Adaptive Variance: An Analogue of Adaptive Smoothness . . . . .	8
4.2	Acceleration on Convex Functions . . . . .	8
4.3	Nonconvex Results for NSD under Adaptive Noise Assumption . . . . .	9
<b>5</b>	<b>Conclusion</b>	<b>10</b>
<b>A</b>	<b>Related Work</b>	<b>15</b>
<b>B</b>	<b>Auxiliary Results</b>	<b>15</b>
B.1	Comparison on adaptive geometry . . . . .	18
<b>C</b>	<b>Proof for the matrix inequality</b>	<b>20</b>
<b>D</b>	<b>Unified proof for adaptive algorithms</b>	<b>25</b>
D.1	Relationship between adaptive algorithms . . . . .	25
D.2	Proof for weighted algorithm . . . . .	26
D.3	Proof for the cumulative and EMA variants . . . . .	32
<b>E</b>	<b>Proof for the accelerated algorithm</b>	<b>34</b>
E.1	Stochastic case without projection . . . . .	34
E.2	Stochastic case with projection . . . . .	38
<b>F</b>	<b>Proof for normalized steepest descent</b>	<b>39</b>
F.1	Convergence rate with respect to adaptive gradient variance . . . . .	41
F.2	Convergence rate under general-norm assumption . . . . .	43
F.3	Dimension dependent lower bound for NSD . . . . .	44

## A RELATED WORK

**Adaptive optimizers.** Matrix-valued preconditioning methods such as Shampoo (Gupta et al., 2018) provide structure-aware updates and a growing body of work refines their preconditioners and practice (e.g., improved analyses/implementations and Adam–Shampoo hybrids) (Morwani et al., 2024; Vyas et al., 2024; Lau et al., 2025; Si et al., 2025). Variants like one-sided Shampoo/ASGO achieve stronger convex rates than earlier methods (Xie et al., 2025b; An et al., 2025), but the nonconvex guarantees are missing which we show in Section 3. The recent work Frans et al. (2025) views adaptive optimizers as a modified matrix whitening technique and points out it can outperform accurate spectral normalization (NSD in our context). They highlight the variance adaptation is previously overlooked while we also want to understand the benefit of adaptive optimizers/geometry.

**NSD.** Cutkosky & Mehta (2020) proves the convergence rate  $O(1/T^{3.5})$  of normalized gradient descent with momentum on nonconvex functions. Pethick et al. (2025); Kovalev (2025b) extend the results to any normalized steepest descent while the smoothness metric becomes smoothness w.r.t. the specific norm NSD uses. Many works analyze popular optimizers by showing the equivalence to NSD under some norm, e.g. Lion is NSD under  $\ell_\infty$ -norm (Sfyraki & Wang, 2025) and Muon is NSD under spectral norm (Chen et al., 2025), and thus apply the classic optimization results. Jiang et al. (2025) achieves a similar result as our Theorem 4.5 but their noise assumption is stronger than ours and only states the result for specific SignGD. The difference between the two kinds of smoothness is also discussed in Balles et al. (2020) in the context for SignGD only.

**Acceleration.** Nesterov acceleration can be viewed as the linear coupling between gradient descent and mirror descent (Kelner et al., 2014; Allen-Zhu & Orecchia, 2014; Cheng et al., 2018). Through this lens, Cutkosky (2019); Kavis et al. (2019); Joulani et al. (2020); Ene et al. (2021) applied it on adaptive optimizers with specific structures like diagonal/coordinate-wise and achieve deterministic  $O(1/T^2)$  convergence rate.

## B AUXILIARY RESULTS

Recall that we define  $P_{\mathcal{H}}(\mathbf{M}) = \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$  for any  $\mathbf{M} \in \mathcal{S}_{++}^d$  in Section 2. In this section, we will prove several important properties of  $P_{\mathcal{H}}$  which will be useful in proving the convergence rate for general well-structured  $\mathcal{H}$ .

It will help interpret these results when keeping in mind that  $P_{\mathcal{H}}(\mathbf{M}) = \mathbf{M}^{1/2}$  when  $\mathcal{H}$  is the PSD matrix cone, i.e., there is no constraint on the structure.

We will first state Lemma A.2 in Xie et al. (2025b), which will be used a lot in our proof.

**Lemma B.1** (Lemma A.2 in Xie et al. (2025b)). *For any  $\mathbf{A}, \mathbf{B} \in \mathcal{H}$ , if  $\langle \mathbf{A}, \mathbf{H} \rangle = \langle \mathbf{B}, \mathbf{H} \rangle$  holds for all  $\mathbf{H} \in \mathcal{H}$ , then it must be  $\mathbf{A} = \mathbf{B}$ .*

**Lemma B.2.** *For any  $\mathbf{M} \in \mathcal{S}_{++}^d$  and any  $\mathbf{H} \in \mathcal{H}$ ,  $\lambda \geq 0$ , it always holds that  $\langle \mathbf{M}, \mathbf{H} \rangle = \langle P_{\mathcal{H}}(\mathbf{M})^2, \mathbf{H} \rangle$  and  $P_{\mathcal{H}}(\mathbf{M} + \lambda \mathbf{I}_d)^2 = P_{\mathcal{H}}(\mathbf{M})^2 + \lambda \mathbf{I}_d$ .*

*Proof.* By property (c) in Proposition A.3 in Xie et al. (2025b), we know that for any  $\mathbf{M} \succ 0$ ,

$$\langle \mathbf{M} - P_{\mathcal{H}}(\mathbf{M})^2, P_{\mathcal{H}}(\mathbf{M})^{-1} \mathbf{H} P_{\mathcal{H}}(\mathbf{M})^{-1} - P_{\mathcal{H}}(\mathbf{M})^{-1} \rangle = 0.$$

Note that  $\mathbf{H} \mapsto P_{\mathcal{H}}(\mathbf{M})^{-1} \mathbf{H} P_{\mathcal{H}}(\mathbf{M})^{-1}$  is a bijection from  $\mathcal{H}$  to  $\mathcal{H}$  by lemma A.1 in Xie et al. (2025b), so we have  $\langle \mathbf{M} - P_{\mathcal{H}}(\mathbf{M})^2, \mathbf{H} - P_{\mathcal{H}}(\mathbf{M})^{-1} \rangle = 0$  for all  $\mathbf{H} \in \mathcal{H}$ .  $\mathbf{H} \mapsto \mathbf{H} - P_{\mathcal{H}}(\mathbf{M})^{-1}$  is also a bijection from  $\mathcal{H}$  to  $\mathcal{H}$ . So we have that  $\langle \mathbf{M} - P_{\mathcal{H}}(\mathbf{M})^2, \mathbf{H} \rangle = 0$  for all  $\mathbf{H} \in \mathcal{H}$ .

For any  $\lambda \geq 0$  and any  $\mathbf{H} \in \mathcal{H}$ , we have that

$$\begin{aligned} \langle P_{\mathcal{H}}(\mathbf{M} + \lambda \mathbf{I}_d)^2, \mathbf{H} \rangle &= \langle \mathbf{M} + \lambda \mathbf{I}_d, \mathbf{H} \rangle \\ &= \langle \mathbf{M}, \mathbf{H} \rangle + \langle \lambda \mathbf{I}_d, \mathbf{H} \rangle \\ &= \langle P_{\mathcal{H}}(\mathbf{M})^2, \mathbf{H} \rangle + \langle \lambda \mathbf{I}_d, \mathbf{H} \rangle \\ &= \langle P_{\mathcal{H}}(\mathbf{M})^2 + \lambda \mathbf{I}_d, \mathbf{H} \rangle. \end{aligned}$$

Note that  $P_{\mathcal{H}}(\mathbf{M} + \lambda \mathbf{I}_d)^2$  and  $P_{\mathcal{H}}(\mathbf{M})^2 + \lambda \mathbf{I}_d$  are both in  $\mathcal{H}$ . We conclude  $P_{\mathcal{H}}(\mathbf{M} + \lambda \mathbf{I}_d)^2 = P_{\mathcal{H}}(\mathbf{M})^2 + \lambda \mathbf{I}_d$  according to Lemma B.1.  $\square$

**Lemma B.3.** For any  $\mathbf{M}_1, \mathbf{M}_2 \in \mathcal{S}_{++}^d$ ,  $P_{\mathcal{H}}(\mathbf{M}_1 + \mathbf{M}_2)^2 = P_{\mathcal{H}}(\mathbf{M}_1)^2 + P_{\mathcal{H}}(\mathbf{M}_2)^2$ . For any  $c \geq 0$ ,  $P_{\mathcal{H}}(c\mathbf{M})^2 = cP_{\mathcal{H}}(\mathbf{M})^2$ . For any  $\mathbf{0} \preceq \mathbf{A} \preceq \mathbf{B}$ , it always holds  $P_{\mathcal{H}}(\mathbf{A}) \preceq P_{\mathcal{H}}(\mathbf{B})$ .

*Proof.* For any  $\mathbf{H} \in \mathcal{H}$ , we have that

$$\begin{aligned} \langle P_{\mathcal{H}}(\mathbf{M}_1 + \mathbf{M}_2)^2, \mathbf{H} \rangle &= \langle \mathbf{M}_1 + \mathbf{M}_2, \mathbf{H} \rangle \\ &= \langle \mathbf{M}_1, \mathbf{H} \rangle + \langle \mathbf{M}_2, \mathbf{H} \rangle \\ &= \langle P_{\mathcal{H}}(\mathbf{M}_1)^2, \mathbf{H} \rangle + \langle P_{\mathcal{H}}(\mathbf{M}_2)^2, \mathbf{H} \rangle \\ &= \langle P_{\mathcal{H}}(\mathbf{M}_1)^2 + P_{\mathcal{H}}(\mathbf{M}_2)^2, \mathbf{H} \rangle. \end{aligned}$$

Then we have that  $P_{\mathcal{H}}(\mathbf{M}_1 + \mathbf{M}_2)^2 = P_{\mathcal{H}}(\mathbf{M}_1)^2 + P_{\mathcal{H}}(\mathbf{M}_2)^2$  from Lemma B.1. For any  $c \geq 0$ , it holds that

$$\begin{aligned} \langle P_{\mathcal{H}}(c\mathbf{M})^2, \mathbf{H} \rangle &= \langle c\mathbf{M}, \mathbf{H} \rangle \\ &= c \langle \mathbf{M}, \mathbf{H} \rangle \\ &= c \langle P_{\mathcal{H}}(\mathbf{M})^2, \mathbf{H} \rangle \\ &= \langle cP_{\mathcal{H}}(\mathbf{M})^2, \mathbf{H} \rangle. \end{aligned}$$

So it also holds that  $P_{\mathcal{H}}(c\mathbf{M})^2 = cP_{\mathcal{H}}(\mathbf{M})^2$  from Lemma B.1.

Moreover, when  $\mathbf{0} \preceq \mathbf{A} \preceq \mathbf{B}$ ,  $P_{\mathcal{H}}(\mathbf{B})^2 = P_{\mathcal{H}}(\mathbf{A})^2 + P_{\mathcal{H}}(\mathbf{B} - \mathbf{A})^2 \succeq P_{\mathcal{H}}(\mathbf{A})^2$ . Then  $P_{\mathcal{H}}(\mathbf{B}) \succeq P_{\mathcal{H}}(\mathbf{A})$  because  $\mathbf{X}^{1/2}$  is a monotone function.  $\square$

**Lemma B.4.** If we define  $\text{proj}_{\mathcal{H}}(\mathbf{M}) = \arg \min_{\mathbf{H} \in \mathcal{H}} \|\mathbf{M} - \mathbf{H}\|_{\mathbb{F}}^2$ , then it holds that  $\text{proj}_{\mathcal{H}}(\mathbf{M}) = P_{\mathcal{H}}(\mathbf{M})^2$ .

*Proof.* We define  $d_{\mathbf{M}}(\mathbf{H}) = \|\mathbf{M} - \mathbf{H}\|_{\mathbb{F}}^2 = \text{Tr}((\mathbf{M} - \mathbf{H})(\mathbf{M} - \mathbf{H})^{\top})$ . Then  $\nabla_{\mathbf{H}} d_{\mathbf{M}}(\mathbf{H}) = 2(\mathbf{H} - \mathbf{M})$ . Because  $\mathcal{H}$  is a cone, by the optimality of  $\text{proj}_{\mathcal{H}}(\mathbf{M})$ , it always holds for any  $\mathbf{H} \in \mathcal{H}$  that

$$0 = \langle \nabla_{\mathbf{H}} d_{\mathbf{M}}(\text{proj}_{\mathcal{H}}(\mathbf{M})), \mathbf{H} - \text{proj}_{\mathcal{H}}(\mathbf{M}) \rangle = 2 \langle \text{proj}_{\mathcal{H}}(\mathbf{M}) - \mathbf{M}, \mathbf{H} - \text{proj}_{\mathcal{H}}(\mathbf{M}) \rangle.$$

For any  $\mathbf{H}' \in \mathcal{H}$ , we can always choose  $\mathbf{H} = \mathbf{H}' + \text{proj}_{\mathcal{H}}(\mathbf{M})$ .  $\mathbf{H} \in \mathcal{H}$  because  $\mathcal{H}$  is closed under matrix addition. Then it always holds that  $\langle \text{proj}_{\mathcal{H}}(\mathbf{M}) - \mathbf{M}, \mathbf{H}' \rangle = 0$ . On the other hand, Lemma B.2 shows it always holds that  $\langle P_{\mathcal{H}}(\mathbf{M})^2 - \mathbf{M}, \mathbf{H}' \rangle = 0$  for any  $\mathbf{H}' \in \mathcal{H}$ . Then we can get  $\text{proj}_{\mathcal{H}}(\mathbf{M}) = P_{\mathcal{H}}(\mathbf{M})^2$  by Lemma B.1.  $\square$

**Lemma B.5.** For any  $\mathbf{A}, \mathbf{B} \in \mathcal{S}_{++}^d$  and any well-structured preconditioner set  $\mathcal{H}$ ,

$$\|P_{\mathcal{H}}(\mathbf{B}) - P_{\mathcal{H}}(\mathbf{A})\|_{\text{op}} \leq \|\mathbf{B} - \mathbf{A}\|_{\text{op}}^{\frac{1}{2}}$$

*Proof.* We have that

$$\begin{aligned} P_{\mathcal{H}}(\mathbf{B}) &= P_{\mathcal{H}}(\mathbf{A} + \mathbf{B} - \mathbf{A}) \preceq P_{\mathcal{H}}(\mathbf{A} + \|\mathbf{B} - \mathbf{A}\|_{\text{op}} \mathbf{I}) \\ &= \left( P_{\mathcal{H}}(\mathbf{A})^2 + \|\mathbf{B} - \mathbf{A}\|_{\text{op}} \mathbf{I} \right)^{\frac{1}{2}} \preceq P_{\mathcal{H}}(\mathbf{A}) + \|\mathbf{B} - \mathbf{A}\|_{\text{op}}^{\frac{1}{2}} \mathbf{I} \end{aligned}$$

The last step is by considering  $\sqrt{\lambda_i^2 + \|\mathbf{B} - \mathbf{A}\|_{\text{op}}} \leq \lambda_i + \sqrt{\|\mathbf{B} - \mathbf{A}\|_{\text{op}}}$  for every eigenvalue  $\lambda_i$  of  $P_{\mathcal{H}}(\mathbf{A})$ . Similarly we can also show  $P_{\mathcal{H}}(\mathbf{A}) \preceq P_{\mathcal{H}}(\mathbf{B}) + \|\mathbf{B} - \mathbf{A}\|_{\text{op}}^{\frac{1}{2}} \mathbf{I}$ , which finishes the proof.  $\square$

**Lemma B.6.**  $\text{Tr}[P_{\mathcal{H}}(\mathbf{X})]$  is a concave function.

*Proof.* For any psd matrices  $\mathbf{A}$  and  $\mathbf{B}$ , we have that

$$\begin{aligned} 2 \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \frac{\mathbf{A} + \mathbf{B}}{2} \right) \right] &= \min_{\mathbf{H} \in \mathcal{H}} \left[ \left\langle \frac{\mathbf{A} + \mathbf{B}}{2}, \mathbf{H}^{-1} \right\rangle + \operatorname{Tr}(\mathbf{H}) \right] \\ &= \min_{\mathbf{H} \in \mathcal{H}} \left[ \frac{\langle \mathbf{A}, \mathbf{H}^{-1} \rangle + \operatorname{Tr}(\mathbf{H})}{2} + \frac{\langle \mathbf{B}, \mathbf{H}^{-1} \rangle + \operatorname{Tr}(\mathbf{H})}{2} \right] \\ &\geq \min_{\mathbf{H} \in \mathcal{H}} \frac{\langle \mathbf{A}, \mathbf{H}^{-1} \rangle + \operatorname{Tr}(\mathbf{H})}{2} + \min_{\mathbf{H} \in \mathcal{H}} \frac{\langle \mathbf{B}, \mathbf{H}^{-1} \rangle + \operatorname{Tr}(\mathbf{H})}{2} \\ &= \operatorname{Tr}[P_{\mathcal{H}}(\mathbf{A})] + \operatorname{Tr}[P_{\mathcal{H}}(\mathbf{B})]. \end{aligned}$$

□

**Lemma B.7.** Let  $\mathbf{A} \succeq 0$  and  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ . It holds

$$\|P_{\mathcal{H}}(\mathbf{A} + \mathbf{x}\mathbf{x}^{\top}) - P_{\mathcal{H}}(\mathbf{A} + \mathbf{y}\mathbf{y}^{\top})\|_{\text{op}} \leq \sqrt{2} \|\mathbf{x} - \mathbf{y}\|_2.$$

*Proof.* Set  $F(\mathbf{X}, \mathbf{Y}) := \operatorname{tr}(\sqrt{\mathbf{X}}\sqrt{\mathbf{Y}})$  for  $\mathbf{X}, \mathbf{Y} \succeq 0$ . By Lieb's concavity theorem [Lieb \(1973\)](#),  $F$  is jointly concave and positively homogeneous, hence superadditive:

$$F(\mathbf{X}_1 + \mathbf{X}_2, \mathbf{Y}_1 + \mathbf{Y}_2) \geq F(\mathbf{X}_1, \mathbf{Y}_1) + F(\mathbf{X}_2, \mathbf{Y}_2).$$

Define  $\Pi_{\mathcal{H}}(\mathbf{H}) = P_{\mathcal{H}}(\mathbf{H})^2$ . Then  $\Pi_{\mathcal{H}}$  is the pinching onto a unital  $*$ -subalgebra  $\mathcal{H}$ .

For any  $\mathbf{A}, \mathbf{P}, \mathbf{Q} \succeq 0$ ,

$$\|P_{\mathcal{H}}(\mathbf{A} + \mathbf{P}) - P_{\mathcal{H}}(\mathbf{A} + \mathbf{Q})\|_{\text{F}}^2 = \operatorname{tr}(\mathbf{A} + \mathbf{P}) + \operatorname{tr}(\mathbf{A} + \mathbf{Q}) - 2 \operatorname{tr}(P_{\mathcal{H}}(\mathbf{A} + \mathbf{P})P_{\mathcal{H}}(\mathbf{A} + \mathbf{Q})).$$

Using  $P_{\mathcal{H}}(\mathbf{Z})^2 = \Pi_{\mathcal{H}}(\mathbf{Z})$ , the cross term equals  $F(\Pi_{\mathcal{H}}(\mathbf{A}) + \Pi_{\mathcal{H}}(\mathbf{P}), \Pi_{\mathcal{H}}(\mathbf{A}) + \Pi_{\mathcal{H}}(\mathbf{Q}))$ , which by superadditivity is at least  $F(\Pi_{\mathcal{H}}(\mathbf{A}), \Pi_{\mathcal{H}}(\mathbf{A})) + F(\Pi_{\mathcal{H}}(\mathbf{P}), \Pi_{\mathcal{H}}(\mathbf{Q})) = \operatorname{tr} \mathbf{A} + \operatorname{tr}(P_{\mathcal{H}}(\mathbf{P})P_{\mathcal{H}}(\mathbf{Q}))$ . Hence

$$\|P_{\mathcal{H}}(\mathbf{A} + \mathbf{P}) - P_{\mathcal{H}}(\mathbf{A} + \mathbf{Q})\|_{\text{F}} \leq \|P_{\mathcal{H}}(\mathbf{P}) - P_{\mathcal{H}}(\mathbf{Q})\|_{\text{F}}. \quad (9)$$

Write  $r = \|\mathbf{x}\|_2$ ,  $s = \|\mathbf{y}\|_2$ ,  $\mathbf{u} = \mathbf{x}/r$ ,  $\mathbf{v} = \mathbf{y}/s$ , and  $\cos \theta = \mathbf{u}^{\top} \mathbf{v}$ . For the unstructured square root,

$$\|\sqrt{\mathbf{x}\mathbf{x}^{\top}} - \sqrt{\mathbf{y}\mathbf{y}^{\top}}\|_{\text{F}}^2 = r^2 + s^2 - 2rs \cos^2 \theta \leq 2(r^2 + s^2 - 2rs \cos \theta) = 2\|\mathbf{x} - \mathbf{y}\|_2^2.$$

Pinching monotonicity of  $F$  yields

$$\|P_{\mathcal{H}}(\mathbf{x}\mathbf{x}^{\top}) - P_{\mathcal{H}}(\mathbf{y}\mathbf{y}^{\top})\|_{\text{F}}^2 = \|\sqrt{\Pi_{\mathcal{H}}(\mathbf{x}\mathbf{x}^{\top})} - \sqrt{\Pi_{\mathcal{H}}(\mathbf{y}\mathbf{y}^{\top})}\|_{\text{F}}^2 \leq \|\sqrt{\mathbf{x}\mathbf{x}^{\top}} - \sqrt{\mathbf{y}\mathbf{y}^{\top}}\|_{\text{F}}^2 \leq 2\|\mathbf{x} - \mathbf{y}\|_2^2.$$

Apply Eq. 9 with  $\mathbf{P} = \mathbf{x}\mathbf{x}^{\top}$ ,  $\mathbf{Q} = \mathbf{y}\mathbf{y}^{\top}$ ,

$$\|P_{\mathcal{H}}(\mathbf{A} + \mathbf{x}\mathbf{x}^{\top}) - P_{\mathcal{H}}(\mathbf{A} + \mathbf{y}\mathbf{y}^{\top})\|_{\text{op}} \leq \|P_{\mathcal{H}}(\mathbf{x}\mathbf{x}^{\top}) - P_{\mathcal{H}}(\mathbf{y}\mathbf{y}^{\top})\|_{\text{F}} \leq \sqrt{2} \|\mathbf{x} - \mathbf{y}\|_2.$$

□

**Lemma B.8.** For any well-structured preconditioner set  $\mathcal{H}$  and any vector  $\mathbf{x} \in \mathbb{R}^d$ , it holds that

$$\|\mathbf{x}\|_{\mathcal{H},*} = \operatorname{Tr}[P_{\mathcal{H}}(\mathbf{x}\mathbf{x}^{\top})].$$

*Proof.* We have that

$$\begin{aligned} \operatorname{Tr}[P_{\mathcal{H}}(\mathbf{x}\mathbf{x}^{\top})] &= \frac{1}{2} \min_{\mathbf{H} \in \mathcal{H}} \mathbf{x}^{\top} \mathbf{H}^{-1} \mathbf{x} + \operatorname{Tr}(\mathbf{H}) \\ &= \frac{1}{2} \min_{\mathbf{H} \in \mathcal{H}, \operatorname{Tr}(\mathbf{H}) \leq 1} \min_{c > 0} c^{-1} \mathbf{x}^{\top} \mathbf{H}^{-1} \mathbf{x} + c \operatorname{Tr}(\mathbf{H}) \\ &= \min_{\mathbf{H} \in \mathcal{H}, \operatorname{Tr}(\mathbf{H}) \leq 1} \sqrt{\mathbf{x}^{\top} \mathbf{H}^{-1} \mathbf{x}}. \end{aligned}$$

[Xie et al. \(2025b\)](#) shows that  $\|\mathbf{x}\|_{\mathcal{H},*} = \min_{\mathbf{H} \in \mathcal{H}, \operatorname{Tr}(\mathbf{H}) \leq 1} \sqrt{\mathbf{x}^{\top} \mathbf{H}^{-1} \mathbf{x}}$  in their proof of lemma 3.3, which finishes the proof.

□

**Lemma B.9.**

$$P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t = \arg \min_{\|\mathbf{x}\|_{\mathcal{H}} \leq 1} \langle \mathbf{x}, \mathbf{g}_t \rangle$$

*Proof.* First we show that it satisfies that  $\|P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t\|_{\mathcal{H}} \leq 1$ . For any  $\mathbf{H} \in \mathcal{H}$  and  $\text{Tr}(\mathbf{H}) \leq 1$ , we have that

$$\begin{aligned} \|P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2 &= \mathbf{g}_t^\top P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{H} P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t \\ &= \text{Tr}(P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{H} P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t \mathbf{g}_t^\top) \\ &= \text{Tr}(P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{H} P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^2) \\ &= \text{Tr}(\mathbf{H}) \leq 1 \end{aligned}$$

Then  $\|P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t\|_{\mathcal{H}} = \sup_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \|P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t\|_{\mathbf{H}} \leq 1$ .

Next, we can employ Lemma B.8 to show that

$$\langle P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t, \mathbf{g}_t \rangle = \text{Tr}(P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)^{-1} \mathbf{g}_t \mathbf{g}_t^\top) = \text{Tr}(P_{\mathcal{H}}(\mathbf{g}_t \mathbf{g}_t^\top)) = \|\mathbf{g}_t\|_{\mathcal{H},*}.$$

Because  $\min_{\|\mathbf{x}\|_{\mathcal{H}} \leq 1} \langle \mathbf{x}, \mathbf{g}_t \rangle = \|\mathbf{g}_t\|_{\mathcal{H},*}$  and there is only one unique vector achieving this optimality, we finish showing the statement.  $\square$

**Proposition B.10.** For any well-structured preconditioner set  $\mathcal{H}$ , let  $\sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}})$  be adaptive gradient variance defined in Definition 4.1. It always hold  $\sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}}) \leq \text{Tr}(P_{\mathcal{H}}(\boldsymbol{\Sigma}))$  when we assume  $\mathbb{E}[\nabla f(\mathbf{x}) - \nabla f_t(\mathbf{x})][\nabla f(\mathbf{x}) - \nabla f_t(\mathbf{x})]^\top \preceq \boldsymbol{\Sigma}$  for any  $t \in \mathcal{T}$ .

*Proof of Proposition B.10.* From Definition 4.1, we can have that

$$\begin{aligned} \sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}})^2 &= \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \sup_{\mathbf{x}, t} \mathbb{E} \|\nabla f(\mathbf{x}) - \nabla f_t(\mathbf{x})\|_{\mathbf{H}^{-1}}^2 \\ &= \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \sup_{\mathbf{x}, t} \mathbb{E} \langle [\nabla f(\mathbf{x}) - \nabla f_t(\mathbf{x})][\nabla f(\mathbf{x}) - \nabla f_t(\mathbf{x})]^\top, \mathbf{H}^{-1} \rangle \\ &\leq \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \langle \boldsymbol{\Sigma}, \mathbf{H}^{-1} \rangle. \end{aligned}$$

By the definition of  $P_{\mathcal{H}}(\boldsymbol{\Sigma})$ , we know that  $\langle \boldsymbol{\Sigma}, P_{\mathcal{H}}(\boldsymbol{\Sigma})^{-1} \rangle = \text{Tr}(P_{\mathcal{H}}(\boldsymbol{\Sigma}))$  and

$$\begin{aligned} \text{Tr}(P_{\mathcal{H}}(\boldsymbol{\Sigma})) &= \frac{1}{2} \min_{\mathbf{H} \in \mathcal{H}} \langle \boldsymbol{\Sigma}, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H}) \\ &\geq \min_{\mathbf{H} \in \mathcal{H}} \sqrt{\langle \boldsymbol{\Sigma}, \mathbf{H}^{-1} \rangle \text{Tr}(\mathbf{H})} \\ &= \min_{\mathbf{H} \in \mathcal{H}} \sqrt{\langle \boldsymbol{\Sigma}, (\mathbf{H} / \text{Tr}(\mathbf{H}))^{-1} \rangle} \\ &= \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \sqrt{\langle \boldsymbol{\Sigma}, \mathbf{H}^{-1} \rangle} \\ &\geq \sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}}), \end{aligned}$$

which finishes the proof.  $\square$

**B.1 COMPARISON ON ADAPTIVE GEOMETRY**

**Proposition 2.5.** For any well-structured preconditioner set  $\mathcal{H} \subseteq \mathcal{S}_+^d$  and any loss function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ , it always holds that  $L_{\|\cdot\|_{\mathcal{H}}}(f) \leq \Lambda_{\mathcal{H}}(f) \leq d \cdot L_{\|\cdot\|_{\mathcal{H}}}(f)$ .

*Proof of Proposition 2.5.* For any  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , by the definition of  $\|\cdot\|_{\mathcal{H}}$  and  $\|\cdot\|_{\mathcal{H},*}$ , it holds for any  $\mathbf{x} \in \mathbb{R}^d$  that  $\|\mathbf{x}\|_{\mathcal{H}} \geq \|\mathbf{x}\|_{\mathbf{H}}$  and  $\|\mathbf{x}\|_{\mathcal{H},*} \leq \|\mathbf{x}\|_{\mathbf{H},*}$ . Then by the definition of  $L_{\|\cdot\|_{\mathcal{H}}}(f)$  in Definition 2.4,

$$L_{\|\cdot\|_{\mathcal{H}}}(f) = \sup_{\mathbf{x}, \mathbf{y}} \frac{\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathcal{H},*}}{\|\mathbf{x} - \mathbf{y}\|_{\mathcal{H}}} \leq \sup_{\mathbf{x}, \mathbf{y}} \frac{\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathbf{H},*}}{\|\mathbf{x} - \mathbf{y}\|_{\mathbf{H}}} = L_{\|\cdot\|_{\mathbf{H}}}(f).$$

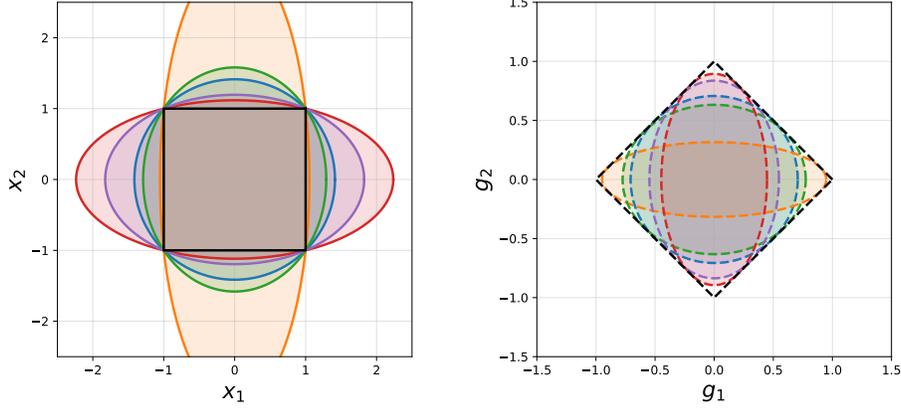


Figure 1: Here we demonstrate the duality between the supremum of the primal norms and the infimum of the corresponding dual norms for any well-structured preconditioner set  $\mathcal{H}$ . In particular, we consider  $\mathcal{H} = \{\text{all diagonal PSD matrices}\}$ , in which case  $\|\cdot\|_{\mathcal{H}} = \|\cdot\|_{\infty}$  and  $\|\cdot\|_{\mathcal{H},*} = \|\cdot\|_1$ . **Left figure:** the  $\|\cdot\|_{\infty}$ -unit ball (black square) in the primal space is the intersection of all  $\|\cdot\|_{\mathbf{H}}$ -unit ball (colored ellipses) for  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , that is,  $\|\cdot\|_{\infty}$  is the supremum of all such primal  $\|\cdot\|_{\mathbf{H}}$  norms. **Right figure:** the  $\|\cdot\|_1$ -unit ball (dashed black square) in the dual space is the union of all  $\|\cdot\|_{\mathbf{H},*}$ -unit balls (dashed ellipses) for  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , that is,  $\|\cdot\|_1$  is the infimum of all such dual  $\|\cdot\|_{\mathbf{H},*}$  norms.

Therefore, further minimizing both sides over  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , we obtain the desired result.

On the other hand, we can choose  $\mathbf{H}^* = \frac{1}{d}\mathbf{I}_d \in \mathcal{H}$  and Xie et al. (2025b) shows that  $\|\cdot\|_{\mathbf{H}^*} = \frac{1}{\sqrt{d}}\|\cdot\|_2$ . Then we have that  $\Lambda_{\mathcal{H}}(f) \leq L_{\|\cdot\|_{\mathbf{H}^*}}(f)$ .

We can also choose  $\mathcal{H}^*$  to be all PSD matrices so that we know  $\mathcal{H} \subseteq \mathcal{H}^*$ . Then we know  $\|\mathbf{x}\|_{\mathcal{H}} = \sup_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \|\mathbf{x}\|_{\mathbf{H}} \leq \sup_{\mathbf{H} \in \mathcal{H}^*, \text{Tr}(\mathbf{H}) \leq 1} \|\mathbf{x}\|_{\mathbf{H}} = \|\mathbf{x}\|_{\mathcal{H}^*}$  and  $\|\mathbf{x}\|_{\mathcal{H},*} = \inf_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \|\mathbf{x}\|_{\mathbf{H}^{-1}} \geq \inf_{\mathbf{H} \in \mathcal{H}^*, \text{Tr}(\mathbf{H}) \leq 1} \|\mathbf{x}\|_{\mathbf{H}^{-1}} = \|\mathbf{x}\|_{\mathcal{H}^*,*}$ . Xie et al. (2025b) also shows that  $\|\mathbf{x}\|_{\mathcal{H}^*} = \|\mathbf{x}\|_{\mathcal{H}^*,*} = \|\mathbf{x}\|_2$ . So we have that

$$\begin{aligned}
L_{\|\cdot\|_{\mathcal{H}}}(f) &= \sup_{\mathbf{x}, \mathbf{y}} \frac{\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathcal{H},*}}{\|\mathbf{x} - \mathbf{y}\|_{\mathcal{H}}} \\
&\geq \sup_{\mathbf{x}, \mathbf{y}} \frac{\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_2}{\|\mathbf{x} - \mathbf{y}\|_2} \\
&= \frac{1}{d} \sup_{\mathbf{x}, \mathbf{y}} \frac{\sqrt{d} \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_2}{\frac{1}{\sqrt{d}} \|\mathbf{x} - \mathbf{y}\|_2} \\
&= \frac{1}{d} \sup_{\mathbf{x}, \mathbf{y}} \frac{\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_{\mathbf{H}^* - 1}}{\|\mathbf{x} - \mathbf{y}\|_{\mathbf{H}^*}} \\
&\geq \frac{1}{d} \Lambda_{\mathcal{H}}(f).
\end{aligned}$$

□

**Proposition B.11.** For any set of stochastic loss functions  $\{f_t\}_{t \in \mathcal{T}}$  and any well-structured preconditioner set  $\mathcal{H}$ , it always holds  $\sigma_{\|\cdot\|_{\mathcal{H},*}}(\{f_t\}_{t \in \mathcal{T}})^2 \leq \sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}})^2 \leq d \cdot \sigma_{\|\cdot\|_{\mathcal{H},*}}(\{f_t\}_{t \in \mathcal{T}})^2$ .

*Proof.* For any  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , it always holds that  $\|\mathbf{u}\|_{\mathcal{H},*}^2 \leq \|\mathbf{u}\|_{\mathbf{H}^{-1}}^2$  for any  $\mathbf{u} \in \mathbb{R}^d$ , and thus

$$\sup_{t \in \mathcal{T}, \mathbf{x} \in \mathbb{R}^d} \mathbb{E}[\|\nabla f_t(\mathbf{x}) - \mathbb{E}[\nabla f_t(\mathbf{x})]\|_{\mathcal{H},*}^2] \leq \sup_{t \in \mathcal{T}, \mathbf{x} \in \mathbb{R}^d} \mathbb{E}[\|\nabla f_t(\mathbf{x}) - \mathbb{E}[\nabla f_t(\mathbf{x})]\|_{\mathbf{H}^{-1}}^2].$$

By minimizing over  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , we conclude that  $\sigma_{\|\cdot\|_{\mathcal{H},*}}(\{f_t\}_{t \in \mathcal{T}})^2 \leq \sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}})^2$ .

Similar with the proof of Proposition 2.5, we can show that  $\sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}})^2 \leq \sup_{\mathbf{x}, t \in [T]} \mathbb{E} d \|\nabla f(\mathbf{x}) - \nabla f_t(\mathbf{x})\|_2^2$  when choosing  $\mathbf{H}^* = \frac{1}{d} \mathbf{I}_d$ . We can also show that  $\sup_{t \in \mathcal{T}, \mathbf{x} \in \mathbb{R}^d} \mathbb{E} \|\nabla f_t(\mathbf{x}) - \mathbb{E}[\nabla f_t(\mathbf{x})]\|_{\mathcal{H},*}^2 \geq \sup_{t \in \mathcal{T}, \mathbf{x} \in \mathbb{R}^d} \mathbb{E} \|\nabla f_t(\mathbf{x}) - \mathbb{E}[\nabla f_t(\mathbf{x})]\|_2^2$ . Then combining them, we have that  $d \sup_{t \in \mathcal{T}, \mathbf{x} \in \mathbb{R}^d} \mathbb{E} \|\nabla f_t(\mathbf{x}) - \mathbb{E}[\nabla f_t(\mathbf{x})]\|_{\mathcal{H},*}^2 \geq \sigma_{\mathcal{H}}(\{f_t\}_{t \in \mathcal{T}})^2$ .  $\square$

## C PROOF FOR THE MATRIX INEQUALITY

**Lemma 3.3.** *Let  $\epsilon \geq 0$  and  $\beta \in (0, 1]$ . For any  $T \in \mathbb{N}$ , consider **any sequence of vectors**  $\mathbf{g}_0, \dots, \mathbf{g}_{T-1} \in \mathbb{R}^d$ . Let  $\mathbf{M}_{-1} = 0$ , and recursively define  $\mathbf{M}_t = \beta \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$  for  $t = 0, \dots, T-1$ . For any well-structured preconditioner set  $\mathcal{H}$ , define  $\mathbf{V}_t = \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$  for each  $t \in [T-1]$ . Then for any  $\mathbf{H} \in \mathcal{H} \cap \mathcal{S}_{++}^d$ , it holds*

$$\sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2 \leq \text{Tr}(\mathbf{H}) \|\mathbf{S}_T\|_{\text{op}} \quad \text{where } \mathbf{S}_T = \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1}.$$

Moreover, there exists an absolute constant  $C_1, C_2 > 0$ , independent of  $d, T, \epsilon, \beta$  and  $\mathcal{H}$ , such that

$$\|\mathbf{S}_T\|_{\text{op}} \leq C_1 \left( 1 + \log \left( 1 + \frac{d}{\epsilon} \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 + d^2(1-\beta)T \right) \right) \left( \frac{1-\beta}{\beta} T + \log \|\mathbf{V}_{T-1}^2 / \epsilon\|_{\text{op}} \right) + C_2.$$

In the special case when  $\mathcal{H}$  is commutative, the above bound can be further improved to

$$\|\mathbf{S}_T\|_{\text{op}} \leq (1-\beta)T + \log \|\mathbf{V}_{T-1}^2 / \epsilon\|_{\text{op}}.$$

*Proof of Lemma 3.3.* For each  $t = 0, \dots, T-1$ , by the definition of  $\{\mathbf{M}_t\}_{t=0}^{T-1}$ , we have

$$\|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2 = \langle \mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1}, \mathbf{g}_t \mathbf{g}_t^\top \rangle = \langle \mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1}, \mathbf{M}_t + \epsilon \mathbf{I}_d \rangle - \beta \langle \mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1}, \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d \rangle.$$

Since  $\mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1} \in \mathcal{H}$ , it follows from Lemma B.2 and the definition of  $\{\mathbf{V}_t\}_{t=0}^{T-1}$  that

$$\begin{aligned} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2 &= \langle \mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1}, P_{\mathcal{H}}(\mathbf{M}_t + \epsilon \mathbf{I}_d)^2 \rangle - \beta \langle \mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1}, P_{\mathcal{H}}(\mathbf{M}_{t-1} + \epsilon \mathbf{I}_d)^2 \rangle \\ &= \langle \mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1}, \mathbf{V}_t^2 \rangle - \beta \langle \mathbf{V}_t^{-1} \mathbf{H} \mathbf{V}_t^{-1}, \mathbf{V}_{t-1}^2 \rangle \\ &= \langle \mathbf{H}, \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \rangle \end{aligned}$$

Then summing over  $t = 0, \dots, T-1$  gives

$$\begin{aligned} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2 &= \sum_{t=0}^{T-1} \langle \mathbf{H}, \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \rangle \\ &\leq \text{Tr}(\mathbf{H}) \left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \right\|_{\text{op}}. \end{aligned}$$

This shows the first part of the lemma. We next bound the spectral norm of  $\sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1}$ .

Note that  $\beta \mathbf{V}_{t-1}^2 = \beta P_{\mathcal{H}}(\mathbf{M}_{t-1} + \epsilon \mathbf{I}_d)^2 = P_{\mathcal{H}}(\beta(\mathbf{M}_{t-1} + \epsilon \mathbf{I}_d))^2$  by Lemma B.3. Since  $\mathbf{M}_t + \epsilon \mathbf{I}_d \succeq \beta(\mathbf{M}_{t-1} + \epsilon \mathbf{I}_d)$ , we further have  $\mathbf{V}_t^2 \succeq \beta \mathbf{V}_{t-1}^2 \succ 0$  by Lemma B.3. Therefore, since  $\lambda_{\min}(\mathbf{V}_t^2) \geq \epsilon$ ,

we can apply Lemma C.1 to get that for any  $C \geq c > 0$ ,

$$\begin{aligned}
\mathbf{V}_t^{-1}(\mathbf{V}_t^2 - \beta\mathbf{V}_{t-1}^2)\mathbf{V}_t^{-1} &\preceq \frac{3(\log C - \log c)}{\pi^2}(\log(\mathbf{V}_t^2) - \log(\beta\mathbf{V}_{t-1}^2)) + \left(\frac{12cd}{\pi^2\epsilon^2}\right. \\
&\quad \left. + \frac{12C^{-1}d}{\pi^2}\right) \text{Tr}(\mathbf{V}_t^2 - \beta\mathbf{V}_{t-1}^2)\mathbf{I}_d. \\
&= \frac{3(\log C - \log c)}{\pi^2} \log \frac{1}{\beta} \cdot \mathbf{I}_d + \frac{3(\log C - \log c)}{\pi^2}(\log(\mathbf{V}_t^2) - \log(\mathbf{V}_{t-1}^2)) \\
&\quad + \left(\frac{12cd}{\pi^2\epsilon^2} + \frac{12C^{-1}d}{\pi^2}\right) \text{Tr}[\mathbf{M}_t + \epsilon\mathbf{I}_d - \beta(\mathbf{M}_{t-1} + \epsilon\mathbf{I}_d)]\mathbf{I}_d \\
&= \frac{3(\log C - \log c)}{\pi^2} \log \frac{1}{\beta} \cdot \mathbf{I}_d + \frac{3(\log C - \log c)}{\pi^2}(\log(\mathbf{V}_t^2) - \log(\mathbf{V}_{t-1}^2)) \\
&\quad + \left(\frac{12cd}{\pi^2\epsilon^2} + \frac{12C^{-1}d}{\pi^2}\right) \text{Tr}[\mathbf{g}_t\mathbf{g}_t^\top + (1-\beta)\epsilon\mathbf{I}_d]\mathbf{I}_d
\end{aligned}$$

where the first equality holds by Lemma B.2. Further summing over  $t = 0, \dots, T-1$  gives

$$\begin{aligned}
\sum_{t=0}^{T-1} \mathbf{V}_t^{-1}(\mathbf{V}_t^2 - \beta\mathbf{V}_{t-1}^2)\mathbf{V}_t^{-1} &\preceq \frac{3(\log C - \log c)}{\pi^2} T \log \frac{1}{\beta} \cdot \mathbf{I}_d + \frac{3(\log C - \log c)}{\pi^2}(\log(\mathbf{V}_{T-1}^2) - \log(\mathbf{V}_{-1}^2)) \\
&\quad + \left(\frac{12cd}{\pi^2\epsilon^2} + \frac{12C^{-1}d}{\pi^2}\right) \sum_{t=0}^{T-1} \text{Tr}[\mathbf{g}_t\mathbf{g}_t^\top + (1-\beta)\epsilon\mathbf{I}_d]\mathbf{I}_d \\
&= \frac{3(\log C - \log c)}{\pi^2} \left(T \log \frac{1}{\beta} \cdot \mathbf{I}_d + \log(\mathbf{V}_{T-1}^2/\epsilon)\right) \\
&\quad + \left(\frac{12cd}{\pi^2\epsilon^2} + \frac{12C^{-1}d}{\pi^2}\right) \left(d(1-\beta)\epsilon T + \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2\right) \mathbf{I}_d
\end{aligned}$$

where the equality holds as  $\mathbf{V}_{-1} = P_{\mathcal{H}}(\epsilon\mathbf{I}_d) = \sqrt{\epsilon}\mathbf{I}_d$ . Note that  $\log(1/\beta) = \log(1 + (1-\beta)/\beta) \leq (1-\beta)/\beta$  for any  $\beta \in (0, 1]$ . Then by triangle inequality for the spectral norm, we have

$$\begin{aligned}
\left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1}(\mathbf{V}_t^2 - \beta\mathbf{V}_{t-1}^2)\mathbf{V}_t^{-1} \right\|_{\text{op}} &\leq \frac{3(\log C - \log c)}{\pi^2} \left(\frac{1-\beta}{\beta} T + \log \|\mathbf{V}_{T-1}^2/\epsilon\|_{\text{op}}\right) \\
&\quad + \left(\frac{12cd}{\pi^2\epsilon^2} + \frac{12C^{-1}d}{\pi^2}\right) \left(d(1-\beta)T\epsilon + \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2\right). \tag{10}
\end{aligned}$$

In particular, we choose

$$c = \frac{(1-\beta)T/\beta + \log \|\mathbf{V}_{T-1}^2/\epsilon\|_{\text{op}}}{4d(d(1-\beta)T\epsilon + \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2)} \epsilon^2, \quad C = \max(\epsilon^2/c, c). \tag{11}$$

With this, the second term on the right-hand side of (10) satisfies that

$$\left(\frac{12cd}{\pi^2\epsilon^2} + \frac{12C^{-1}d}{\pi^2}\right) \left(d(1-\beta)T\epsilon + \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2\right) \leq \frac{6}{\pi^2} \left(\frac{1-\beta}{\beta} T + \log \|\mathbf{V}_{T-1}^2/\epsilon\|_{\text{op}}\right).$$

Plugging this back into (10), we obtain

$$\left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1}(\mathbf{V}_t^2 - \beta\mathbf{V}_{t-1}^2)\mathbf{V}_t^{-1} \right\|_{\text{op}} \leq \frac{6 + 3 \log(C/c)}{\pi^2} \left(\frac{1-\beta}{\beta} T + \log \|\mathbf{V}_{T-1}^2/\epsilon\|_{\text{op}}\right).$$

Moreover, by the choice of  $C, c$  in (11),

$$\log \frac{C}{c} = \log(\max(\epsilon^2/c^2, 1)) = 2 \log \left( \max \left( \frac{4d(d(1-\beta)T\epsilon + \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2)}{(1-\beta)T/\beta + \log \|\mathbf{V}_{T-1}^2/\epsilon\|_{\text{op}}}, 1 \right) \right).$$

For convenience, denote  $A = (1-\beta)T/\beta + \log \|\mathbf{V}_{T-1}^2/\epsilon\|_{\text{op}}$  and  $B = 4d(d(1-\beta)T\epsilon + \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2)$ . Then combining the above two displays, we get

$$\left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \right\|_{\text{op}} \leq \frac{6}{\pi^2} A + \frac{3}{\pi^2} A \log \left( \max \left( \frac{B}{A}, 1 \right) \right).$$

When  $A > 1$ , we always have  $\max(B/A, 1) \leq 1 + B$ , which gives rise to the following bound

$$\left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \right\|_{\text{op}} \leq \frac{6}{\pi^2} (1 + \log(1 + B)) A.$$

When  $A < 1$ , using the fact that  $x \log(1/x) \leq 1/e$ , we have

$$A \log \left( \max \left( \frac{B}{A}, 1 \right) \right) \leq A \log(1 + B) - A \log A \leq A \log(1 + B) + 1/e.$$

In this case, the bound becomes

$$\left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \right\|_{\text{op}} \leq \frac{6}{\pi^2} \left[ (1 + \log(1 + B)) A + 1/e \right].$$

Combining the above two case, we can conclude that there exists universal constants  $C_1, C_2 > 0$  such that

$$\left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \right\|_{\text{op}} \leq C_1 (1 + \log(1 + B)) A + C_2.$$

This completes the proof for general well-structured preconditioner set  $\mathcal{H}$ .

In this special case where matrices in  $\mathcal{H}$  are diagonal, we have that

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} &= (1 - \beta) T \mathbf{I}_d + \beta \sum_{t=0}^{T-1} (\mathbf{I}_d - \mathbf{V}_t^{-1} \mathbf{V}_{t-1}^2 \mathbf{V}_t^{-1}) \\ &\preceq (1 - \beta) T \mathbf{I}_d + \beta \sum_{t=0}^{T-1} \log(\mathbf{V}_t \mathbf{V}_{t-1}^{-2} \mathbf{V}_t) \end{aligned}$$

where the inequality follows from the fact that  $\mathbf{I}_d - \mathbf{A}^{-1} \preceq \log \mathbf{A}$  for any  $\mathbf{A} \succeq 0$ . Further note that as  $\mathbf{V}_{t-1}, \mathbf{V}_t$  are diagonal matrices, they commute with each other. Then using the fact that  $\log(\mathbf{A}\mathbf{B}) = \log \mathbf{A} + \log \mathbf{B}$  when  $\mathbf{A}$  and  $\mathbf{B}$  commute, we obtain

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} &\preceq (1 - \beta) T \mathbf{I}_d + \beta \sum_{t=0}^{T-1} [\log(\mathbf{V}_t^2) - \log(\mathbf{V}_{t-1}^2)] \\ &= (1 - \beta) T \mathbf{I}_d + \log(\mathbf{V}_{T-1}^2/\epsilon). \end{aligned}$$

Therefore, in this case it holds that

$$\left\| \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1} \right\|_{\text{op}} \leq (1 - \beta) T + \log \|\mathbf{V}_{T-1}^2/\epsilon\|_{\text{op}}.$$

This completes the proof.  $\square$

**Lemma C.1.** For any positive definite matrices  $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{d \times d}$  such that  $\mathbf{Y} \preceq \mathbf{X}$ , it holds for any  $0 \leq c \leq C$  that

$$\mathbf{X}^{-1/2} (\mathbf{X} - \mathbf{Y}) \mathbf{X}^{-1/2} \preceq \frac{3(\log C - \log c)}{\pi^2} (\log \mathbf{X} - \log \mathbf{Y}) + \left( \frac{12cd}{\pi^2 \lambda_{\min}(\mathbf{X})^2} + \frac{12C^{-1}d}{\pi^2} \right) \text{Tr}(\mathbf{X} - \mathbf{Y}) \cdot \mathbf{I}_d.$$

*Proof of Lemma C.1.* For any  $\delta \in (0, 1)$ , since the matrix logarithm is operator concave, it holds that

$$\log((1 - \delta)\mathbf{X} + \delta\mathbf{Y}) \succeq (1 - \delta)\log \mathbf{X} + \delta \log \mathbf{Y}.$$

Reorganizing the above inequality yields that for all  $\delta \in (0, 1)$ ,

$$\log \mathbf{X} - \log \mathbf{Y} \succeq -\frac{\log(\mathbf{X} + \delta(\mathbf{Y} - \mathbf{X})) - \log \mathbf{X}}{\delta}.$$

Taking the limit  $\delta \rightarrow 0$ , we obtain

$$\log \mathbf{X} - \log \mathbf{Y} \succeq \partial \log(\mathbf{X})[\mathbf{X} - \mathbf{Y}].$$

The proof is completed by further applying Lemma C.2 with  $\mathbf{A} = \mathbf{X} - \mathbf{Y}$ .  $\square$

**Lemma C.2.** For any positive definite matrix  $\mathbf{X} \in \mathbb{R}^{d \times d}$  and any positive semi-definite matrix  $\mathbf{A} \in \mathbb{R}^{d \times d}$ , it holds for any  $0 \leq c \leq C$  that

$$\mathbf{X}^{-1/2} \mathbf{A} \mathbf{X}^{-1/2} \preceq \frac{3(\log C - \log c)}{\pi^2} \partial \log(\mathbf{X})[\mathbf{A}] + \left( \frac{12cd}{\pi^2 \lambda_{\min}(\mathbf{X})^2} + \frac{12C^{-1}d}{\pi^2} \right) \text{Tr}(\mathbf{A}) \cdot \mathbf{I}_d. \quad (12)$$

*Proof of Lemma C.2.* We consider the following expansion of the matrix logarithm:

$$\log(\mathbf{X} + \delta \mathbf{A}) = \log \mathbf{X} + \int_0^\infty (\mathbf{X} + z\mathbf{I})^{-1} \cdot \delta \mathbf{A} \cdot (\mathbf{X} + z\mathbf{I})^{-1} dz + O(\delta^2).$$

This implies that

$$\partial \log(\mathbf{X})[\mathbf{A}] = \lim_{\delta \rightarrow 0} \frac{\log(\mathbf{X} + \delta \mathbf{A}) - \log \mathbf{X}}{\delta} = \int_0^\infty \frac{\mathbf{I}}{\mathbf{X} + z\mathbf{I}} \cdot \mathbf{A} \cdot \frac{\mathbf{I}}{\mathbf{X} + z\mathbf{I}} dz.$$

Let  $\mathbf{v} \in \mathbb{R}^d$  be an arbitrary unit vector. Note that both  $\mathbf{X}^{-1/2} \mathbf{A} \mathbf{X}^{-1/2}$  and  $\partial \log(\mathbf{X})[\mathbf{A}]$  are linear in  $\mathbf{A}$ , and thus we first consider  $\mathbf{A} = \mathbf{u}\mathbf{u}^\top$  for any unit vector  $\mathbf{u} \in \mathbb{R}^d$ . We define the following two quantities

$$f(\mathbf{u}, \mathbf{v}) = \mathbf{v}^\top \partial \log(\mathbf{X})[\mathbf{u}\mathbf{u}^\top] \mathbf{v} = \mathbf{v}^\top \int_0^\infty (\mathbf{X} + z\mathbf{I}_d)^{-1} \mathbf{u}\mathbf{u}^\top (\mathbf{X} + z\mathbf{I}_d)^{-1} dz \mathbf{v}, \quad (13)$$

$$g(\mathbf{u}, \mathbf{v}) = \mathbf{v}^\top \mathbf{X}^{-1/2} \mathbf{u}\mathbf{u}^\top \mathbf{X}^{-1/2} \mathbf{v}. \quad (14)$$

Now suppose that the SVD of  $\mathbf{X}$  is  $\mathbf{X} = \mathbf{U} \text{diag}(\lambda_1, \dots, \lambda_d) \mathbf{U}^\top$  where  $\mathbf{U} \in \mathbb{R}^{d \times d}$  is an orthogonal matrix. Then  $\mathbf{X}^{-1/2} = \mathbf{U} \text{diag}(\lambda_1^{-1/2}, \dots, \lambda_d^{-1/2}) \mathbf{U}^\top$  and  $(\mathbf{X} + z\mathbf{I})^{-1} = \mathbf{U} \text{diag}(1/(\lambda_1 + z), \dots, 1/(\lambda_d + z)) \mathbf{U}^\top$ . Writing  $\tilde{\mathbf{u}} = \mathbf{U}\mathbf{u}$  and  $\tilde{\mathbf{v}} = \mathbf{U}\mathbf{v}$ , then  $f(\mathbf{u}, \mathbf{v})$  becomes

$$\begin{aligned} f(\mathbf{u}, \mathbf{v}) &= \int_0^\infty \tilde{\mathbf{v}}^\top \text{diag}\left(\frac{1}{\lambda_1 + z}, \dots, \frac{1}{\lambda_d + z}\right) \tilde{\mathbf{u}} \tilde{\mathbf{u}}^\top \text{diag}\left(\frac{1}{\lambda_1 + z}, \dots, \frac{1}{\lambda_d + z}\right) \tilde{\mathbf{v}} dz \\ &= \int_0^\infty \langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (1/(\lambda_1 + z), \dots, 1/(\lambda_d + z)) \rangle^2 dz. \end{aligned}$$

Similarly,  $g(\mathbf{u}, \mathbf{v})$  becomes

$$\begin{aligned} g(\mathbf{u}, \mathbf{v}) &= \tilde{\mathbf{v}}^\top \text{diag}(\lambda_1^{-1/2}, \dots, \lambda_d^{-1/2}) \tilde{\mathbf{u}} \tilde{\mathbf{u}}^\top \text{diag}(\lambda_1^{-1/2}, \dots, \lambda_d^{-1/2}) \tilde{\mathbf{v}} \\ &= \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (1/\lambda_1^{-1/2}, \dots, 1/\lambda_d^{-1/2}) \right\rangle^2. \end{aligned}$$

Now applying the fact that  $\lambda^{-1/2} = \frac{1}{\pi} \int_0^\infty \frac{z^{-1/2}}{\lambda + z} dz$  for any  $\lambda > 0$ , we have

$$g(\mathbf{u}, \mathbf{v}) = \left( \frac{1}{\pi} \int_0^\infty \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (z^{-1/2}/(\lambda_1 + z), \dots, z^{-1/2}/(\lambda_d + z)) \right\rangle dz \right)^2.$$

To further bound  $g(\mathbf{u}, \mathbf{v})$ , we split the integral into three parts:  $[0, c]$ ,  $[c, C]$ , and  $[C, \infty)$  where  $c > 0$  and  $C > 0$  are constants to be determined later. That is,

$$\begin{aligned} g(\mathbf{u}, \mathbf{v}) &\leq \frac{3}{\pi^2} \left( \int_0^c \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (z^{-1/2}/(\lambda_1 + z), \dots, z^{-1/2}/(\lambda_d + z)) \right\rangle dz \right)^2 \\ &\quad + \frac{3}{\pi^2} \left( \int_c^C \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (z^{-1/2}/(\lambda_1 + z), \dots, z^{-1/2}/(\lambda_d + z)) \right\rangle dz \right)^2 \\ &\quad + \frac{3}{\pi^2} \left( \int_C^\infty \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (z^{-1/2}/(\lambda_1 + z), \dots, z^{-1/2}/(\lambda_d + z)) \right\rangle dz \right)^2 \\ &=: g_1(\mathbf{u}, \mathbf{v}) + g_2(\mathbf{u}, \mathbf{v}) + g_3(\mathbf{u}, \mathbf{v}) \end{aligned}$$

where we apply the triangle inequality and denote the three terms on the right-hand side by  $g_1(\mathbf{u}, \mathbf{v})$ ,  $g_2(\mathbf{u}, \mathbf{v})$ ,  $g_3(\mathbf{u}, \mathbf{v})$  respectively. We control each term separately.

First, for  $g_1(\mathbf{u}, \mathbf{v})$ ,

$$\begin{aligned} g_1(\mathbf{u}, \mathbf{v}) &= \frac{3}{\pi^2} \left( \int_0^c \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (1/(\lambda_1 + z), \dots, 1/(\lambda_d + z)) \right\rangle z^{-1/2} dz \right)^2 \\ &\leq \frac{3}{\pi^2} \left( \int_0^c \|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2 \left( \sum_{i=1}^d \frac{1}{(\lambda_i + z)^2} \right)^{1/2} z^{-1/2} dz \right)^2 \\ &\leq \frac{3}{\pi^2} \left( \frac{\|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2 \sqrt{d}}{\min_{i \in [d]} \lambda_i} \int_0^c z^{-1/2} dz \right)^2 \\ &= \frac{12 \|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2^2 \cdot cd}{\pi^2 \cdot \min_{i \in [d]} \lambda_i^2} \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz inequality and in the second inequality we apply  $\frac{1}{\lambda_i + z} \leq \frac{1}{\lambda_i}$  as each  $\lambda_i$  is positive. Next, for the integral over  $[c, C]$ , we have

$$\begin{aligned} g_2(\mathbf{u}, \mathbf{v}) &= \frac{3}{\pi^2} \left( \int_c^C \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (1/(\lambda_1 + z), \dots, 1/(\lambda_d + z)) \right\rangle z^{-1/2} dz \right)^2 \\ &\leq \frac{3}{\pi^2} \left( \int_c^C \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (1/(\lambda_1 + z), \dots, 1/(\lambda_d + z)) \right\rangle^2 dz \right) \left( \int_c^C z^{-1} dz \right) \\ &\leq \frac{3(\log C - \log c)}{\pi^2} \left( \int_0^\infty \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (1/(\lambda_1 + z), \dots, 1/(\lambda_d + z)) \right\rangle^2 dz \right) \\ &= \frac{3(\log C - \log c)}{\pi^2} f(\mathbf{u}, \mathbf{v}) \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz inequality, and in the second inequality we relax the domain of the integral to  $[0, \infty)$ , which exactly gives us  $f(\mathbf{u}, \mathbf{v})$ . Finally, for  $g_3(\mathbf{u}, \mathbf{v})$ ,

$$\begin{aligned} g_3(\mathbf{u}, \mathbf{v}) &= \frac{3}{\pi^2} \left( \int_C^\infty \left\langle \tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}, (1/(\lambda_1 + z), \dots, 1/(\lambda_d + z)) \right\rangle z^{-1/2} dz \right)^2 \\ &\leq \frac{3}{\pi^2} \left( \int_C^\infty \|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2 \left( \sum_{i=1}^d \frac{1}{(\lambda_i + z)^2} \right)^{1/2} z^{-1/2} dz \right)^2 \\ &\leq \frac{3}{\pi^2} \left( \int_C^\infty \|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2 \sqrt{d} z^{-3/2} dz \right)^2 \\ &= \frac{12 \|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2^2 \cdot C^{-1} d}{\pi^2} \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz inequality and in the second inequality we apply  $\frac{1}{\lambda_i + z} \leq \frac{1}{z}$  as each  $\lambda_i$  is positive. Collecting the above bounds, we obtain

$$g(\mathbf{u}, \mathbf{v}) \leq \frac{3(\log C - \log c)}{\pi^2} f(\mathbf{u}, \mathbf{v}) + \frac{12 \|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2^2 \cdot cd}{\pi^2 \cdot \min_{i \in [d]} \lambda_i^2} + \frac{12 \|\tilde{\mathbf{u}} \odot \tilde{\mathbf{v}}\|_2^2 \cdot C^{-1} d}{\pi^2}. \quad (15)$$

Now for any general positive semi-definite matrix  $\mathbf{A} \in \mathbb{R}^{d \times d}$  with eigendecomposition  $\mathbf{A} = \sum_{i=1}^d \alpha_i \mathbf{u}_i \mathbf{u}_i^\top$ , we apply the bound (15) to each  $\mathbf{u}_i$  in the eigendecomposition and sum over all  $i$  to get

$$\begin{aligned} \sum_{i=1}^d \alpha_i g(\mathbf{u}_i, \mathbf{v}) &\leq \frac{3(\log C - \log c)}{\pi^2} \sum_{i=1}^d \alpha_i f(\mathbf{u}_i, \mathbf{v}) + \frac{12cd}{\pi^2 \cdot \min_{i \in [d]} \lambda_i^2} \sum_{i=1}^d \alpha_i \|\tilde{\mathbf{u}}_i \odot \tilde{\mathbf{v}}\|_2^2 \\ &\quad + \frac{12C^{-1}d}{\pi^2} \sum_{i=1}^d \alpha_i \|\tilde{\mathbf{u}}_i \odot \tilde{\mathbf{v}}\|_2^2. \end{aligned}$$

Note that  $\|\tilde{\mathbf{u}}_i \odot \tilde{\mathbf{v}}\|_2^2 \leq \|\tilde{\mathbf{u}}\|_2^2 \|\tilde{\mathbf{v}}\|_2^2 = 1$  as both  $\tilde{\mathbf{u}}$  and  $\tilde{\mathbf{v}}$  are unit vectors. Therefore, we further have

$$\begin{aligned} \sum_{i=1}^d \alpha_i g(\mathbf{u}_i, \mathbf{v}) &\leq \frac{3(\log C - \log c)}{\pi^2} \sum_{i=1}^d \alpha_i f(\mathbf{u}_i, \mathbf{v}) + \frac{12cd}{\pi^2 \cdot \min_{i \in [d]} \lambda_i^2} \sum_{i=1}^d \alpha_i + \frac{12C^{-1}d}{\pi^2} \sum_{i=1}^d \alpha_i \\ &= \frac{3(\log C - \log c)}{\pi^2} \sum_{i=1}^d \alpha_i f(\mathbf{u}_i, \mathbf{v}) + \frac{12cd}{\pi^2 \cdot \min_{i \in [d]} \lambda_i^2} \text{Tr}(\mathbf{A}) + \frac{12C^{-1}d}{\pi^2} \text{Tr}(\mathbf{A}). \end{aligned}$$

Recall the definition of  $f(\mathbf{u}, \mathbf{v})$  and  $g(\mathbf{u}, \mathbf{v})$  in (13) and (14). Then we have shown that for any unit vector  $\mathbf{v} \in \mathbb{R}^d$ ,

$$\mathbf{v}^\top \mathbf{X}^{-1/2} \mathbf{A} \mathbf{X}^{-1/2} \mathbf{v} \leq \frac{3(\log C - \log c)}{\pi^2} \mathbf{v}^\top \partial \log(\mathbf{X})[\mathbf{A}] \mathbf{v} + \frac{12cd}{\pi^2 \cdot \min_{i \in [d]} \lambda_i^2} \text{Tr}(\mathbf{A}) + \frac{12C^{-1}d}{\pi^2} \text{Tr}(\mathbf{A}).$$

Therefore, we conclude that

$$\mathbf{X}^{-1/2} \mathbf{A} \mathbf{X}^{-1/2} \preceq \frac{3(\log C - \log c)}{\pi^2} \partial \log(\mathbf{X})[\mathbf{A}] + \left( \frac{12cd}{\pi^2 \lambda_{\min}(\mathbf{X})^2} + \frac{12C^{-1}d}{\pi^2} \right) \text{Tr}(\mathbf{A}) \cdot \mathbf{I}_d.$$

□

## D UNIFIED PROOF FOR ADAPTIVE ALGORITHMS

### D.1 RELATIONSHIP BETWEEN ADAPTIVE ALGORITHMS

- *Cumulative accumulation* (Algorithm 4) maintains the plain sum of past squared gradients, thereby giving equal weight to the entire gradient history. A famous example in this category is AdaGrad.
- *EMA accumulation* (Algorithm 5) computes an exponential moving average of past gradients, which yields a stationary estimate of recent gradient statistics. Notable examples include Adam and RMSProp.
- *Weighted accumulation* (Algorithm 6) applies geometrically decaying weights to past gradients, which differs from EMA accumulation in that it does not normalize the weights to sum to one.

---

#### Algorithm 4 General Adaptive Cumulative Optimization Algorithm

---

**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , convex cone  $\mathcal{H} \subset \mathcal{S}_+$ ,  $\beta$

**Input:** initialization  $\mathbf{x}_0$ , stochastic loss functions  $\{f_t\}_{t=1}^T : \mathbb{R}^{d_L \times d_R} \rightarrow \mathbb{R}$

$\mathbf{M}_0 \leftarrow \mathbf{0}$

**for**  $t = 1, 2, \dots, T$ :

$\mathbf{g}_t \leftarrow \nabla f_t(\mathbf{x}_{t-1})$

$\mathbf{M}_t \leftarrow \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$

$\mathbf{V}_t \leftarrow \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$

$\mathbf{x}_t \leftarrow \mathbf{x}_{t-1} - \eta \mathbf{V}_t^{-1} \mathbf{g}_t$

**return**  $\mathbf{x}_T$

---

**Algorithm 5** General Adaptive EMA Optimization Algorithm

---

**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , convex cone  $\mathcal{H} \subset \mathcal{S}_+$ ,  $\beta$   
**Input:** initialization  $\mathbf{x}_0$ , stochastic loss functions  $\{f_t\}_{t=1}^T : \mathbb{R}^{d_L \times d_R} \rightarrow \mathbb{R}$   
 $\mathbf{M}_0 \leftarrow \mathbf{0}$   
**for**  $t = 1, 2, \dots, T$  :  
 $\mathbf{g}_t \leftarrow \nabla f_t(\mathbf{x}_{t-1})$   
 $\mathbf{M}_t \leftarrow \beta \mathbf{M}_{t-1} + (1 - \beta) \mathbf{g}_t \mathbf{g}_t^\top$   
 $\mathbf{V}_t \leftarrow \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$   
 $\mathbf{x}_t \leftarrow \mathbf{x}_{t-1} - \eta \mathbf{V}_t^{-1} \mathbf{g}_t$   
**return**  $\mathbf{x}_T$

---

**Algorithm 6** General Adaptive Weighted Optimization Algorithm

---

**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , convex cone  $\mathcal{H} \subset \mathcal{S}_+$ ,  $\beta$   
**Input:** initialization  $\mathbf{x}_0$ , stochastic loss functions  $\{f_t\}_{t=1}^T : \mathbb{R}^{d_L \times d_R} \rightarrow \mathbb{R}$   
 $\mathbf{M}_0 \leftarrow \mathbf{0}$   
**for**  $t = 1, 2, \dots, T$  :  
 $\mathbf{g}_t \leftarrow \nabla f_t(\mathbf{x}_{t-1})$   
 $\mathbf{M}_t \leftarrow \beta \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$   
 $\mathbf{V}_t \leftarrow \arg \min_{\mathbf{H} \in \mathcal{H}} \langle \mathbf{M}_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$   
 $\mathbf{x}_t \leftarrow \mathbf{x}_{t-1} - \eta \mathbf{V}_t^{-1} \mathbf{g}_t$   
**return**  $\mathbf{x}_T$

---

We also put the definition of one-sided Shampoo (Xie et al., 2025b) in Algorithm 7.

**Algorithm 7** One-sided Shampoo

---

**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , initial  $\mathbf{M}_0, \mathbf{L}_0 = \mathbf{0}$   
**Input:** initialization  $\mathbf{x}_0$ , stochastic loss functions  $\{f_t\}_{t=1}^T : \mathbb{R}^{d_L \times d_R} \rightarrow \mathbb{R}$   
**for**  $t = 1, 2, \dots, T$  :  
 $\mathbf{G}_t \leftarrow \nabla f_t(\mathbf{X}_{t-1})$   
 $\mathbf{L}_t \leftarrow \mathbf{L}_{t-1} + \mathbf{G}_t \mathbf{G}_t^\top$   
 $\mathbf{X}_t \leftarrow \mathbf{X}_{t-1} - \eta (\mathbf{L}_t + \epsilon \mathbf{I}_{d_L})^{-\frac{1}{2}} \mathbf{G}_t$   
**return**  $\mathbf{x}_T$

---

## D.2 PROOF FOR WEIGHTED ALGORITHM

Before presenting the main theorems, we first clarify our gradient noise assumption below. Assumption D.1 is stronger than the conventional noise assumption when it assumes the condition holds almost surely. It is required in proving Lemma D.3. Lemma D.5 can still hold when we only assume the covariance matrix is bounded.

**Assumption D.1.** For any  $t \in [T]$  and any  $\mathbf{x} \in \mathbb{R}^d$ ,  $\mathbb{E}[\nabla f_t(\mathbf{x})] = \nabla f(\mathbf{x})$  where the expectation is taken with respect to the randomness in the loss  $f_t$ . Moreover, there exists  $\Sigma \succeq 0$  such that for all  $t \geq 0$ ,  $-\Sigma \preceq \nabla f(\mathbf{x}) \nabla f(\mathbf{x})^\top - \nabla f_t(\mathbf{x}) \nabla f_t(\mathbf{x})^\top \preceq \Sigma$ .

With Assumption D.1 in place, Theorem D.2 presents the general result for weighted adaptive algorithms on stochastic nonconvex functions, which is then specialized to the cumulative and EMA variants in Theorem D.7 and D.8, respectively. When plugging  $\Sigma = \mathbf{0}$  in the stochastic rate, we can get the deterministic result in Section 3.2.

In this section, we will define  $\mathbf{H}^* = \arg \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} L_{\|\cdot\|_{\mathbf{H}}}(f)$ . For all the  $\mathbf{M}_t$  and  $\mathbf{V}_t$  in this section except the proof of Theorem D.7 and Theorem D.8, they are defined as in Algorithm 6, i.e.,  $\mathbf{M}_t = \beta \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$ . We will define  $\bar{\mathbf{g}}_t = \nabla f(\mathbf{x}_t)$  and

$$\tilde{\mathbf{M}}_t = \beta \mathbf{M}_{t-1} + \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top + \Sigma, \quad \tilde{\mathbf{V}}_t = P_{\mathcal{H}}(\tilde{\mathbf{M}}_t + \epsilon \mathbf{I}_d). \quad (16)$$

Then it always holds  $\mathbf{M}_t \preceq \tilde{\mathbf{M}}_t$  because of Assumption D.1. By Lemma B.3, it also holds  $\mathbf{V}_t \preceq \tilde{\mathbf{V}}_t$ .

**Theorem D.2.** For any  $\epsilon \geq 0$ ,  $\beta \in (0, 1]$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 1 with well-structured preconditioner set  $\mathcal{H}$ , where the update of  $\mathbf{M}_t$  follows the weighted version, i.e.,  $\mathbf{M}_t = \beta \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$  for all  $t \in [T]$ . Let  $\Lambda_{\mathcal{H}}(f)$  be the adaptive smoothness of the loss  $f$  according to Definition 2.4. Then under Assumption D.1, it holds that

$$\mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \right] \leq \frac{\sqrt{\sum_{i=0}^{T-1} \beta^{i/2}}}{T} \xi + \frac{\sqrt{\text{Tr} [P_{\mathcal{H}}((\sum_{i=0}^{T-1} \beta^i) \boldsymbol{\Sigma} + \epsilon \mathbf{I}_d)]}}{\sqrt{T}} \sqrt{\xi}.$$

where  $\xi$  is given by

$$\xi = \frac{2\Delta_0}{\eta} + \eta \Lambda_{\mathcal{H}}(f) \|\mathbf{S}_T\|_{\text{op}} + \sqrt{2d} \|\boldsymbol{\Sigma}\|_{\text{op}}^{1/2} \|\mathbf{S}_T\|_{\text{op}} \quad (17)$$

and  $\mathbf{S}_T = \mathbb{E} \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1}$ .

For general well-structured preconditioner set,  $\|\mathbf{S}_T\|_{\text{op}} = \tilde{O}(\log(d)[(1-\beta)T/\beta + \log(d)])$ . When the preconditioner set only has diagonal matrices,  $\|\mathbf{S}_T\|_{\text{op}} = (1-\beta)T + \tilde{O}(1)$ .

*Proof of Theorem D.2.* With the shorthand  $\bar{\mathbf{g}}_t = \nabla f(\mathbf{x}_t)$ , we need to bound

$$\mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \right] = \mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \|\bar{\mathbf{g}}_t\|_{\mathcal{H},*} \right] = \mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \text{Tr}[P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)] \right] \quad (18)$$

where the second equality holds by Lemma B.8. By employing Lemma D.4 and Lemma D.5, we have that

$$\begin{aligned} & \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)] \right)^2 \\ & \leq \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] \right) \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t] \right) \\ & \leq \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] \right) \left( T \cdot \text{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{T-1} \beta^i \right) \boldsymbol{\Sigma} + \epsilon \mathbf{I}_d \right) \right] + \left( \sum_{i=0}^{T-1} \beta^{i/2} \right) \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top \tilde{\mathbf{V}}_t^{-1}] \right). \end{aligned} \quad (19)$$

It then suffices to bound the sum of  $\text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top]$ . To this end, we apply Lemma D.3 to get

$$\mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] \leq 2 \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] + \sqrt{2 \|\boldsymbol{\Sigma}\|_{\text{op}}} \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \bar{\mathbf{g}}_t^\top]. \quad (20)$$

Here, the second term on the right-hand side can be controlled by applying Lemma 3.3. For the first term on the right-hand side, we need to apply the descent lemma. Specifically, recall that  $\mathbf{H}^* = \arg \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} L_{\|\cdot\|_{\mathbf{H}}}(f)$ , and then by second-order Taylor expansion,

$$\begin{aligned} f(\mathbf{x}_{t+1}) & \leq f(\mathbf{x}_t) + \langle \nabla f(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \frac{\Lambda_{\mathcal{H}}(f)}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_{\mathbf{H}^*}^2 \\ & = f(\mathbf{x}_t) - \eta \langle \bar{\mathbf{g}}_t, \mathbf{V}_t^{-1} \mathbf{g}_t \rangle + \frac{\eta^2 \Lambda_{\mathcal{H}}(f)}{2} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2. \end{aligned}$$

By taking expectation on both sides and summing over  $t = 0, 1, \dots, T-1$ , we get

$$\mathbb{E}[f(\mathbf{x}_T) - f(\mathbf{x}_0)] \leq -\eta \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] + \frac{\eta^2 \Lambda_{\mathcal{H}}(f)}{2} \mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2.$$

Rearranging the above inequality, we have

$$\begin{aligned} \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] &\leq \frac{\mathbb{E}[f(\mathbf{x}_0) - f(\mathbf{x}_T)]}{\eta} + \frac{\eta \Lambda_{\mathcal{H}}(f)}{2} \mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 \\ &\leq \frac{\Delta_0}{\eta} + \frac{\eta \Lambda_{\mathcal{H}}(f)}{2} \mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 \end{aligned} \quad (22)$$

where the second inequality holds by the definition  $\Delta_0 = f(\mathbf{x}_0) - \min_{\mathbf{x}} f(\mathbf{x})$ . Therefore, combining (21) and (22) yields

$$\mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \tilde{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] \leq \frac{2\Delta_0}{\eta} + \eta \Lambda_{\mathcal{H}}(f) \mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 + \sqrt{2\|\boldsymbol{\Sigma}\|_{\text{op}}} \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \mathbf{g}_t^\top]. \quad (23)$$

Now, we apply Lemma 3.3 to both  $\mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2$  and  $\mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \mathbf{g}_t^\top] = \mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_2^2$ , which gives

$$\begin{aligned} \mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 &\leq \text{Tr}[\mathbf{H}^*] \cdot \|\mathbf{S}_T\|_{\text{op}} = \|\mathbf{S}_T\|_{\text{op}}, \\ \mathbb{E} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_2^2 &\leq d \|\mathbf{S}_T\|_{\text{op}} \end{aligned}$$

where  $\mathbf{S}_T = \mathbb{E} \sum_{t=0}^{T-1} \mathbf{V}_t^{-1} (\mathbf{V}_t^2 - \beta \mathbf{V}_{t-1}^2) \mathbf{V}_t^{-1}$ . Based on this, we further define

$$\xi = \frac{2}{\eta} [f(\mathbf{x}_0) - \min f(\mathbf{x})] + \eta \Lambda_{\mathcal{H}}(f) \|\mathbf{S}_T\|_{\text{op}} + \sqrt{2\|\boldsymbol{\Sigma}\|_{\text{op}}} d \|\mathbf{S}_T\|_{\text{op}}. \quad (24)$$

Applying these bounds to (23) gives

$$\mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \tilde{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] \leq \xi. \quad (25)$$

Now we can plug (25) into (20), and hence it follows from (18) that

$$\begin{aligned} \mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \right] &= \frac{1}{T} \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)] \\ &\leq \frac{\sqrt{\sum_{i=0}^{T-1} \beta^{i/2}}}{T} \xi + T^{-1/2} \text{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{T-1} \beta^i \right) \boldsymbol{\Sigma} + \epsilon \mathbf{I}_d \right) \right]^{\frac{1}{2}} \sqrt{\xi}. \end{aligned}$$

It remains to provide a concrete bound for  $\|\mathbf{S}_T\|_{\text{op}}$ , which follows from Lemma 3.3 and Lemma D.6. Specifically, we know from Lemma 3.3 that

$$\|\mathbf{S}_T\|_{\text{op}} \leq C_1 \left( \log \left[ \frac{d}{\epsilon} \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 + d^2(1-\beta)T \right] + 1 \right) \left( \frac{(1-\beta)T}{\beta} + \log \left\| \frac{\mathbf{V}_{T-1}^2}{\epsilon} \right\|_{\text{op}} \right) + C_2. \quad (26)$$

Moreover, by Lemma D.6, we know that

$$\begin{aligned} \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 &= \text{poly}(T, \|\boldsymbol{\Sigma}\|_{\text{op}}, \|\nabla f(\mathbf{x}_0)\|_2, \Lambda_{\mathcal{H}}(f), d, \eta) \\ \|\mathbf{V}_{T-1}\|_{\text{op}} &= \text{poly}(T, \|\boldsymbol{\Sigma}\|_{\text{op}}, \|\nabla f(\mathbf{x}_0)\|_2, \Lambda_{\mathcal{H}}(f), d, \eta, \epsilon). \end{aligned}$$

Then the first term in Eq. 26 satisfies

$$\begin{aligned} \log \left[ \frac{d}{\epsilon} \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 + d^2(1-\beta)T \right] + 1 &= 1 + \log [\text{poly}(T, \|\boldsymbol{\Sigma}\|_{\text{op}}, \|\nabla f(\mathbf{x}_0)\|_2, \Lambda_{\mathcal{H}}(f), d, \eta, 1-\beta, \epsilon^{-1})] \\ &= \tilde{O}(\log d) \end{aligned}$$

where  $\tilde{O}(\cdot)$  hides logarithm dependence on all problem parameters except  $d$ . Similarly, the second term in Eq. 26 satisfies

$$\frac{(1-\beta)T}{\beta} + \log \left\| \frac{\mathbf{V}_{T-1}^2}{\epsilon} \right\|_{\text{op}} = \frac{(1-\beta)T}{\beta} + \tilde{O}(\log d)$$

Plugging these two bounds into Eq. 26 yields that for any well-structured  $\mathcal{H}$ ,

$$\|\mathbf{S}_T\|_{\text{op}} = \tilde{O} \left( \log d \cdot \left( \frac{(1-\beta)T}{\beta} + \log d \right) \right) \quad (27)$$

When  $\mathcal{H}$  only has diagonal matrices, we can improve Lemma D.6 with results in Xie et al. (2025a). The dependence on  $d$  can be improved such that  $\log(\mathbf{V}_{T-1}^2/\epsilon) = \tilde{O}(1)$ . Then Lemma 3.3 gives us

$$\|\mathbf{S}_T\|_{\text{op}} \leq (1-\beta)T + \log \left\| \frac{\mathbf{V}_{T-1}^2}{\epsilon} \right\|_{\text{op}} = (1-\beta)T + \tilde{O}(1).$$

This completes the proof.  $\square$

**Lemma D.3.** *Under the setting of Theorem D.7, let  $\tilde{\mathbf{V}}_t$  be as defined in Eq. 16. Then it holds that*

$$\sum_{t=0}^{T-1} \mathbb{E}_t \text{Tr}[\mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] \geq \frac{1}{2} \sum_{t=0}^{T-1} \mathbb{E}_t \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] - \frac{\sqrt{2\|\boldsymbol{\Sigma}\|_{\text{op}}}}{2} \sum_{t=0}^{T-1} \mathbb{E}_t \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \mathbf{g}_t^\top].$$

*Proof.* First we can compute the gap by replacing  $\mathbf{V}_t$  with  $\tilde{\mathbf{V}}_t$  as follows

$$\begin{aligned} \left| \text{Tr} \left[ (\mathbf{V}_t^{-1} - \tilde{\mathbf{V}}_t^{-1}) \mathbf{g}_t \bar{\mathbf{g}}_t^\top \right] \right| &= \left| \text{Tr} \left[ \tilde{\mathbf{V}}_t^{-1} (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top \right] \right| \\ &\leq \frac{1}{2} \text{Tr} \left[ \tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top \right] + \frac{1}{2} \text{Tr} \left[ (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \tilde{\mathbf{V}}_t^{-1} (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top \mathbf{V}_t^{-1} \right], \end{aligned} \quad (28)$$

where the second inequality follows from the fact that  $\text{Tr}[\mathbf{A}^\top \mathbf{B}] \leq \frac{1}{2} \text{Tr}[\mathbf{A}\mathbf{A}^\top] + \frac{1}{2} \text{Tr}[\mathbf{B}\mathbf{B}^\top]$  with  $\mathbf{A} = \tilde{\mathbf{V}}_t^{-\frac{1}{2}} \bar{\mathbf{g}}_t$  and  $\mathbf{B} = \tilde{\mathbf{V}}_t^{-\frac{1}{2}} (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \mathbf{V}_t^{-1} \mathbf{g}_t$ . Note that since  $\tilde{\mathbf{V}}_t \succeq \mathbf{V}_t$ , it holds that  $0 \preceq \tilde{\mathbf{V}}_t - \mathbf{V}_t \preceq \tilde{\mathbf{V}}_t$ . Then we can upper bound the second term on the right-hand side of (28) as

$$\begin{aligned} \text{Tr} \left[ (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \tilde{\mathbf{V}}_t^{-1} (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top \mathbf{V}_t^{-1} \right] &\leq \text{Tr} \left[ (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top \mathbf{V}_t^{-1} \right] \\ &\leq \|\tilde{\mathbf{V}}_t - \mathbf{V}_t\|_{\text{op}} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] \\ &= \|P_{\mathcal{H}}(\tilde{\mathbf{M}}_t) - P_{\mathcal{H}}(\mathbf{M}_t)\|_{\text{op}} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] \end{aligned}$$

Further applying Lemma B.5 yields

$$\begin{aligned} \text{Tr} \left[ (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \tilde{\mathbf{V}}_t^{-1} (\tilde{\mathbf{V}}_t - \mathbf{V}_t) \mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top \mathbf{V}_t^{-1} \right] &\leq \|\tilde{\mathbf{M}}_t - \mathbf{M}_t\|_{\text{op}}^{1/2} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] \\ &\leq \sqrt{2\|\boldsymbol{\Sigma}\|_{\text{op}}} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] \end{aligned} \quad (29)$$

where the second inequality holds because  $\mathbf{0} \preceq \tilde{\mathbf{M}}_t - \mathbf{M}_t = \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top + \boldsymbol{\Sigma} - \mathbf{g}_t \bar{\mathbf{g}}_t^\top \preceq 2\boldsymbol{\Sigma}$  based on Assumption D.1. Then plugging (29) back into (28), we get

$$\begin{aligned} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-1} \mathbf{g}_t \bar{\mathbf{g}}_t^\top] &\geq \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] - \sum_{t=0}^{T-1} \left| \text{Tr} \left[ (\mathbf{V}_t^{-1} - \tilde{\mathbf{V}}_t^{-1}) \mathbf{g}_t \bar{\mathbf{g}}_t^\top \right] \right| \\ &\geq \frac{1}{2} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] - \frac{\sqrt{2\|\boldsymbol{\Sigma}\|_{\text{op}}}}{2} \sum_{t=0}^{T-1} \text{Tr}[\mathbf{V}_t^{-2} \mathbf{g}_t \bar{\mathbf{g}}_t^\top]. \end{aligned}$$

$\square$

**Lemma D.4.** Under the setting of Theorem D.7, let  $\tilde{\mathbf{V}}_t$  be as defined in Eq. 16. Then it holds that

$$\left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] \right) \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t] \right) \geq \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)] \right)^2.$$

*Proof of Lemma D.4.* For convenience, denote  $\mathbf{A}_t = P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)$  and  $\mathbf{B}_t = \tilde{\mathbf{V}}_t^{\frac{1}{2}}$ . Then

$$\sum_{t=0}^{T-1} \text{Tr}[P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)] = \sum_{t=0}^{T-1} \text{Tr}[\mathbf{A}_t] = \sum_{t=0}^{T-1} \langle \mathbf{A}_t, \mathbf{I} \rangle = \sum_{t=0}^{T-1} \langle \mathbf{B}_t^{-1} \mathbf{A}_t, \mathbf{B}_t \rangle.$$

Now taking expectation on both sides and applying Cauchy-Schwarz inequality, we get

$$\mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)] \leq \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{B}_t^{-1} \mathbf{A}_t\|_F \|\mathbf{B}_t\|_F] \leq \sum_{t=0}^{T-1} (\mathbb{E}\|\mathbf{B}_t^{-1} \mathbf{A}_t\|_F^2)^{1/2} (\mathbb{E}\|\mathbf{B}_t\|_F^2)^{1/2}$$

Applying Cauchy-Schwarz inequality again, we further have

$$\begin{aligned} \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)] &\leq \left( \sum_{t=0}^{T-1} \mathbb{E}\|\mathbf{B}_t^{-1} \mathbf{A}_t\|_F^2 \right)^{1/2} \left( \sum_{t=0}^{T-1} \mathbb{E}\|\mathbf{B}_t\|_F^2 \right)^{1/2} \\ &= \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top)^2] \right)^{1/2} \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t] \right)^{1/2} \\ &= \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t^{-1} \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top] \right)^{1/2} \left( \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t] \right)^{1/2} \end{aligned}$$

where the second equality follows from Lemma B.2 as  $\tilde{\mathbf{V}}_t \in \mathcal{H}$ . This completes the proof.  $\square$

**Lemma D.5.** Under the setting of Theorem D.7, let  $\tilde{\mathbf{V}}_t$  be as defined in Eq. 16. Then it holds that

$$\mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\tilde{\mathbf{V}}_t] \leq T \cdot \text{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{T-1} \beta^i \right) \boldsymbol{\Sigma} + \epsilon \mathbf{I}_d \right) \right] + \left( \sum_{i=0}^{T-1} \beta^{i/2} \right) \mathbb{E} \sum_{t=0}^{T-1} \text{Tr}[\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top \tilde{\mathbf{V}}_t^{-1}]. \quad (30)$$

*Proof of Lemma D.5.* Recall that  $\tilde{\mathbf{V}}_t = P_{\mathcal{H}}(\tilde{\mathbf{M}}_t + \epsilon \mathbf{I}_d) \in \mathcal{H}$ . Then applying Lemma B.2, we get

$$\text{Tr}[\tilde{\mathbf{V}}_t] = \text{Tr}[\tilde{\mathbf{V}}_t^2 \tilde{\mathbf{V}}_t^{-1}] = \text{Tr}[(P_{\mathcal{H}}(\tilde{\mathbf{M}}_t + \epsilon \mathbf{I}_d))^2 \tilde{\mathbf{V}}_t^{-1}] = \text{Tr}[(\tilde{\mathbf{M}}_t + \epsilon \mathbf{I}_d) \tilde{\mathbf{V}}_t^{-1}]. \quad (31)$$

Further plugging in  $\tilde{\mathbf{M}}_t = \beta \mathbf{M}_{t-1} + \bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top + \boldsymbol{\Sigma}$ , we have

$$\begin{aligned} \text{Tr}[\tilde{\mathbf{V}}_t] &= \text{Tr}[(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top + \boldsymbol{\Sigma} + \beta \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d) \tilde{\mathbf{V}}_t^{-1}] \\ &= \text{Tr}[(\boldsymbol{\Sigma} + \beta \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d) \tilde{\mathbf{V}}_t^{-1}] + \text{Tr}[\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top \tilde{\mathbf{V}}_t^{-1}] \end{aligned}$$

Applying Lemma B.2 again, we further have

$$\begin{aligned} \text{Tr}[\tilde{\mathbf{V}}_t] &= \text{Tr}[P_{\mathcal{H}}(\boldsymbol{\Sigma} + \beta \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d)^2 \tilde{\mathbf{V}}_t^{-1}] + \text{Tr}[\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top \tilde{\mathbf{V}}_t^{-1}] \\ &\leq \text{Tr}[P_{\mathcal{H}}(\boldsymbol{\Sigma} + \beta \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d)] + \text{Tr}[\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top \tilde{\mathbf{V}}_t^{-1}] \end{aligned} \quad (32)$$

where the inequality holds because  $\tilde{\mathbf{V}}_t = P_{\mathcal{H}}(\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top + \boldsymbol{\Sigma} + \beta \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d) \succeq P_{\mathcal{H}}(\boldsymbol{\Sigma} + \beta \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d)$ . Next, we will further control the right-hand side of the above inequality by recursive expansion.

For notational convenience, for any  $1 \leq s < t$ , denote

$$\mathbf{A}_s = \left( \sum_{i=0}^{s-1} \beta^i \right) \boldsymbol{\Sigma} + \beta^s \mathbf{M}_{t-s} + \epsilon \mathbf{I}_d.$$

Then for any  $1 \leq s < t$ , we have

$$\begin{aligned} \mathbb{E} \operatorname{Tr}[P_{\mathcal{H}}(\mathbf{A}_s)] &= \mathbb{E} \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{s-1} \beta^i \right) \Sigma + \beta^{s+1} \mathbf{M}_{t-s-1} + \beta^s \mathbf{g}_{t-s} \mathbf{g}_{t-s}^\top + \epsilon \mathbf{I}_d \right) \right] \\ &\leq \mathbb{E} \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \mathbb{E}_{t-s-1} \left( \sum_{i=0}^{s-1} \beta^i \right) \Sigma + \beta^{s+1} \mathbf{M}_{t-s-1} + \beta^s \mathbf{g}_{t-s} \mathbf{g}_{t-s}^\top + \epsilon \mathbf{I}_d \right) \right] \end{aligned}$$

where the inequality holds because  $\operatorname{Tr}[P_{\mathcal{H}}(\mathbf{X})]$  is concave in  $\mathbf{X}$  by Lemma B.6. Then since  $\mathbb{E} \mathbf{g}_{t-s} \mathbf{g}_{t-s}^\top \preceq \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top + \Sigma$  by Assumption D.1, we further have

$$\begin{aligned} \mathbb{E} \operatorname{Tr}[P_{\mathcal{H}}(\mathbf{A}_s)] &\leq \mathbb{E} \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^s \beta^i \right) \Sigma + \beta^{s+1} \mathbf{M}_{t-s-1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top + \epsilon \mathbf{I}_d \right) \right] \\ &= \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top)]. \end{aligned}$$

Applying the same trick as in (31) to  $\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top$ , we obtain

$$\begin{aligned} \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_s)] &\leq \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top)^2 P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top)^{-1}] \\ &= \mathbb{E} \operatorname{Tr} [(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top) P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top)^{-1}] \\ &= \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_{s+1})^2 P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top)^{-1}] \\ &\quad + \mathbb{E} \operatorname{Tr} [\beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top)^{-1}] \\ &\leq \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_{s+1})] + \mathbb{E} \operatorname{Tr} [\beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top)^{-1}] \quad (33) \end{aligned}$$

where the second equality is again by Lemma B.2 and the second inequality follows from the fact that  $P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top) \succeq P_{\mathcal{H}}(\mathbf{A}_{s+1})$ . In addition, note that

$$\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top = \beta^s \left( \frac{\sum_{i=0}^s \beta^i}{\beta^s} \Sigma + \beta \mathbf{M}_{t-s-1} + \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top \right) + \epsilon \mathbf{I}_d \succeq \beta^s (\tilde{\mathbf{M}}_{t-s} + \epsilon \mathbf{I}_d).$$

Thus it follows from Lemma B.3 that

$$P_{\mathcal{H}}(\mathbf{A}_{s+1} + \beta^s \bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top) \succeq P_{\mathcal{H}}(\beta^s (\tilde{\mathbf{M}}_{t-s} + \epsilon \mathbf{I}_d)) = \beta^{s/2} P_{\mathcal{H}}(\tilde{\mathbf{M}}_{t-s} + \epsilon \mathbf{I}_d) = \beta^{s/2} \tilde{\mathbf{V}}_{t-s}. \quad (34)$$

Now combining (33) and (34) yields

$$\mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_s)] \leq \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_{s+1})] + \beta^{s/2} \mathbb{E} \operatorname{Tr} [\bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top \tilde{\mathbf{V}}_{t-s}^{-1}].$$

Further telescoping over  $s = 1, \dots, t-1$ , we obtain

$$\begin{aligned} \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\Sigma + \beta \mathbf{M}_{t-1} + \epsilon \mathbf{I}_d)] &= \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_1)] \\ &\leq \mathbb{E} \operatorname{Tr} [P_{\mathcal{H}}(\mathbf{A}_t)] + \sum_{s=1}^{t-1} \beta^{s/2} \mathbb{E} \operatorname{Tr} [\bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top \tilde{\mathbf{V}}_{t-s}^{-1}] \\ &= \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{t-1} \beta^i \right) \Sigma + \epsilon \mathbf{I}_d \right) \right] + \sum_{s=1}^{t-1} \beta^{s/2} \mathbb{E} \operatorname{Tr} [\bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top \tilde{\mathbf{V}}_{t-s}^{-1}] \quad (35) \end{aligned}$$

Now, plugging (35) back into (32), we obtain

$$\begin{aligned} \mathbb{E} \operatorname{Tr} [\tilde{\mathbf{V}}_t] &\leq \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{t-1} \beta^i \right) \Sigma + \epsilon \mathbf{I}_d \right) \right] + \sum_{s=0}^{t-1} \beta^{s/2} \mathbb{E} \operatorname{Tr} [\bar{\mathbf{g}}_{t-s} \bar{\mathbf{g}}_{t-s}^\top \tilde{\mathbf{V}}_{t-s}^{-1}] \\ &= \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{t-1} \beta^i \right) \Sigma + \epsilon \mathbf{I}_d \right) \right] + \sum_{s=1}^t \beta^{(t-s)/2} \mathbb{E} \operatorname{Tr} [\bar{\mathbf{g}}_s \bar{\mathbf{g}}_s^\top \tilde{\mathbf{V}}_s^{-1}]. \end{aligned}$$

Finally, summing over  $t = 0, 1, \dots, T-1$ , we get

$$\begin{aligned} \mathbb{E} \sum_{t=0}^{T-1} \operatorname{Tr} [\tilde{\mathbf{V}}_t] &\leq \sum_{t=0}^{T-1} \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{t-1} \beta^i \right) \Sigma + \epsilon \mathbf{I}_d \right) \right] + \sum_{t=0}^{T-1} \sum_{s=1}^t \beta^{(t-s)/2} \mathbb{E} \operatorname{Tr} [\bar{\mathbf{g}}_s \bar{\mathbf{g}}_s^\top \tilde{\mathbf{V}}_s^{-1}] \\ &\leq T \cdot \operatorname{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{T-1} \beta^i \right) \Sigma + \epsilon \mathbf{I}_d \right) \right] + \left( \sum_{i=0}^{T-1} \beta^{i/2} \right) \mathbb{E} \sum_{t=0}^{T-1} \operatorname{Tr} [\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top \tilde{\mathbf{V}}_t^{-1}]. \end{aligned}$$

This completes the proof.  $\square$

**Lemma D.6.** *Under the setting of Theorem D.2, for any fixed initialization  $\mathbf{x}_0$ , let  $\mathbf{g}_0, \dots, \mathbf{g}_{T-1}$  and  $\mathbf{V}_0, \dots, \mathbf{V}_{T-1}$  be given by Algorithm 6. Then the following hold with probability 1:*

$$\begin{aligned} \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 &\leq 2(\|\boldsymbol{\Sigma}\|_{\text{op}} + \|\nabla f(\mathbf{x}_0)\|_2^2 + \Lambda_{\mathcal{H}}(f)^2 d\eta^2)T^3, \\ \|\mathbf{V}_{T-1}\|_{\text{op}} &\leq \sqrt{2(\|\boldsymbol{\Sigma}\|_{\text{op}} + \|\nabla f(\mathbf{x}_0)\|_2^2 + \Lambda_{\mathcal{H}}(f)^2 d\eta^2)T^3 + \epsilon d}. \end{aligned}$$

*Proof.* In this proof, we will define  $\mathbf{H}^* \in \arg \min_{\mathbf{H} \in \mathcal{H}, -\mathbf{H} \preceq \nabla^2 f(\mathbf{x}) \preceq \mathbf{H}} \text{Tr}(\mathbf{H})$ .

We first control  $\sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2$ . According to Assumption D.1, we have

$$\sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 = \sum_{t=0}^{T-1} \|\mathbf{g}_t \mathbf{g}_t^\top\|_{\text{op}} \leq \sum_{t=0}^{T-1} \left( \|\boldsymbol{\Sigma}\|_{\text{op}} + \|\bar{\mathbf{g}}_t \bar{\mathbf{g}}_t^\top\|_{\text{op}} \right) = T \|\boldsymbol{\Sigma}\|_{\text{op}} + \sum_{t=0}^{T-1} \|\bar{\mathbf{g}}_t\|_2^2 \quad (36)$$

So it suffices to control the sum of  $\|\bar{\mathbf{g}}_t\|_2^2$ . By triangle inequality,  $\|\bar{\mathbf{g}}_t\|_2^2 \leq 2\|\bar{\mathbf{g}}_0\|_2^2 + 2\|\bar{\mathbf{g}}_t - \bar{\mathbf{g}}_0\|_2^2$ , and thus we only need to bound the distance  $\bar{\mathbf{g}}_t - \bar{\mathbf{g}}_0$ . To this end, since  $-\mathbf{H}^* \preceq \nabla^2 f(\mathbf{x}) \preceq \mathbf{H}^*$  for all  $\mathbf{x}$ , we have

$$\|\bar{\mathbf{g}}_t - \bar{\mathbf{g}}_0\|_2 = \|\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_0)\|_2 \leq \|\mathbf{H}^*\|_{\text{op}} \|\mathbf{x}_t - \mathbf{x}_0\|_2 \leq \Lambda_{\mathcal{H}}(f) \sum_{s=0}^{t-1} \|\mathbf{x}_{s+1} - \mathbf{x}_s\|_2. \quad (37)$$

Moreover, we can control each  $\|\mathbf{x}_{s+1} - \mathbf{x}_s\|_2$  as follows:

$$\begin{aligned} \|\mathbf{x}_{s+1} - \mathbf{x}_s\|_2^2 &= \eta^2 \|\mathbf{V}_s^{-1} \mathbf{g}_s\|_2^2 = \eta^2 \text{Tr}(\mathbf{g}_s^\top \mathbf{V}_s^{-2} \mathbf{g}_s) \\ &= \eta^2 \text{Tr}(P_{\mathcal{H}}(\mathbf{M}_s)^{-2} \mathbf{g}_s \mathbf{g}_s^\top) \\ &= \eta^2 \text{Tr}(P_{\mathcal{H}}(\mathbf{M}_s)^{-2} P_{\mathcal{H}}(\mathbf{g}_s \mathbf{g}_s^\top)^2) \end{aligned}$$

where the last equality holds by Lemma B.2. Then since  $\mathbf{M}_s \succeq \mathbf{g}_s \mathbf{g}_s^\top$ , we have  $P_{\mathcal{H}}(\mathbf{M}_s) \succeq P_{\mathcal{H}}(\mathbf{g}_s \mathbf{g}_s^\top)$  by Lemma B.3, and it follows that

$$\|\mathbf{x}_{s+1} - \mathbf{x}_s\|_2^2 \leq \eta^2 \text{Tr}(\mathbf{I}_d) = d\eta^2.$$

Plugging this back into Eq. 37, we get  $\|\bar{\mathbf{g}}_t - \bar{\mathbf{g}}_0\|_2 \leq \eta t \Lambda_{\mathcal{H}}(f) \sqrt{d}$ , so  $\|\bar{\mathbf{g}}_t\|_2^2 \leq 2\|\bar{\mathbf{g}}_0\|_2^2 + 2\eta^2 t^2 \Lambda_{\mathcal{H}}(f)^2 d$ . Now applying this to Eq. 36, we obtain

$$\begin{aligned} \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 &\leq T \|\boldsymbol{\Sigma}\|_{\text{op}} + \sum_{t=0}^{T-1} (2\|\bar{\mathbf{g}}_0\|_2^2 + 2\eta^2 t^2 \Lambda_{\mathcal{H}}(f)^2 d) \\ &\leq 2(\|\boldsymbol{\Sigma}\|_{\text{op}} + \|\bar{\mathbf{g}}_0\|_2^2 + \Lambda_{\mathcal{H}}(f)^2 d\eta^2)T^3. \end{aligned} \quad (38)$$

Finally, we can bound  $\|\mathbf{V}_{T-1}\|_{\text{op}}$  as follows:

$$\|\mathbf{V}_{T-1}\|_{\text{op}}^2 = \|\mathbf{V}_{T-1}^2\|_{\text{op}} = \|P_{\mathcal{H}}(\mathbf{M}_{T-1})^2\|_{\text{op}} \leq \text{Tr}(P_{\mathcal{H}}(\mathbf{M}_{T-1})^2) = \text{Tr}(\mathbf{M}_{T-1})$$

where we apply Lemma B.2 in the last equality. Note that  $\text{Tr}(\mathbf{M}_{T-1}) = \sum_{t=0}^{T-1} \|\mathbf{g}_t\|_2^2 + \epsilon d$ , so it follows from (38) that

$$\|\mathbf{V}_{T-1}\|_{\text{op}} \leq \sqrt{2(\|\boldsymbol{\Sigma}\|_{\text{op}} + \|\bar{\mathbf{g}}_0\|_2^2 + \Lambda_{\mathcal{H}}(f)^2 d\eta^2)T^3 + \epsilon d}.$$

This completes the proof.  $\square$

### D.3 PROOF FOR THE CUMULATIVE AND EMA VARIANTS

We will derive the convergence rate of cumulative optimizers and EMA optimizers by reducing from Theorem D.2. In the stochastic case, the dependence on  $T$  of both convergence rate is  $\tilde{O}(T^{-1/4})$ , matching previous results on specific examples AdaGrad and Adam (Xie et al., 2025a; Li et al., 2025).

**Theorem D.7.** For any  $\epsilon \geq 0$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 1 with well-structured preconditioner set  $\mathcal{H}$ , where the update of  $\mathbf{M}_t$  follows the cumulative version, i.e.,  $\mathbf{M}_t = \mathbf{M}_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$  for all  $t \in [T]$ . Let  $\Lambda_{\mathcal{H}}(f)$  be the adaptive smoothness of the loss  $f$  according to Definition 2.4. Then under Assumption D.1, it holds that

$$\mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \right] \leq \frac{1}{\sqrt{T}} \xi + \frac{1}{T^{1/4}} \text{Tr}[P_{\mathcal{H}}(\boldsymbol{\Sigma} + \frac{\epsilon}{T} \mathbf{I}_d)]^{\frac{1}{2}} \sqrt{\xi}$$

where  $\xi$  is given by

$$\xi = \tilde{O} \left( \frac{\Delta_0}{\eta} + (\eta \cdot \Lambda_{\mathcal{H}}(f) + d \|\boldsymbol{\Sigma}\|_{\text{op}}^{1/2}) \log^2 d \right).$$

Moreover, when setting the learning rate  $\eta = \sqrt{\frac{\Delta_0}{\Lambda_{\mathcal{H}}(f) \log^2 d}}$ , it holds that  $\xi = \tilde{O}(\sqrt{\Delta_0 \cdot \Lambda_{\mathcal{H}}(f)} \log d + d \|\boldsymbol{\Sigma}\|_{\text{op}}^{1/2} \log^2 d)$ .

*Proof of Theorem D.7.* The desired result directly follows from Theorem D.2 with  $\beta = 1$ , where we additionally apply the bound  $\text{Tr}[P_{\mathcal{H}}(T\boldsymbol{\Sigma} + \epsilon \mathbf{I}_d)] \leq \sqrt{T} \text{Tr}[P_{\mathcal{H}}(\boldsymbol{\Sigma} + \frac{\epsilon}{T} \mathbf{I}_d)]$  using Lemma B.3.  $\square$

**Theorem D.8.** For any  $\epsilon \geq 0$ ,  $\beta \in (0, 1)$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 1 with well-structured preconditioner set  $\mathcal{H}$ , where the update of  $\mathbf{M}_t$  follows the EMA version, i.e.,  $\mathbf{M}_t = \beta \mathbf{M}_{t-1} + (1 - \beta) \mathbf{g}_t \mathbf{g}_t^\top$  for all  $t \in [T]$ . Let  $\Lambda_{\mathcal{H}}(f)$  be the adaptive smoothness of the loss  $f$  according to Definition 2.4. Then under Assumption D.1, it holds that

$$\mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \right] \leq \frac{1}{\sqrt{T}} \kappa + \frac{1}{T^{1/4}} \text{Tr}[P_{\mathcal{H}}(\boldsymbol{\Sigma} + \epsilon \mathbf{I}_d)]^{\frac{1}{2}} \sqrt{\kappa}$$

where  $\kappa$  is given by

$$\kappa = \tilde{O} \left( \frac{\Delta_0}{\eta \sqrt{T}} + \frac{\eta \cdot \Lambda_{\mathcal{H}}(f) + d \sqrt{1 - \beta} \|\boldsymbol{\Sigma}\|_{\text{op}}^{1/2}}{\sqrt{T}} \left( \frac{1}{\beta} T + \frac{\log d}{1 - \beta} \right) \log d \right).$$

Moreover, when setting  $1 - \beta = \Theta(\frac{\log d}{T})$  and  $\eta = \sqrt{\frac{\Delta_0}{\Lambda_{\mathcal{H}}(f) \cdot T \log^2 d}}$ , it holds that

$$\kappa = \tilde{O} \left( \left( \sqrt{\Delta_0 \cdot \Lambda_{\mathcal{H}}(f)} + d \|\boldsymbol{\Sigma}\|_{\text{op}}^{1/2} \right) \log d \right).$$

When there is no noise, i.e.,  $f_t \equiv f$ , it holds that

$$\mathbb{E} \left[ \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \right] \leq \frac{1}{\sqrt{T}} \kappa + \frac{\sqrt{d} \epsilon^{1/4}}{T^{1/4}} \sqrt{\kappa}$$

where  $\kappa$  is given by

$$\kappa = \tilde{O} \left( \frac{\Delta_0}{\eta \sqrt{T}} + \frac{\eta \cdot \Lambda_{\mathcal{H}}(f)}{\sqrt{T}} \left( \frac{1}{\beta} T + \frac{\log d}{1 - \beta} \right) \log d \right).$$

Moreover, when setting  $1 - \beta = \Theta(\frac{\log d}{T})$  and  $\eta = \sqrt{\frac{\Delta_0}{\Lambda_{\mathcal{H}}(f) \cdot T \log^2 d}}$ , it holds that

$$\kappa = \tilde{O} \left( \sqrt{\Delta_0 \cdot \Lambda_{\mathcal{H}}(f)} \log d \right).$$

*Proof of Theorem D.8.* As mentioned in Section 3, Algorithm 5 is equivalent to Algorithm 6 with  $\epsilon/(1 - \beta)$  in place of  $\epsilon$  and  $\eta/\sqrt{1 - \beta}$  in place of  $\eta$ . Therefore, the convergence rate of Algorithm 5 follows from Theorem D.2 with this change of hyperparameters. Specifically,

$$\frac{1}{T} \mathbb{E} \sum_{t=0}^{T-1} \|\bar{\mathbf{g}}_t\|_{\mathcal{H},*} \leq \frac{\sqrt{\sum_{i=0}^{T-1} \beta^{i/2}}}{T} \xi + \frac{\sqrt{\text{Tr}[P_{\mathcal{H}}((\sum_{i=0}^{T-1} \beta^i) \boldsymbol{\Sigma} + \frac{\epsilon}{1 - \beta} \mathbf{I}_d)]}}{\sqrt{T}} \sqrt{\xi} \quad (39)$$

where  $\xi$  satisfies

$$\xi = \tilde{O} \left( \frac{\Delta_0 \sqrt{1-\beta}}{\eta} + \left( \frac{\eta}{\sqrt{1-\beta}} \Lambda_{\mathcal{H}}(f) + d \|\Sigma\|_{\text{op}}^{1/2} \right) \left( \frac{1-\beta}{\beta} T + \log d \right) \log d \right). \quad (40)$$

We can further simplify the result for  $\beta < 1$  as follows. Note that  $\sum_{i=0}^{T-1} \beta^i \leq 1/(1-\beta)$  and

$$\sum_{i=0}^{T-1} \beta^{i/2} \leq \frac{1}{1-\sqrt{\beta}} = \frac{1+\sqrt{\beta}}{1-\beta} \leq \frac{2}{1-\beta}.$$

It then follows from Lemma B.3 that

$$\text{Tr} \left[ P_{\mathcal{H}} \left( \left( \sum_{i=0}^{T-1} \beta^i \right) \Sigma + \frac{\epsilon}{1-\beta} \mathbf{I}_d \right) \right] \leq \text{Tr} \left[ P_{\mathcal{H}} \left( \frac{1}{1-\beta} (\Sigma + \epsilon \mathbf{I}_d) \right) \right] = \frac{1}{\sqrt{1-\beta}} \text{Tr} [P_{\mathcal{H}}(\Sigma + \epsilon \mathbf{I}_d)].$$

Applying these to (39), we get

$$\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\bar{\mathbf{g}}_t\|_{\mathcal{H},*} \leq \frac{\sqrt{2}}{T} \frac{\xi}{\sqrt{1-\beta}} + \frac{\sqrt{\text{Tr}[P_{\mathcal{H}}(\Sigma + \epsilon \mathbf{I}_d)]}}{\sqrt{T}} \sqrt{\frac{\xi}{\sqrt{1-\beta}}}.$$

Denote  $\kappa = \xi / \sqrt{(1-\beta)T}$ . Then we have

$$\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\bar{\mathbf{g}}_t\|_{\mathcal{H},*} \leq \frac{\sqrt{2}}{\sqrt{T}} \kappa + \frac{\sqrt{\text{Tr}[P_{\mathcal{H}}(\Sigma + \epsilon \mathbf{I}_d)]}}{T^{1/4}} \sqrt{\kappa}.$$

By (40), we know that  $\kappa$  satisfies that

$$\kappa = \tilde{O} \left( \frac{\Delta_0}{\eta \sqrt{T}} + \frac{\eta \Lambda_{\mathcal{H}}(f) + d \sqrt{1-\beta} \|\Sigma\|_{\text{op}}^{1/2}}{\sqrt{T}} \left( \frac{1}{\beta} T + \frac{\log d}{1-\beta} \right) \log d \right).$$

This completes the proof.  $\square$

## E PROOF FOR THE ACCELERATED ALGORITHM

In this section, we will define  $\mathbf{H}^* = \arg \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} L_{\|\cdot\|_{\mathbf{H}}}(f)$ . Our proof follows the strategy developed by Kovalev (2025a). For completeness, we reproduce some of the arguments to make the exposition self-contained.

### E.1 STOCHASTIC CASE WITHOUT PROJECTION

**Theorem 4.3.** *For a well-structured preconditioner set  $\mathcal{H}$ , let  $f$  be a convex loss function whose  $\mathcal{H}$ -smoothness constant is  $\Lambda_{\mathcal{H}}(f) \in (0, \infty)$  according to Definition 2.4. For  $\epsilon > 0$ ,  $T > 0$ , consider Algorithm 2 with  $\alpha_t = 2/(t+2)$  for  $t = 0, 1, \dots, T-1$ . Suppose  $\mathbf{x}^*$  is the global minima and  $\max_{t=0,1,\dots,T-1} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathcal{H}} \leq D$  for some  $D > 0$  and Assumption 4.2 holds with adaptive gradient variance  $\sigma_{\mathcal{H}}(\{f_t\}_{t=1}^T)^2 \leq \sigma_{\mathcal{H}}^2$  for some  $\sigma_{\mathcal{H}} \in [0, \infty)$ . Then it holds that*

$$\begin{aligned} \mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] &\leq \frac{2D^2\epsilon}{\eta(T+1)^2} \mathbb{E} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left( \frac{D^2}{2\eta} - \frac{\eta}{2} \right) \mathbb{E} \frac{4}{(T+1)^2} \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \\ &\quad + \frac{2\eta^2}{(T+1)^2} \cdot \Lambda_{\mathcal{H}}(f) \cdot \tilde{O}(\log^2 d) + \frac{\eta}{T^{1/2}} \sigma_{\mathcal{H}} \cdot \tilde{O}(\log d). \end{aligned}$$

Moreover, when choosing learning rate  $\eta = D$ , the convergence rate becomes

$$\mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] = \tilde{O} \left( \frac{\Lambda_{\mathcal{H}}(f) D^2 \log^2 d + d \sqrt{\epsilon} D}{T^2} + \frac{\sigma_{\mathcal{H}} D \log d}{\sqrt{T}} \right).$$

*Proof of Theorem 4.3.* Combining the results of Lemma E.2 and Lemma E.3, we get

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{E}[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] &\leq \mathbb{E}\left[\frac{D^2\epsilon}{2\eta} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left(\frac{D^2}{2\eta} - \frac{\eta}{2}\right) \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t\right] \\ &\quad + \mathbb{E}\left[\eta \sum_{t=0}^{T-1} \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle\right] \\ &\quad + \mathbb{E}\left[\frac{\Lambda_{\mathcal{H}}(f)\eta^2}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2\right]. \end{aligned} \quad (41)$$

It remains to further bound the last two terms.

For the last term on the right-hand side of (41), we apply Lemma 3.3 to get

$$\sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 \leq \text{Tr}[\mathbf{H}^*] \cdot \tilde{O}(\log^2 d) = \Lambda_{\mathcal{H}}(f) \cdot \tilde{O}(\log^2 d) \quad (42)$$

where the second step follows from the definition of  $\mathbf{H}^*$ .

For the third term on the right-hand side of (41), consider any  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}(\mathbf{H}) \leq 1$ , and then it follows from the Cauchy-Schwarz inequality that

$$\begin{aligned} \mathbb{E}\left|\sum_{t=0}^{T-1} \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle\right| &\leq \sum_{t=0}^{T-1} (\mathbb{E}\|\mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t)\|_2^2)^{1/2} (\mathbb{E}\|\mathbf{V}_t^{-1} \mathbf{g}_t\|_2^2)^{1/2} \\ &\leq \left(\mathbb{E}\sum_{t=0}^{T-1} \|\mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t)\|_{\mathbf{H}^{-1}}^2\right)^{1/2} \left(\mathbb{E}\sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2\right)^{1/2}. \end{aligned}$$

Here we similarly have  $\mathbb{E}\sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}}^2 = \text{Tr}[\mathbf{H}] \cdot \tilde{O}(\log^2 d) \leq \tilde{O}(\log^2 d)$  according to Lemma 3.3. Then we further minimize over  $\mathbf{H} \in \mathcal{H}$  with  $\text{Tr}[\mathbf{H}] \leq 1$  to get

$$\begin{aligned} &\min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \mathbb{E}\sum_{t=0}^{T-1} \|\mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t)\|_{\mathbf{H}^{-1}}^2 \\ &= \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \mathbb{E}\sum_{t=0}^{T-1} \frac{1}{\alpha_t^2} \|\nabla f_t(\alpha_t \mathbf{x}_t + (1 - \alpha_t) \bar{\mathbf{x}}_t) - \nabla f(\alpha_t \mathbf{x}_t + (1 - \alpha_t) \bar{\mathbf{x}}_t)\|_{\mathbf{H}^{-1}}^2 \\ &\leq \sigma_{\mathcal{H}}^2 \sum_{t=0}^{T-1} \frac{1}{\alpha_t^2} \end{aligned}$$

where the last equation follows from Assumption 4.2 and the condition on  $\sigma_{\mathcal{H}}$ . Therefore, we obtain

$$\mathbb{E}\left|\sum_{t=0}^{T-1} \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle\right| \leq \sigma_{\mathcal{H}} \left(\sum_{t=0}^{T-1} \alpha_t^{-2}\right)^{1/2} \cdot \tilde{O}(\log^2 d) \leq \sigma_{\mathcal{H}} T^{3/2} \cdot \tilde{O}(\log^2 d) \quad (43)$$

where we have plugged in the choice of  $\alpha_t = 2/(t+2)$ .

Finally, combining (41), (42) and (43), we obtain

$$\begin{aligned} \mathbb{E}\sum_{t=0}^{T-1} (f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)) &\leq \frac{D^2\epsilon}{2\eta} \mathbb{E}\text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left(\frac{D^2}{2\eta} - \frac{\eta}{2}\right) \mathbb{E}\sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \\ &\quad + \frac{\eta^2}{2} \Lambda_{\mathcal{H}}(f) \cdot \tilde{O}(\log^2 d) + \eta \sigma_{\mathcal{H}} T^{3/2} \cdot \tilde{O}(\log d) \end{aligned}$$

The proof is then completed by further applying Lemma E.1.  $\square$

**Lemma E.1.** *Under the setting of Theorem 4.3, it holds that*

$$\mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] \leq \frac{4}{(T+1)^2} \sum_{t=0}^{T-1} \mathbb{E}[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] \quad (44)$$

*Proof of Lemma E.1.* Plugging in the definition of  $f^{\alpha_t, \bar{\mathbf{x}}_t}$  in (8), we have

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{E}[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1})] &= \sum_{t=0}^{T-1} \alpha_t^{-2} \mathbb{E}[f(\alpha_t \mathbf{x}^* + (1-\alpha_t)\bar{\mathbf{x}}_t) - f(\alpha_t \mathbf{x}_{t+1} + (1-\alpha_t)\bar{\mathbf{x}}_t)] \\ &\leq \sum_{t=0}^{T-1} \alpha_t^{-2} \mathbb{E}[\alpha_t f(\mathbf{x}^*) + (1-\alpha_t)f(\bar{\mathbf{x}}_t) - f(\bar{\mathbf{x}}_{t+1})] \\ &= \sum_{t=0}^{T-1} \alpha_t^{-2} \mathbb{E}[(1-\alpha_t)(f(\bar{\mathbf{x}}_t) - f(\mathbf{x}^*)) - (f(\bar{\mathbf{x}}_{t+1}) - f(\mathbf{x}^*))] \\ &= \sum_{t=1}^{T-1} (\alpha_t^{-2}(1-\alpha_t) - \alpha_{t-1}^{-2}) \mathbb{E}[f(\bar{\mathbf{x}}_t) - f(\mathbf{x}^*)] \\ &\quad + \alpha_0^{-2}(1-\alpha_0) \mathbb{E}[f(\bar{\mathbf{x}}_0) - f(\mathbf{x}^*)] - \alpha_{T-1}^{-2} \mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] \end{aligned}$$

where the inequality follows from the convexity of  $f$ . Plugging in the choice of  $\alpha_t = 2/(t+2)$ , we further obtain

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{E}[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1})] &= \frac{(T+1)^2}{4} \mathbb{E}[f(\mathbf{x}^*) - f(\bar{\mathbf{x}}_T)] - \frac{1}{4} \sum_{t=1}^{T-1} \mathbb{E}[f(\bar{\mathbf{x}}_t) - f(\mathbf{x}^*)] \\ &\leq \frac{(T+1)^2}{4} \mathbb{E}[f(\mathbf{x}^*) - f(\bar{\mathbf{x}}_T)]. \end{aligned}$$

This completes the proof.  $\square$

**Lemma E.2.** *Under the setting of Theorem 4.3, it holds that*

$$\sum_{t=0}^{T-1} \mathbb{E}\left[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)\right] \leq \mathbb{E}\left[\frac{D^2\epsilon}{2\eta} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left(\frac{D^2}{2\eta} + \frac{\eta}{2}\right) \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t\right].$$

*Proof of Lemma E.2.* For every  $t = 0, 1, \dots, T-1$ , by the convexity of  $f^{\alpha_t, \bar{\mathbf{x}}_t}$ , it follows from the Taylor expansion of  $f^{\alpha_t, \bar{\mathbf{x}}_t}$  at  $\mathbf{x}_t$  that

$$f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*) \leq \langle \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}^* \rangle = \langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle + \langle \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle$$

Since  $\mathbf{g}_t$  is an unbiased estimate of  $\nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t)$ , taking expectation on both sides yields

$$\mathbb{E}[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] \leq \mathbb{E}[\langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle]$$

By definition  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{V}_t^{-1} \mathbf{g}_t$ , so it follows from Lemma E.4 that

$$\langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle \leq \frac{1}{2\eta} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_t}^2 - \frac{1}{2\eta} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|_{\mathbf{V}_t}^2 + \frac{\eta}{2} \|\mathbf{g}_t\|_{\mathbf{V}_t^{-1}}^2$$

Define  $\mathbf{V}_{-1} = 0$  for convenience. Then combining the above two equations yields

$$\begin{aligned} \mathbb{E}[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] &\leq \mathbb{E}\left[\frac{1}{2\eta} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_t}^2 - \frac{1}{2\eta} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|_{\mathbf{V}_t}^2 + \frac{\eta}{2} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2\right] \\ &= \mathbb{E}\left[\frac{1}{2\eta} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_{t-1}}^2 - \frac{1}{2\eta} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|_{\mathbf{V}_t}^2\right] \\ &\quad + \mathbb{E}\left[\frac{1}{2\eta} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_t - \mathbf{V}_{t-1}}^2 + \frac{\eta}{2} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2\right]. \end{aligned}$$

Summing over  $t$ , we obtain

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{E} [f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] &\leq \mathbb{E} \left[ -\frac{1}{2\eta} \|\mathbf{x}_T - \mathbf{x}^*\|_{\mathbf{V}_{T-1}}^2 + \frac{1}{2\eta} \sum_{t=0}^{T-1} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_t - \mathbf{V}_{t-1}}^2 \right] \\ &\quad + \mathbb{E} \left[ \frac{\eta}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2 \right]. \end{aligned}$$

Since  $\mathbf{V}_t \succeq \mathbf{V}_{t-1}$ , we have  $\|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_t - \mathbf{V}_{t-1}}^2 \leq \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathcal{H}}^2 \text{Tr}(\mathbf{V}_t - \mathbf{V}_{t-1}) \leq D^2 \cdot \text{Tr}(\mathbf{V}_t - \mathbf{V}_{t-1})$ , where the second inequality holds because of the assumption that  $\max_{t \in [T]} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathcal{H}} \leq D$ . Therefore, we further have

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{E} [f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] &\leq \mathbb{E} \left[ \frac{D^2}{2\eta} \sum_{t=0}^{T-1} \text{Tr}(\mathbf{V}_t - \mathbf{V}_{t-1}) + \frac{\eta}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2 \right] \\ &= \mathbb{E} \left[ \frac{D^2}{2\eta} \text{Tr}(\mathbf{V}_{T-1}) + \frac{\eta}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2 \right] \\ &= \mathbb{E} \left[ \frac{D^2}{2\eta} \langle \mathbf{M}_{T-1} + \epsilon \mathbf{I}_d, \mathbf{V}_{T-1}^{-1} \rangle + \frac{\eta}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2 \right] \\ &= \mathbb{E} \left[ \frac{D^2}{2\eta} \epsilon \cdot \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \frac{D^2}{2\eta} \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_{T-1}^{-1} \mathbf{g}_t + \frac{\eta}{2} \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \right]. \end{aligned}$$

Since  $\mathbf{V}_t \preceq \mathbf{V}_{T-1}$  for all  $t \in [T-1]$ , we have  $\mathbf{g}_t^\top \mathbf{V}_{T-1}^{-1} \mathbf{g}_t \leq \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t$  for all  $t \in [T-1]$ . Consequently,

$$\sum_{t=0}^{T-1} \mathbb{E} [f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] \leq \mathbb{E} \left[ \frac{D^2 \epsilon}{2\eta} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left( \frac{D^2}{2\eta} + \frac{\eta}{2} \right) \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \right].$$

This completes the proof.  $\square$

**Lemma E.3.** *Under the setting of Theorem 4.3, it holds that*

$$\begin{aligned} \sum_{t=0}^{T-1} (f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)) &\leq \sum_{t=0}^{T-1} (f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)) - \eta \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2 \\ &\quad + \eta \sum_{t=0}^{T-1} \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle + \frac{\Lambda_{\mathcal{H}}(f) \eta^2}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2. \end{aligned}$$

*Proof of Lemma E.3.* By Definition 2.4, we have that

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\Lambda_{\mathcal{H}}(f)}{2} \|\mathbf{x} - \mathbf{y}\|_{\mathbf{H}^*}^2.$$

By the definition of  $f^{\alpha_t, \bar{\mathbf{x}}_t}$  in (8),

$$\begin{aligned} f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1}) &\leq f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) + \langle \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \frac{\Lambda_{\mathcal{H}}(f)}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_{\mathbf{H}^*}^2 \\ &= f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - \eta \langle \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle + \frac{\Lambda_{\mathcal{H}}(f) \eta^2}{2} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 \\ &= f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - \eta \langle \mathbf{g}_t, \mathbf{V}_t^{-1} \mathbf{g}_t \rangle + \eta \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle + \frac{\Lambda_{\mathcal{H}}(f) \eta^2}{2} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 \end{aligned}$$

where we plug in the update rule for  $\mathbf{x}_{t+1}$  in Algorithm 2. Summing over  $t$  yields

$$\begin{aligned} \sum_{t=0}^{T-1} (f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)) &\leq \sum_{t=0}^{T-1} (f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)) - \eta \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2 \\ &\quad + \eta \sum_{t=0}^{T-1} \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle + \frac{\Lambda_{\mathcal{H}}(f) \eta^2}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2. \end{aligned}$$

This completes the proof.  $\square$

The following lemma is a standard result for mirror descent.

**Lemma E.4** (Lemma 5, Gupta et al. 2017). *For any  $\mathbf{x}_0 \in \mathcal{X}$ ,  $\mathbf{g} \in \mathbb{R}^d$  and  $\mathbf{H} \succeq 0$ , let  $\mathbf{x}_1 = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{g}, \mathbf{x}_0 \rangle + \|\mathbf{x} - \mathbf{x}_0\|_{\mathbf{H}}^2$ . Then it holds that*

$$\langle \mathbf{g}, \mathbf{x}_0 - \mathbf{x}_1 \rangle \leq \frac{1}{2} \|\mathbf{H}^{-1} \mathbf{g}\|_{\mathbf{H}}^2 + \frac{1}{2} \|\mathbf{x}_0 - \mathbf{x}_1\|_{\mathbf{H}}^2.$$

Moreover, for any  $\mathbf{x} \in \mathcal{X}$ , it holds that

$$\begin{aligned} \langle \mathbf{g}, \mathbf{x}_1 - \mathbf{x} \rangle &\leq \frac{1}{2} \|\mathbf{x}_0 - \mathbf{x}\|_{\mathbf{H}}^2 - \frac{1}{2} \|\mathbf{x}_1 - \mathbf{x}\|_{\mathbf{H}}^2 - \frac{1}{2} \|\mathbf{x}_0 - \mathbf{x}_1\|_{\mathbf{H}}^2, \\ \langle \mathbf{g}, \mathbf{x}_0 - \mathbf{x} \rangle &\leq \frac{1}{2} \|\mathbf{x}_0 - \mathbf{x}\|_{\mathbf{H}}^2 - \frac{1}{2} \|\mathbf{x}_1 - \mathbf{x}\|_{\mathbf{H}}^2 + \frac{1}{2} \|\mathbf{H}^{-1} \mathbf{g}\|_{\mathbf{H}}^2. \end{aligned}$$

## E.2 STOCHASTIC CASE WITH PROJECTION

---

### Algorithm 8 Accelerated Adaptive Algorithm with Projection

---

**Hyperparam:**  $\epsilon \geq 0$ , total steps  $T$ , learning rate  $\eta$ , convex cone  $\mathcal{H} \subset \mathcal{S}_+$ , radius  $D$ ?

**Input:** initialization  $\mathbf{x}_0 = \bar{\mathbf{x}}_0$ , a sequence of positive constants  $\alpha_0, \dots, \alpha_T \in (0, 1]$

```

 $M_{-1} \leftarrow \mathbf{0}$ 
for  $t = 0, 1, 2, \dots, T$  :
   $\mathbf{g}_t \leftarrow \nabla f_t^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t)$  where  $f_t^{\alpha_t, \bar{\mathbf{x}}_t}$  is defined in (8)
   $M_t \leftarrow M_{t-1} + \mathbf{g}_t \mathbf{g}_t^\top$ 
   $V_t \leftarrow \arg \min_{\mathbf{H} \in \mathcal{H}} \langle M_t + \epsilon \mathbf{I}_d, \mathbf{H}^{-1} \rangle + \text{Tr}(\mathbf{H})$ 
   $\mathbf{x}_{t+1/2} \leftarrow \mathbf{x}_t - \eta V_t^{-1} \mathbf{g}_t$ 
   $\mathbf{x}_{t+1} \leftarrow \arg \min_{\|\mathbf{x}\|_{\mathcal{H}} \leq D} \|\mathbf{x} - \mathbf{x}_{t+1/2}\|_{V_t}^2$ 
   $\bar{\mathbf{x}}_{t+1} \leftarrow \alpha_t \mathbf{x}_{t+1/2} + (1 - \alpha_t) \bar{\mathbf{x}}_t$ 
return  $\bar{\mathbf{x}}_T$ 

```

---

Theorem E.5 shows the convergence rate of Algorithm 8, which is the same as Algorithm 2 while removing the dependence on prior knowledge of  $D$ .

**Theorem E.5.** *For a well-structured preconditioner set  $\mathcal{H}$ , let  $f$  be a convex loss function whose  $\mathcal{H}$ -smoothness constant is  $\Lambda_{\mathcal{H}}(f) \in (0, \infty)$  according to Definition 2.4. For  $\epsilon > 0$ ,  $T > 0$ , consider Algorithm 8 with  $\alpha_t = 2/(t+2)$  for  $t = 0, 1, \dots, T-1$ . Suppose the global minima  $\mathbf{x}^*$  satisfies that  $\|\mathbf{x}^*\|_{\mathcal{H}} \leq D$  and Assumption 4.2 holds with adaptive gradient variance  $\sigma_{\mathcal{H}}(\{f_t\}_{t=1}^T)^2 \leq \sigma_{\mathcal{H}}^2$  for some  $\sigma_{\mathcal{H}} \in [0, \infty)$ . Then it holds that*

$$\begin{aligned} \mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] &\leq \frac{2D^2\epsilon}{\eta(T+1)^2} \mathbb{E} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left( \frac{D^2}{2\eta} - \frac{\eta}{2} \right) \frac{4}{(T+1)^2} \mathbb{E} \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \\ &\quad + \frac{2\eta^2}{(T+1)^2} \cdot \Lambda_{\mathcal{H}}(f) \cdot \tilde{O}(\log^2 d) + \frac{\eta}{T^{1/2}} \sigma_{\mathcal{H}} \cdot \tilde{O}(\log d). \end{aligned}$$

Moreover, when choosing learning rate  $\eta = D$ , the convergence rate becomes

$$\mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] = \tilde{O} \left( \frac{\Lambda_{\mathcal{H}}(f) D^2 \log^2 d + d\sqrt{\epsilon}D}{T^2} + \frac{\sigma_{\mathcal{H}} D \log d}{\sqrt{T}} \right).$$

*Proof of Theorem E.5.* Combining the results of Lemma E.7 and Lemma E.8, we get

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{E}[f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1/2}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] &\leq \mathbb{E} \left[ \frac{D^2\epsilon}{2\eta} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left( \frac{D^2}{2\eta} - \frac{\eta}{2} \right) \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \right] \\ &\quad + \mathbb{E} \left[ \eta \sum_{t=0}^{T-1} \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle \right] \\ &\quad + \mathbb{E} \left[ \frac{\Lambda_{\mathcal{H}}(f) \eta^2}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2 \right]. \end{aligned}$$

The rest of the proof will be the same as the proof of Theorem 4.3.  $\square$

**Lemma E.6.** *Under the setting of Theorem E.5, it holds that*

$$\mathbb{E}[f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*)] \leq \frac{4}{(T+1)^2} \sum_{t=0}^{T-1} \mathbb{E} [f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1/2}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)] \quad (45)$$

*Proof of Lemma E.6.* The proof is the same as the proof of Lemma E.1 with  $\mathbf{x}_{t+1/2}$  in place of  $\mathbf{x}_{t+1}$ .  $\square$

Define  $\mathbf{V}_{-1} = 0$  for convenience.

**Lemma E.7.** *Under the setting of Theorem E.5, it holds that*

$$\sum_{t=0}^{T-1} \mathbb{E} \left[ f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*) \right] \leq \mathbb{E} \left[ \frac{D^2 \epsilon}{2\eta} \text{Tr}(\mathbf{V}_{T-1}^{-1}) + \left( \frac{D^2}{2\eta} + \frac{\eta}{2} \right) \sum_{t=0}^{T-1} \mathbf{g}_t^\top \mathbf{V}_t^{-1} \mathbf{g}_t \right].$$

*Proof of Lemma E.7.* The proof is the same as the proof of Lemma E.2, except that controlling  $\langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle$  requires some extra work. Specifically, by definition  $\mathbf{x}_{t+1/2} = \mathbf{x}_t - \eta \mathbf{V}_t^{-1} \mathbf{g}_t$ , so it follows from Lemma E.4 that

$$\langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle \leq \frac{1}{2\eta} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_t}^2 - \frac{1}{2\eta} \|\mathbf{x}_{t+1/2} - \mathbf{x}^*\|_{\mathbf{V}_t}^2 + \frac{\eta}{2} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2.$$

Since  $\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}, \|\mathbf{x}\|_{\mathcal{H}} \leq D} \|\mathbf{x} - \mathbf{x}_{t+1/2}\|_{\mathbf{V}_t}^2$  and  $\|\mathbf{x}^*\|_{\mathcal{H}} \leq D$ , it holds that  $\frac{d}{d\lambda} \|\mathbf{x}_{t+1} + \lambda(\mathbf{x}^* - \mathbf{x}_{t+1}) - \mathbf{x}_{t+1/2}\|_{\mathbf{V}_t}^2 \Big|_{\lambda=0} \geq 0$ , which implies that  $(\mathbf{x}^* - \mathbf{x}_{t+1})^\top \mathbf{V}_t (\mathbf{x}_{t+1} - \mathbf{x}_{t+1/2}) \geq 0$ . Then it follows that  $\|\mathbf{x}_{t+1/2} - \mathbf{x}^*\|_{\mathbf{V}_t}^2 = \|\mathbf{x}_{t+1} - \mathbf{x}^*\|_{\mathbf{V}_t}^2 + 2(\mathbf{x}^* - \mathbf{x}_{t+1})^\top \mathbf{V}_t (\mathbf{x}_{t+1} - \mathbf{x}_{t+1/2}) + \|\mathbf{x}_{t+1} - \mathbf{x}_{t+1/2}\|_{\mathbf{V}_t}^2 \geq \|\mathbf{x}_{t+1} - \mathbf{x}^*\|_{\mathbf{V}_t}^2$ . Therefore, applying this to the above inequality yields

$$\langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle \leq \frac{1}{2\eta} \|\mathbf{x}_t - \mathbf{x}^*\|_{\mathbf{V}_t}^2 - \frac{1}{2\eta} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|_{\mathbf{V}_t}^2 + \frac{\eta}{2} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2.$$

Then the rest of the proof is the same.  $\square$

**Lemma E.8.** *Under the setting of Theorem E.5, it holds that*

$$\begin{aligned} \sum_{t=0}^{T-1} (f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_{t+1/2}) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)) &\leq \sum_{t=0}^{T-1} (f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t) - f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}^*)) - \eta \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{V}_t}^2 \\ &\quad + \eta \sum_{t=0}^{T-1} \langle \mathbf{g}_t - \nabla f^{\alpha_t, \bar{\mathbf{x}}_t}(\mathbf{x}_t), \mathbf{V}_t^{-1} \mathbf{g}_t \rangle + \frac{\eta^2}{2} \sum_{t=0}^{T-1} \|\mathbf{V}_t^{-1} \mathbf{g}_t\|_{\mathbf{H}^*}^2. \end{aligned}$$

*Proof of Lemma E.8.* The proof is the same as the proof of Lemma E.3 with  $\mathbf{x}_{t+1/2}$  in place of  $\mathbf{x}_{t+1}$ .  $\square$

## F PROOF FOR NORMALIZED STEEPEST DESCENT

Suppose that  $\|\cdot\|$  is a norm on  $\mathbb{R}^d$  and  $\|\cdot\|_*$  is its dual norm. The following descent lemma characterizes the progress of one step of steepest descent under  $\|\cdot\|$ .

**Lemma F.1** (Descent lemma). *Let  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  be a differentiable function. For any  $\mathbf{m} \in \mathbb{R}^d$ , let  $\mathbf{u} = \arg \max_{\|\mathbf{v}\| \leq 1} \langle \mathbf{m}, \mathbf{v} \rangle$ . Then for any  $\eta > 0$ , it holds that*

$$f(\mathbf{x} - \eta \mathbf{u}) \leq f(\mathbf{x}) - \eta \|\nabla f(\mathbf{x})\|_* + 2\eta \|\nabla f(\mathbf{x}) - \mathbf{m}\|_* + \frac{\eta^2}{2} L_{\|\cdot\|}(f).$$

*Proof of Lemma F.1.* By definition of the smoothness in Definition 2.3, we have

$$\begin{aligned} f(\mathbf{x} - \eta \mathbf{u}) &\leq f(\mathbf{x}) - \nabla f(\mathbf{x})^\top (\eta \mathbf{u}) + \frac{\|\eta \mathbf{u}\|^2}{2} L_{\|\cdot\|}(f) \\ &= f(\mathbf{x}) - \eta \langle \nabla f(\mathbf{x}) - \mathbf{m}, \mathbf{u} \rangle - \eta \langle \mathbf{m}, \mathbf{u} \rangle + \frac{\eta^2}{2} L_{\|\cdot\|}(f) \\ &\leq f(\mathbf{x}) + \eta \|\nabla f(\mathbf{x}) - \mathbf{m}\|_* - \eta \|\mathbf{m}\|_* + \frac{\eta^2}{2} L_{\|\cdot\|}(f) \end{aligned}$$

where in the last step we use  $\|\mathbf{u}\| = 1$  and  $\langle \mathbf{m}, \mathbf{u} \rangle = \|\mathbf{m}\|_*$  by the definition of  $\mathbf{u}$ . Then we further apply the triangle inequality  $\|\mathbf{m}\|_* \geq \|\nabla f(\mathbf{x})\|_* - \|\nabla f(\mathbf{x}) - \mathbf{m}\|_*$  to get

$$f(\mathbf{x} - \eta \mathbf{u}) \leq f(\mathbf{x}) - \eta \|\nabla f(\mathbf{x})\|_* + 2\eta \|\nabla f(\mathbf{x}) - \mathbf{m}\|_* + \frac{\eta^2}{2} L_{\|\cdot\|}(f).$$

This concludes the proof.  $\square$

**Proposition F.2.** For any  $\epsilon \geq 0$ ,  $\alpha \in (0, 1)$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 3 with  $\|\cdot\|$ . Let  $L_{\|\cdot\|}(f)$  be the smoothness of the loss  $f$  w.r.t.  $\|\cdot\|$  according to Definition 2.3. Suppose Assumption 4.2 holds. Then it holds that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_* &\leq \frac{\Delta_0}{\eta T} + \frac{L_{\|\cdot\|}(f)\eta}{2} + \frac{2}{\alpha T} \mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_* + \frac{2(1-\alpha)\eta}{\alpha} L_{\|\cdot\|}(f) \\ &\quad + \frac{2\alpha}{T} \sum_{t=0}^{T-1} \mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_* \end{aligned}$$

Proposition F.2. Apply Lemma F.1 with  $\|\cdot\|$  yields that

$$f(\mathbf{x}_{t+1}) - f(\mathbf{x}_t) \leq -\eta \|\nabla f(\mathbf{x}_t)\|_* + \frac{L_{\|\cdot\|}(f)\eta^2}{2} + 2\eta \|\nabla f(\mathbf{x}_t) - \mathbf{m}_t\|_*. \quad (46)$$

Rearranging the telescoping sum of Eq. 46 from  $t = 0$  to  $T - 1$ , we have that

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_* \leq \frac{1}{\eta T} [f(\mathbf{x}_0) - \min f(\mathbf{x})] + \frac{L_{\|\cdot\|}(f)\eta}{2} + \frac{2}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t) - \mathbf{m}_t\|_*. \quad (47)$$

By the update rule of Algorithm 3, we have  $\mathbf{m}_t = (1 - \alpha)\mathbf{m}_{t-1} + \alpha\mathbf{g}_t$ , so

$$\mathbf{m}_t - \nabla f(\mathbf{x}_t) = (1 - \alpha)(\mathbf{m}_{t-1} - \nabla f(\mathbf{x}_{t-1})) + (1 - \alpha)(\nabla f(\mathbf{x}_{t-1}) - \nabla f(\mathbf{x}_t)) + \alpha(\mathbf{g}_t - \nabla f(\mathbf{x}_t)).$$

Unrolling the above recursion, we get that

$$\begin{aligned} \mathbf{m}_t - \nabla f(\mathbf{x}_t) &= (1 - \alpha)^t (\mathbf{m}_0 - \nabla f(\mathbf{x}_0)) \\ &\quad + \sum_{i=1}^t (1 - \alpha)^{t-i} [(1 - \alpha)(\nabla f(\mathbf{x}_{i-1}) - \nabla f(\mathbf{x}_i)) + \alpha(\mathbf{g}_i - \nabla f(\mathbf{x}_i))]. \end{aligned}$$

and

$$\begin{aligned} \mathbb{E} \|\nabla f(\mathbf{x}_t) - \mathbf{m}_t\|_* &\leq (1 - \alpha)^t \mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_* + \sum_{i=1}^t (1 - \alpha)^{t-i+1} \|\nabla f(\mathbf{x}_{i-1}) - \nabla f(\mathbf{x}_i)\|_* \\ &\quad + \alpha \mathbb{E} \left\| \sum_{i=1}^t (1 - \alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_* \end{aligned} \quad (48)$$

We know that

$$\|\nabla f(\mathbf{x}_{i-1}) - \nabla f(\mathbf{x}_i)\|_* \leq L_{\|\cdot\|}(f) \|\mathbf{x}_{i-1} - \mathbf{x}_i\| \leq \eta L_{\|\cdot\|}(f). \quad (49)$$

Then plugging (49) into (48), we can get that

$$\begin{aligned} \mathbb{E} \|\nabla f(\mathbf{x}_t) - \mathbf{m}_t\|_* &\leq (1 - \alpha)^t \mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_* + (1 - \alpha)\eta L_{\|\cdot\|}(f) \sum_{i=1}^t (1 - \alpha)^{t-i} \\ &\quad + \alpha \mathbb{E} \left\| \sum_{i=1}^t (1 - \alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_* \\ &\leq (1 - \alpha)^t \mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_* + \frac{1 - \alpha}{\alpha} \eta L_{\|\cdot\|}(f) \\ &\quad + \alpha \mathbb{E} \left\| \sum_{i=1}^t (1 - \alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_* \end{aligned}$$

Plugging it in Eq. 47, we have that

$$\begin{aligned}
\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_* &\leq \frac{\Delta_0}{\eta T} + \frac{L_{\|\cdot\|}(f)\eta}{2} + \frac{2}{T} \sum_{t=0}^{T-1} [(1-\alpha)^t \mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_*] \\
&\quad + \frac{2}{T} \sum_{t=0}^{T-1} \left[ \frac{1-\alpha}{\alpha} \eta L_{\|\cdot\|}(f) + \alpha \mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_* \right] \\
&= \frac{\Delta_0}{\eta T} + \frac{L_{\|\cdot\|}(f)\eta}{2} + \frac{2}{\alpha T} \mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_* + \frac{2(1-\alpha)\eta}{\alpha} L_{\|\cdot\|}(f) \\
&\quad + \frac{2\alpha}{T} \sum_{t=0}^{T-1} \mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_*
\end{aligned}$$

□

### F.1 CONVERGENCE RATE WITH RESPECT TO ADAPTIVE GRADIENT VARIANCE

The proof of Theorem 4.5 largely borrows from Kovalev (2025b), upon which we change the dependence on adaptive noise and improve some coefficient constant.

**Theorem 4.5.** *Let  $\mathcal{H}$  be a well-structured preconditioner set. For any  $\epsilon \geq 0$ ,  $\alpha \in (0, 1)$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 3 with  $\|\cdot\| = \|\cdot\|_{\mathcal{H}}$  and  $\mathbf{m}_0 = \nabla f_0(\mathbf{x}_0)$ . Let  $L_{\|\cdot\|_{\mathcal{H}}}(f)$  be the smoothness of the loss  $f$  w.r.t.  $\|\cdot\|_{\mathcal{H}}$  according to Definition 2.3. Suppose Assumption 4.2 holds with adaptive gradient variance  $\sigma_{\mathcal{H}}(\{f_t\}_{t=1}^T)^2 \leq \sigma_{\mathcal{H}}^2$ . Then it holds that*

$$\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} \leq \frac{\Delta_0}{\eta T} + \frac{2\eta}{\alpha} L_{\|\cdot\|}(f) + \frac{2\sigma_{\mathcal{H}}}{\alpha T} + 2\sigma_{\mathcal{H}}\sqrt{\alpha}.$$

Let  $a_0 = \sqrt{\Delta_0 L_{\|\cdot\|}(f) / \sigma_{\mathcal{H}}}$ . If  $a_0 < 1$ , then

- When  $T < a_0^{-6}$ , we choose  $\alpha = T^{-2/3}$  and  $\eta = \sqrt{\frac{\Delta_0}{L_{\|\cdot\|_{\mathcal{H}}}(f)}} T^{-5/12}$ . The rate is  $O(\sigma_{\mathcal{H}} T^{-1/3})$ .
- When  $T \geq a_0^{-6}$ , we choose  $\alpha = \frac{a_0}{\sqrt{T}}$ ,  $\eta = \frac{\Delta_0^{3/4} T^{-3/4}}{L_{\|\cdot\|_{\mathcal{H}}}(f)^{1/4} \sigma_{\mathcal{H}}^{1/2}}$ . The rate is  $O\left(\frac{(\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f))^{1/4} \sqrt{\sigma_{\mathcal{H}}}}{T^{1/4}}\right)$ .

If  $a_0 \geq 1$ , then

- When  $T \leq a_0^2$ , we choose  $\alpha = 1$  and  $\eta = \sqrt{\frac{\Delta_0}{L_{\|\cdot\|_{\mathcal{H}}}(f)}} T^{-1/2}$ . The rate is  $O\left(\sqrt{\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f)} T^{-1/2}\right)$ .
- When  $T \geq a_0^2$ , we choose  $\alpha = \frac{a_0}{\sqrt{T}}$  and  $\eta = \frac{\Delta_0^{3/4} T^{-3/4}}{L_{\|\cdot\|_{\mathcal{H}}}(f)^{1/4} \sigma_{\mathcal{H}}^{1/2}}$ . The rate is  $O\left(\frac{(\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f))^{1/4} \sqrt{\sigma_{\mathcal{H}}}}{T^{1/4}}\right)$ .

*Proof of Theorem 4.5.* By applying Proposition F.2 with  $\|\cdot\| = \|\cdot\|_{\mathcal{H}}$ , we have that

$$\begin{aligned}
\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} &\leq \frac{\Delta_0}{\eta T} + \frac{L_{\|\cdot\|}(f)\eta}{2} + \frac{2}{\alpha T} \mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_{\mathcal{H},*} + \frac{2(1-\alpha)\eta}{\alpha} L_{\|\cdot\|}(f) \\
&\quad + \frac{2}{T} \sum_{t=0}^{T-1} \mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} \alpha (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_{\mathcal{H},*}. \tag{50}
\end{aligned}$$

We only need to bound  $\mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_{\mathcal{H},*}$  and  $\mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} \alpha (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_{\mathcal{H},*}$ . For notational simplicity, we define  $\omega_i = (1-\alpha)^{t-i} \alpha$  and  $\delta_i = \mathbf{g}_i - \nabla f(\mathbf{x}_i)$ . By Assumption 4.2, we

know that  $\mathbb{E}[\delta_j | \delta_i] = 0$  for  $j > 1$ . Then we have that

$$\begin{aligned}
\mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} \alpha (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_{\mathcal{H},*}^2 &= \mathbb{E} \left\| \sum_{i=1}^t \omega_i \delta_i \right\|_{\mathcal{H},*}^2 \\
&= \mathbb{E} \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \left( \sum_{i=1}^t \omega_i \delta_i \right) \mathbf{H}^{-1} \left( \sum_{i=1}^t \omega_i \delta_i \right) \\
&\leq \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \mathbb{E} \left( \sum_{i=1}^t \omega_i \delta_i \right) \mathbf{H}^{-1} \left( \sum_{i=1}^t \omega_i \delta_i \right) \\
&= \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \mathbb{E} \sum_{i=1}^t \omega_i^2 \delta_i^\top \mathbf{H}^{-1} \delta_i \\
&= \min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \mathbb{E} \sum_{i=1}^t \omega_i^2 \|\delta_i\|_{\mathbf{H}^{-1}}^2.
\end{aligned}$$

Now by the condition on  $\sigma_{\mathcal{H}}$  and the definition of the adaptive gradient variance in Definition 4.1, we know that  $\min_{\mathbf{H} \in \mathcal{H}, \text{Tr}(\mathbf{H}) \leq 1} \mathbb{E} \|\delta_i\|_{\mathbf{H}^{-1}}^2 \leq \sigma_{\mathcal{H}}^2$  for each  $i \in [t]$ . Therefore, we further have

$$\mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} \alpha (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_{\mathcal{H},*}^2 \leq \sigma_{\mathcal{H}}^2 \sum_{i=1}^t \omega_i^2 \leq \sigma_{\mathcal{H}}^2 \alpha \quad (51)$$

where the second inequality holds by the definition of  $\omega_i$  and the condition that  $\alpha \in (0, 1)$ . Similarly, we also have

$$\mathbb{E} \|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_{\mathcal{H},*}^2 = \mathbb{E} \|\mathbf{g}_0 - \nabla f(\mathbf{x}_0)\|_{\mathcal{H},*}^2 \leq \sigma_{\mathcal{H}}^2. \quad (52)$$

Plugging in Eq. 51 and Eq. 52 into Eq. 50, we can get

$$\begin{aligned}
\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} &\leq \frac{\Delta_0}{\eta T} + \frac{\eta L_{\|\cdot\|_{\mathcal{H}}}(f)}{2} + \frac{2\sigma_{\mathcal{H}}}{\alpha T} + \frac{2(1-\alpha)\eta}{\alpha} L_{\|\cdot\|_{\mathcal{H}}}(f) + 2\sqrt{\alpha}\sigma_{\mathcal{H}} \\
&\leq \frac{\Delta_0}{\eta T} + \frac{2\eta}{\alpha} L_{\|\cdot\|_{\mathcal{H}}}(f) + \frac{2\sigma_{\mathcal{H}}}{\alpha T} + 2\sqrt{\alpha}\sigma_{\mathcal{H}}.
\end{aligned} \quad (53)$$

Next, we will choose appropriate hyperparameters  $\eta$  and  $\alpha$  to achieve the optimal rate.

For any fixed  $\alpha < 1$ , to balance the first two terms on the right-hand side of (53), we choose

$$\eta = \sqrt{\frac{\Delta_0 \alpha}{2T L_{\|\cdot\|_{\mathcal{H}}}(f)}}.$$

Also, for simplicity, denote  $a_0 = \sqrt{\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f)}/\sigma_{\mathcal{H}}$ . Then the convergence rate becomes

$$\begin{aligned}
\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_{\mathcal{H},*} &\leq 2\sqrt{\frac{2\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f)}{\alpha T}} + \frac{2\sigma_{\mathcal{H}}}{\alpha T} + 2\sqrt{\alpha}\sigma_{\mathcal{H}} \\
&\leq 2\sqrt{2}\sigma_{\mathcal{H}}(f) \left( \frac{a_0}{\sqrt{\alpha T}} + \frac{1}{\alpha T} + \sqrt{\alpha} \right).
\end{aligned} \quad (54)$$

We further stratify the values of  $\alpha$ ,  $a_0$  and  $T$  to minimize the right-hand side:

- When  $\alpha \leq T^{-2/3}$ ,  $\frac{1}{\alpha T} \geq \sqrt{\alpha}$ . We only need to minimize  $\frac{1}{\alpha T} + \frac{a_0}{\sqrt{\alpha} \sqrt{T}}$ , which is achieved at  $\alpha = T^{-2/3}$  because it is a monotone decreasing function in  $\alpha$ . So the convergence rate is always no better than that in the case of  $\alpha \geq T^{-2/3}$ .
- When  $\alpha \geq T^{-2/3}$ ,  $\frac{1}{\alpha T} \leq \sqrt{\alpha}$ . We only need to minimize  $\sqrt{\alpha} + \frac{a_0}{\sqrt{\alpha T}}$ . Then  $\frac{a_0}{\sqrt{T}}$  is the global minimizer but there are constraints  $T^{-2/3} \leq \alpha \leq 1$ .

- When  $a_0 \leq T^{-1/6}$ , we have that  $\frac{a_0}{\sqrt{T}} \leq T^{-2/3}$ . Then we need to choose  $\alpha = T^{-2/3}$  and the rate is  $\sigma_{\mathcal{H}}T^{-1/3} + \sqrt{\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f)}T^{-1/6} \leq 2\sigma_{\mathcal{H}}T^{-1/3}$ .
- When  $a_0 \geq \sqrt{T}$ , we have that  $\frac{a_0}{\sqrt{T}} \geq 1$ . Then we need to choose  $\alpha = 1$  and the rate is  $\sigma_{\mathcal{H}} + \sqrt{\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f)}T^{-1/2} \leq 2\sqrt{\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f)}T^{-1/2}$ .
- When  $T^{-1/6} \leq a_0 \leq \sqrt{T}$ , we have that  $T^{-2/3} \leq \frac{a_0}{\sqrt{T}} \leq 1$ . We can choose  $\alpha = \frac{a_0}{\sqrt{T}}$ . The rate becomes  $(\Delta_0 L_{\|\cdot\|_{\mathcal{H}}}(f))^{1/4} \sqrt{\sigma_{\mathcal{H}}}T^{-1/4}$ .

This completes the proof.  $\square$

## F.2 CONVERGENCE RATE UNDER GENERAL-NORM ASSUMPTION

**Definition F.3** (Distortion of norms). *For any target norm  $\|\cdot\|$  and a reference norm  $\|\cdot\|_{\text{ref}}$  on a finite-dimensional vector space, the corresponding distortion of  $\|\cdot\|$  with respect to  $\|\cdot\|_{\text{ref}}$  is defined as*

$$\psi(\|\cdot\|, \|\cdot\|_{\text{ref}}) := \sup_{\mathbf{x} \neq 0} \frac{\|\mathbf{x}\|}{\|\mathbf{x}\|_{\text{ref}}} \cdot \sup_{\mathbf{x} \neq 0} \frac{\|\mathbf{x}\|_{\text{ref}}}{\|\mathbf{x}\|}. \quad (55)$$

Since all norms on the finite-dimensional space are equivalent,  $\psi(\|\cdot\|, \|\cdot\|_{\text{ref}})$  is always finite.

**Theorem 4.6.** *For any  $\epsilon \geq 0$ ,  $\alpha \in (0, 1)$ ,  $\eta > 0$ , and  $T \in \mathbb{N}$ , let  $\{\mathbf{x}_t\}_{t=0}^T$  be the iterates of Algorithm 3 with any norm  $\|\cdot\|$ . Suppose Assumption 4.2 holds with the gradient variance  $\sigma_{\|\cdot\|_*}(\{f_t\}_{t=1}^T)^2 \leq \sigma_{\|\cdot\|_*}^2$  for some  $\sigma_{\|\cdot\|_*} \in [0, \infty)$ . Then it holds that*

$$\mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_* \leq \frac{\Delta_0}{\eta T} + \frac{2\eta}{\alpha} L_{\|\cdot\|}(f) + \frac{2}{\alpha T} \sigma_{\|\cdot\|_*} + 2\sigma_{\|\cdot\|_*} \cdot \min\left(1, \alpha^{1/2} \psi(\|\cdot\|_*, \|\cdot\|_2)\right)$$

where  $\psi(\|\cdot\|_*, \|\cdot\|_2) = \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_*}{\|\mathbf{x}\|_2} \cdot \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_2}{\|\mathbf{x}\|_*}$  measures the distortion between the two norms.

*Proof of Theorem 4.6.* Since  $\mathbb{E}[\|\mathbf{m}_0 - \nabla f(\mathbf{x}_0)\|_*] \leq \sigma_{\|\cdot\|_*}$ , it follows from Proposition F.2 that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_* &\leq \frac{\Delta_0}{\eta T} + \frac{L_{\|\cdot\|}(f)\eta}{2} + \frac{2\sigma_{\|\cdot\|_*}}{\alpha T} + \frac{2(1-\alpha)\eta}{\alpha} L_{\|\cdot\|}(f) \\ &\quad + \frac{2\alpha}{T} \sum_{t=0}^{T-1} \mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_*. \end{aligned} \quad (56)$$

There are two ways to control the last terms, and we discuss them separately below.

For the first approach, we directly apply triangle inequality to get

$$\begin{aligned} \alpha \cdot \mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_* &\leq \mathbb{E} \sum_{i=1}^t \alpha (1-\alpha)^{t-i} \|\mathbf{g}_i - \nabla f(\mathbf{x}_i)\|_* \\ &\leq \sum_{i=1}^t \alpha (1-\alpha)^{t-i} [\mathbb{E} \|\mathbf{g}_i - \nabla f(\mathbf{x}_i)\|_*^2]^{1/2} \leq \sigma_{\|\cdot\|_*}. \end{aligned} \quad (57)$$

For the second approach, we first convert  $\|\cdot\|_*$  to  $\|\cdot\|_2$  and get

$$\begin{aligned} \mathbb{E} \left[ \left\| \alpha \sum_{i=1}^t (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_*^2 \right] &\leq \left( \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_*}{\|\mathbf{x}\|_2} \right)^2 \cdot \mathbb{E} \left[ \left\| \sum_{i=1}^t \alpha (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_2^2 \right] \\ &= \left( \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_*}{\|\mathbf{x}\|_2} \right)^2 \cdot \sum_{i=1}^t \alpha^2 (1-\alpha)^{2(t-i)} \cdot \mathbb{E} [\|\mathbf{g}_i - \nabla f(\mathbf{x}_i)\|_2^2] \\ &\leq \left( \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_*}{\|\mathbf{x}\|_2} \right)^2 \left( \sup_{\mathbf{x}} \frac{\|\mathbf{x}\|_2}{\|\mathbf{x}\|_*} \right)^2 \cdot \sum_{i=0}^t (1-\alpha)^{2i} \alpha^2 \mathbb{E} [\|\mathbf{g}_i - \nabla f(\mathbf{x}_i)\|_*^2] \\ &\leq \psi(\|\cdot\|_*, \|\cdot\|_2)^2 \cdot \sigma_{\|\cdot\|_*}^2 \cdot \alpha. \end{aligned}$$

Using Jensen's inequality, this further implies that

$$\alpha \cdot \mathbb{E} \left\| \sum_{i=1}^t (1-\alpha)^{t-i} (\mathbf{g}_i - \nabla f(\mathbf{x}_i)) \right\|_* \leq \alpha^{1/2} \psi(\|\cdot\|_*, \|\cdot\|_2) \cdot \sigma_{\|\cdot\|_*}. \quad (58)$$

Now, by taking the minimum over the two bounds in (57) and (58), we obtain from (56) that

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_* &\leq \frac{\Delta_0}{\eta T} + \frac{L_{\|\cdot\|}(f)\eta}{2} + \frac{2\sigma_{\|\cdot\|_*}}{\alpha T} + \frac{2(1-\alpha)\eta}{\alpha} L_{\|\cdot\|}(f) \\ &\quad + 2\sigma_{\|\cdot\|_*} \cdot \min\left(1, \alpha^{1/2} \psi(\|\cdot\|_*, \|\cdot\|_2)\right). \end{aligned}$$

This completes the proof.  $\square$

### F.3 DIMENSION DEPENDENT LOWER BOUND FOR NSD

In this subsection, we construct the example for the lower bound in Theorem 4.7.

**Lemma F.4.** Fix any  $\epsilon > 0$ . For any distinct points  $x_0, x_1, \dots, x_T \in \mathbb{R}$  where  $T + 1 \leq \frac{1}{4\epsilon^2}$ , there exists a function  $p : \mathbb{R} \rightarrow \mathbb{R}$ , such that its derivative  $p'$  is 1-Lipschitz,  $p(x_0) - \inf_x p(x) \leq 1$  and  $p'(x_t) = -\epsilon$  for  $t = 0, 1, \dots, T$ .

*Proof of Lemma F.4.* We use  $g(x)$  to denote  $p'(x)$  in the proof. For any set of points  $\{x_t\}_{t=0}^T \subseteq \mathbb{R}$ , we will construct  $g(x)$  such that  $g(x) = -\epsilon$  for  $x < x_0$  and  $g(x + \frac{1}{\epsilon}) = g(x)$  for  $x \geq x_0$ . We first explain what it requires to construct this function  $g$ : For such a  $g$ , because of the periodicity, we can always find  $x_0 = y_0 \leq y_1 \leq \dots \leq y_{T'} < x_0 + \frac{1}{\epsilon}$  for  $T' \leq T$  so that for any  $x_t \geq x_0$ , there is a  $y_s$  such that there exists  $n \in \mathbb{N}$  satisfying  $x_t = y_s + \frac{1}{\epsilon} \cdot n$  and therefore  $g(x_t) = g(y_s)$ . For any  $x_t < x_0$ ,  $g(x_t) = -\epsilon$  by definition. Also for  $x \leq x_0$ , it holds that  $p(x_0) \leq p(x)$  and  $g(x)$  is 1-Lipschitz. Then it suffices to construct  $g(x)$  in the interval  $[x_0, x_0 + \frac{1}{\epsilon}]$  to satisfy the following conditions

- $g(x)$  is 1-Lipschitz in  $[x_0, x_0 + \frac{1}{\epsilon}]$ .
- For any  $x > x_0$ , it can be written as  $x = x_0 + \delta + n\epsilon^{-1}$  for some  $\delta \in [0, \epsilon^{-1})$  and non-negative integer  $n$ , and thus  $p(x) - p(x_0) = \int_{x_0}^{x_0+\delta} g(x)dx + n \int_{x_0}^{x_0+\epsilon^{-1}} g(x)dx$ . Therefore, to ensure that  $p(x_0) - \inf_x p(x) \leq 1$ , we require  $p(x_0) - p(x_0 + \delta) = \int_{x_0}^{x_0+\delta} g(x)dx \geq -1$  for any  $0 \leq \delta \leq \frac{1}{\epsilon}$  and  $\int_{x_0}^{x_0+\frac{1}{\epsilon}} g(x)dx \geq 0$ .
- $g(y_s) = -\epsilon$  for  $s = 0, 1, \dots, T'$ .

We will further define  $y_{T'+1} = x_0 + \frac{1}{\epsilon}$  in the following construction.

For each  $t = 0, 1, \dots, T'$ , we construct  $g$  on the interval  $[y_t, y_{t+1}]$  as follows:

$$g(x) = \begin{cases} -\epsilon + (x - y_t), & x \in [y_t, (y_t + y_{t+1})/2]; \\ -\epsilon - (x - y_{t+1}), & x \in ((y_t + y_{t+1})/2, y_{t+1}]. \end{cases}$$

It is straightforward to see that  $g$  is 1-Lipschitz on  $[x_0, x_0 + \epsilon^{-1}]$  and  $g(y_t) = -\epsilon$  for  $t = 0, 1, \dots, T' + 1$ . Also, note that  $g(x) \geq -\epsilon$  for all  $x \in [x_0, x_0 + \epsilon^{-1}]$ . Thus, for any  $0 \leq \delta \leq \frac{1}{\epsilon}$ , it always holds that  $\int_{x_0}^{x_0+\delta} g(x)dx \geq -\epsilon\delta \geq -1$ . Then it only remains to show  $\int_{x_0}^{x_0+\frac{1}{\epsilon}} g(x)dx \geq 0$ .

By direct calculation, we have

$$\begin{aligned}
\int_{x_0}^{x_0 + \frac{1}{\epsilon}} g(x) dx &= \sum_{t=0}^{T'} \int_{y_t}^{y_{t+1}} g(x) dx = \sum_{t=0}^{T'} \left( -\epsilon(y_{t+1} - y_t) + \frac{(y_{t+1} - y_t)^2}{4} \right) \\
&= -1 + \frac{1}{4} \sum_{t=0}^{T'} (y_{t+1} - y_t)^2 \\
&\geq -1 + \frac{1}{4} \frac{(\sum_{t=0}^{T'} (y_{t+1} - y_t))^2}{T' + 1} \\
&= -1 + \frac{1}{4\epsilon^2(T' + 1)}
\end{aligned}$$

where the inequality holds by Cauchy-Schwarz inequality. Since  $T' + 1 \leq T + 1 \leq \frac{1}{4\epsilon^2}$ , it immediately follows that  $\int_{x_0}^{x_0 + \frac{1}{\epsilon}} g(x) dx \geq 0$ . This completes the proof.  $\square$

**Theorem 4.7.** *For any fixed  $\Delta_0, L, \sigma^2, d, T$ , learning rate  $\eta$ , and any averaging parameter  $\alpha$ , there exists a loss function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ , a sequence of stochastic iid loss functions  $f_0, f_1, \dots, f_{T-1}$  and an initialization  $\mathbf{x}_0$  satisfying the following conditions: (1)  $f(\mathbf{x}_0) - \inf_{\mathbf{x}} f(\mathbf{x}) = \Delta_0$  and  $L_{\|\cdot\|_\infty}(f) \leq L$ ; (2) For any  $\mathbf{x} \in \mathbb{R}^d$ , it holds that  $\mathbb{E}[\nabla f_t(\mathbf{x})] = \nabla f(\mathbf{x})$  and  $\mathbb{E}[\|\nabla f_t(\mathbf{x}) - \nabla f(\mathbf{x})\|_1^2] \leq \sigma^2$ .*

When running Algorithm 3 with  $\|\cdot\| = \|\cdot\|_\infty$ , learning rate  $\eta$ , averaging parameter  $\alpha$  and initialization  $\mathbf{x}_0 = \mathbf{0}$ , it holds that

$$\mathbb{E} \left[ \min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_1 \right] = \min \{ e^{-2} 5^{-\frac{1}{4}} (dL\Delta_0\sigma^2)^{\frac{1}{4}} T^{-\frac{1}{2}}, e^{-2} 5^{-\frac{1}{2}} \sigma \}$$

*Proof of Theorem 4.7.* Inspired by Lemma 1 in Chewi et al. (2023), we first rescale the original problem to a parameter-free scaling for simplicity. Indeed, for any  $f, \{f_t\}_{t=0}^{T-1}, \mathbf{x}_0$  satisfying the conditions (a) and (b), we can define  $h(\mathbf{x}) = \Delta_0^{-1} \cdot f(\sqrt{\Delta_0/L}\mathbf{x})$  and  $h_t(\mathbf{x}) = \Delta_0^{-1} \cdot f_t(\sqrt{\Delta_0/L}\mathbf{x})$ . Then it can be verified that  $h$  satisfies that  $h(\mathbf{x}_0) - \inf_{\mathbf{x}} h(\mathbf{x}) = 1$  and  $L_{\|\cdot\|_\infty}(h) \leq 1$ . Therefore, it suffices to construct  $h$  and associated stochastic loss functions  $h_0, h_1, \dots, h_{T-1}$  such that  $h(\mathbf{x}_0) - \inf_{\mathbf{x}} h(\mathbf{x}) = 1$ ,  $L_{\|\cdot\|_\infty}(h) \leq 1$  and  $\mathbb{E}[\|\nabla h_t(\mathbf{x}) - \nabla h(\mathbf{x})\|_1] \leq \sigma/\sqrt{L\Delta_0}$ , and then the construction is completed with the reverse transform  $f(\mathbf{x}) = \Delta_0 \cdot h(\sqrt{L/\Delta_0}\mathbf{x})$  and  $f_t(\mathbf{x}) = \Delta_0 \cdot h_t(\sqrt{L/\Delta_0}\mathbf{x})$ .

Correspondingly, we can rescale the learning rate and the iterates of Algorithm 3 on the loss function  $f$  to transform  $\{\mathbf{x}_t\}_{t=0}^T$  to be iterates obtained by running Algorithm 3 on the loss function  $h$ : It can be easily checked that  $\mathbf{x}_t = \sqrt{\Delta_0/L}\tilde{\mathbf{x}}_t$ , where  $\{\tilde{\mathbf{x}}_t\}_{t=0}^{T-1}$  are the iterates obtained by running Algorithm 3 on  $h_0, \dots, h_{T-1}$  with learning rate  $\eta\sqrt{L/\Delta_0}$  and momentum factor  $\alpha$ . Then if it holds that  $\mathbb{E}[\min_{t=0}^{T-1} \|\nabla h(\tilde{\mathbf{x}}_t)\|_1] = \min \{ e^{-2} 5^{-\frac{1}{4}} (d\sigma'^2)^{1/4} T^{-1/2}, e^{-2} 5^{-\frac{1}{2}} \sigma' \}$  where  $\sigma' = \sigma/\sqrt{L\Delta_0}$ , we immediately have

$$\begin{aligned}
\mathbb{E} \left[ \min_{t=0,1,\dots,T-1} \|\nabla f(\mathbf{x}_t)\|_1 \right] &= \mathbb{E} \left[ \min_{t=0,1,\dots,T-1} \sqrt{L\Delta_0} \cdot \|\nabla h(\tilde{\mathbf{x}}_t)\|_1 \right] \\
&= \min \{ e^{-2} 5^{-\frac{1}{4}} (dL\Delta_0\sigma^2)^{\frac{1}{4}} T^{-\frac{1}{2}}, e^{-2} 5^{-\frac{1}{2}} \sigma \}
\end{aligned}$$

which gives the desired result in Theorem 4.7. Given the above discussion, we only need to construct a loss function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  and associated stochastic loss functions  $f_0, f_1, \dots, f_{T-1}$  such that  $f(\mathbf{x}_0) - \inf_{\mathbf{x}} f(\mathbf{x}) = 1$ ,  $L_{\|\cdot\|_\infty}(f) \leq 1$  and  $\mathbb{E}[\|\nabla f_t(\mathbf{x}) - \nabla f(\mathbf{x})\|_1^2] \leq \sigma^2$  for all  $t = 0, \dots, T-1$ , and show that

$$\mathbb{E} \left[ \min_{t=0,\dots,T-1} \|\nabla f(\mathbf{x}_t)\|_1 \right] = \min \left( e^{-2} 5^{-\frac{1}{4}} \frac{d^{\frac{1}{4}} \sigma^{\frac{1}{2}}}{T^{\frac{1}{2}}}, e^{-2} 5^{-\frac{1}{2}} \sigma \right)$$

Below we present a construction for such a hard instance.

**Construction of loss functions.** We consider initialization  $\mathbf{x}_0 = \mathbf{0}$ . For convenience, we define the target quantity as

$$\epsilon = \min \left( \frac{d^{1/4} \sigma^{1/2}}{5^{1/4} T^{1/2}}, \frac{\sigma}{\sqrt{5}} \right). \quad (59)$$

We first construct a sequence of independently random noise vectors  $\boldsymbol{\delta}_0, \boldsymbol{\delta}_1, \dots, \boldsymbol{\delta}_{T-1}$  as follows. For constant  $C = \frac{\sigma^2}{5\epsilon}$  and constant  $\theta = \frac{5\epsilon^2}{d\sigma^2}$ , each random vector  $\boldsymbol{\delta}_t = -C[R_{t,1}, \dots, R_{t,d}]$  where iid variables  $R_{t,i} \sim \text{Bernoulli}(\theta)$ . The Bernoulli distribution is well-defined because  $\theta \leq \frac{1}{d} \leq 1$  by the choice of  $\epsilon$  in Eq. 59. Then we know that  $\mathbb{E}[\boldsymbol{\delta}_t] = -C\theta \mathbf{1}_d = -\frac{\epsilon}{d} \mathbf{1}_d$ ,

$$\mathbb{E}[\|\boldsymbol{\delta}_t\|_1] = d\mathbb{E}[|\delta_{t,1}|] = Cd\theta = \epsilon$$

and

$$\begin{aligned} \mathbb{E}[\|\boldsymbol{\delta}_t - \mathbb{E}[\boldsymbol{\delta}_t]\|_1^2] &= C^2 \mathbb{E} \left[ \left( \sum_{i=1}^d |R_{t,i} - \theta| \right)^2 \right] \\ &= C^2 \left[ d\mathbb{E}[|R_{t,1} - \theta|^2] + d(d-1)\mathbb{E}[|R_{t,1} - \theta|^2]^2 \right] \\ &= C^2 \left[ d[\theta(1-\theta)^2 + (1-\theta)\theta^2] + d(d-1)[2\theta(1-\theta)]^2 \right] \\ &= C^2 \left[ d\theta(1-\theta) + d(d-1)4\theta^2(1-\theta)^2 \right] \\ &\leq C^2 (d\theta + 4d^2\theta^2) \\ &\leq 5C^2 d\theta = \sigma^2 \end{aligned} \quad (60)$$

The last inequality holds because of  $d\theta \leq 1$  by the definition of  $\theta$ . And we can verify that  $\mathbb{E}[\|\boldsymbol{\delta}_t\|_1] = \epsilon$  and  $\mathbb{E}[\|\boldsymbol{\delta}_t - \mathbb{E}[\boldsymbol{\delta}_t]\|_1^2] \leq 5C^2 d\theta = \sigma^2$  by plugging in  $C = \frac{\sigma^2}{5\epsilon}$  and  $\theta = \frac{5\epsilon^2}{d\sigma^2}$ .

Next, we construct the loss function  $f$ . Consider the set of points  $S_{\eta, \epsilon} = \{0, \eta, 2\eta, \dots, (N-1)\eta\}$  with  $N = \lfloor \frac{1}{4\epsilon^2} \rfloor$ . We then invoke Lemma F.4 with set  $S_{\eta, \epsilon}$  and level  $-\epsilon$  to obtain a one-dimensional function  $p$ , such that its derivative  $p'$  is 1-Lipschitz,  $p(0) - \inf_x p(x) \leq 1$  and  $p'(k\eta) = -\epsilon$  for all  $k \leq N$ . Now we define  $f: \mathbb{R}^d \rightarrow \mathbb{R}$  as

$$f(\mathbf{x}) = \frac{1}{d} \sum_{i=1}^d p(x_i) \quad (61)$$

Leveraging the properties of  $p$ , we can show that for any  $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ ,

$$\begin{aligned} f(\mathbf{x}_0) - \inf_{\mathbf{x}} f(\mathbf{x}) &\leq \frac{1}{d} \sum_{i=1}^d (p(0) - \inf_x p(x)) \leq 1, \\ \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}')\|_1 &= \frac{1}{d} \sum_{i=1}^d |p'(x_i) - p'(x'_i)| \leq \|\mathbf{x} - \mathbf{x}'\|_\infty, \end{aligned}$$

where the last inequality holds because  $p'$  is 1-Lipschitz. This shows that  $f$  satisfies the desired condition (a). Then for each  $t = 0, 1, \dots, T-1$ , we define the stochastic loss function  $f_t$  as

$$f_t(\mathbf{x}) = f(\mathbf{x}) + \langle \boldsymbol{\delta}_t, \mathbf{x} \rangle + \frac{\epsilon}{d} \langle \mathbf{1}_d, \mathbf{x} \rangle. \quad (62)$$

Its gradient is given by

$$\nabla f_t(\mathbf{x}) = \nabla f(\mathbf{x}) + \boldsymbol{\delta}_t + \frac{\epsilon}{d} \mathbf{1}_d.$$

By the previous construction of  $\boldsymbol{\delta}_t$ , it is clear that  $\mathbb{E}[\nabla f_t(\mathbf{x})] = \nabla f(\mathbf{x})$ , and it follows from (60) that the variance of  $\nabla f_t(\mathbf{x})$  under  $\|\cdot\|_1$  satisfies  $\mathbb{E}[\|\nabla f_t(\mathbf{x}) - \nabla f(\mathbf{x})\|_1^2] = \mathbb{E}[\|\boldsymbol{\delta}_t - \mathbb{E}[\boldsymbol{\delta}_t]\|_1^2] \leq \sigma^2$ . Therefore, the stochastic loss functions  $f_0, f_1, \dots, f_{T-1}$  satisfy the condition (b). Next, we proceed to establish the lower bound for  $\mathbb{E}[\min_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|_1]$ , where  $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{T-1}$  are obtained by running Algorithm 3 on the constructed loss functions  $f_0, f_1, \dots, f_{T-1}$  with  $\mathbf{x}_0 = \mathbf{0}$ .

**Lower bound on gradient norm.** First note that by the construction of the loss functions and the choice  $\mathbf{x}_0 = 0$ , the coordinate-wise dynamics have the same distribution across coordinates, i.e.,  $(x_{0,i}, x_{1,i}, \dots, x_{T-1,i})$  has the same distribution as  $(x_{0,j}, x_{1,j}, \dots, x_{T-1,j})$  for all  $i, j \in [d]$ . Utilizing this observation, we have

$$\begin{aligned} \mathbb{E} \left[ \min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_1 \right] &= \mathbb{E} \left[ \min_{t \in [T]} \frac{1}{d} \sum_{i=1}^d |p'(x_{t,i})| \right] \\ &\geq \frac{1}{d} \sum_{i=1}^d \mathbb{E} \left[ \min_{t \in [T]} |p'(x_{t,i})| \right] \\ &= \mathbb{E} \left[ \min_{t \in [T]} |p'(x_{t,1})| \right]. \end{aligned} \quad (63)$$

So it suffices to focus on the dynamics in the first coordinate.

By the construction of the function  $p$  in the previous step, we have  $p'(0) = -\epsilon$ . Therefore, when  $x_{t,1} = 0$ , we have  $[\nabla f_t(\mathbf{x})]_1 = p'(x_{t,1})/d + \delta_{t,1} + \epsilon/d = \delta_{t,1}$ . Consequently, since  $x_{0,1} = 0$ , the first coordinate of  $\mathbf{x}_t$  will remain to be zero until at some step  $\delta_{t,1} \neq 0$ , in which case we must have  $\delta_{t,1} = -C$ . This implies that

$$\min\{t \in [T] : x_{t,1} \neq 0\} - 1 = \min \left\{ t \in [T] : [\nabla f_t(\mathbf{x}_t)]_1 = -C \right\} =: \tau,$$

and this stopping time  $\tau$  follows a geometric distribution with parameter  $\theta = \frac{5\epsilon^2}{d\sigma^2}$  again by the distribution of  $\delta_t$ . Hence,

$$\begin{aligned} \mathbb{P} \left[ \tau \geq \frac{1}{\theta} \right] &= \mathbb{P} \left[ \tau \geq \lceil \frac{1}{\theta} \rceil \right] = (1 - \theta)^{\lceil \frac{1}{\theta} \rceil} \\ &\geq \exp \left( -\theta \lceil \frac{1}{\theta} \rceil \right) \\ &\geq e^{-1-\theta} \geq e^{-2}. \end{aligned}$$

Now conditioned on the event that  $\tau \geq \frac{1}{\theta}$ , we can show that

$$\tau + N \geq \frac{1}{\theta} + \frac{1}{4\epsilon^2} \geq 2\sqrt{\frac{1}{\theta}} \cdot \frac{1}{4\epsilon^2} = \frac{1}{\sqrt{\theta}\epsilon} = \frac{\sqrt{d}\sigma}{\sqrt{5}\epsilon^2},$$

which is no smaller than  $T$  by the choice of  $\epsilon$  in Eq. 59. Therefore, we have  $x_{t,1} = 0$  for all  $t = 0, 1, \dots, \tau$  and there are at most  $T - \tau$  distinct points among  $x_{0,1}, x_{1,1}, \dots, x_{T,1}$ . Then since  $|x_{t+1} - x_t|$  is either 0 or  $\eta$  by the update rule of Algorithm 3 with  $\|\cdot\| = \|\cdot\|_\infty$ , we see that  $x_{0,1}, x_{1,1}, \dots, x_{T,1} \in \{0, \eta, 2\eta, \dots, (N-1)\eta\}$  on the event  $\tau \geq \frac{1}{\theta}$ , in which case we have  $p'(x_{t,1}) = -\epsilon$  for all  $t = 0, 1, \dots, T$  by our construction of  $p$ . Leveraging this, it follows from Eq. 63 that

$$\begin{aligned} \mathbb{E} \left[ \min_{t \in [T]} \|\nabla f(\mathbf{x}_t)\|_1 \right] &\geq \mathbb{E} \left[ \min_{t \in [T]} |p'(x_{t,1})| \right] \geq \mathbb{P} \left[ \tau \geq \frac{1}{\theta} \right] \cdot \mathbb{E} \left[ \min_{t \in [T]} |p'(x_{t,1})| \mid \tau \geq \frac{1}{\theta} \right] \\ &\geq e^{-2}\epsilon \end{aligned}$$

This gives the desired lower bound and completes the proof.  $\square$