

# Optimistic Online Mirror Descent for Bridging Stochastic and Adversarial Online Convex Optimization

**Sijia Chen**

CHENSJ@LAMDA.NJU.EDU.CN

*National Key Laboratory for Novel Software Technology, Nanjing University, China*

**Yu-Jie Zhang**

YUJIE.ZHANG@MS.K.U-TOKYO.AC.JP

*The University of Tokyo, Chiba, Japan*

**Wei-Wei Tu**

TUWWCN@GMAIL.COM

*Artificial Productivity Inc., Beijing, China*

**Peng Zhao**

ZHAOP@LAMDA.NJU.EDU.CN

**Lijun Zhang\***

ZHANGLJ@LAMDA.NJU.EDU.CN

*National Key Laboratory for Novel Software Technology, Nanjing University, China*

*School of Artificial Intelligence, Nanjing University, China*

**Editor:** Francesco Orabona

## Abstract

The stochastically extended adversarial (SEA) model, introduced by [Sachs et al. \(2022\)](#), serves as an interpolation between stochastic and adversarial online convex optimization. Under the smoothness condition on expected loss functions, it is shown that the expected *static regret* of optimistic follow-the-regularized-leader (FTRL) depends on the cumulative stochastic variance  $\sigma_{1:T}^2$  and the cumulative adversarial variation  $\Sigma_{1:T}^2$  for convex functions. [Sachs et al. \(2022\)](#) also provide a regret bound based on the maximal stochastic variance  $\sigma_{\max}^2$  and the maximal adversarial variation  $\Sigma_{\max}^2$  for strongly convex functions. Inspired by their work, we investigate the theoretical guarantees of optimistic online mirror descent (OMD) for the SEA model with smooth expected loss functions. For convex and smooth functions, we obtain the *same*  $\mathcal{O}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$  regret bound, but with a relaxation of the convexity requirement from individual functions to expected functions. For strongly convex and smooth functions, we establish an  $\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2)))$  bound, *better* than their  $\mathcal{O}((\sigma_{\max}^2 + \Sigma_{\max}^2) \log T)$  result. For exp-concave and smooth functions, our approach yields a *new*  $\mathcal{O}(d \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2))$  bound. Moreover, we introduce the first expected *dynamic regret* guarantee for the SEA model with convex and smooth expected functions, which is more favorable than static regret bounds in non-stationary environments. Furthermore, we expand our investigation to scenarios with non-smooth expected loss functions and propose novel algorithms built upon optimistic OMD with an implicit update, successfully attaining both static and dynamic regret guarantees.

## 1 Introduction

Online convex optimization (OCO) is a fundamental framework for online learning and has been applied in a variety of real-world applications such as spam filtering and portfolio

---

\*. Corresponding author.

management (Hazan, 2016). OCO problems can be mainly divided into two categories: adversarial online convex optimization (adversarial OCO) (Zinkevich, 2003; Hazan et al., 2007) and stochastic online convex optimization (SCO) (Nemirovski et al., 2009; Hazan and Kale, 2011; Lan, 2012). Adversarial OCO assumes that the loss functions are chosen arbitrarily or adversarially and the goal is to minimize the regret. SCO assumes that the loss functions are independently and identically distributed (i.i.d.), and the goal is to minimize the excess risk. Although the two models have been extensively studied (Shalev-Shwartz et al., 2009; Hazan, 2016; Orabona, 2019), in real scenarios the nature is not always completely adversarial or stochastic, but often lies somewhere in between.

### 1.1 The Stochastically Extended Adversarial Model

The stochastically extended adversarial (SEA) model is introduced by Sachs et al. (2022) as an intermediate problem setup between adversarial OCO and SCO. In round  $t \in [T]$ , the learner selects a decision  $\mathbf{x}_t$  from a convex feasible domain  $\mathcal{X} \subseteq \mathbb{R}^d$ , and nature chooses a distribution  $\mathfrak{D}_t$  from a set of distributions. Then, the learner suffers a loss  $f_t(\mathbf{x}_t)$ , where the individual function (also called random function)  $f_t$  is sampled from the distribution  $\mathfrak{D}_t$ . The distributions are allowed to vary over time, and by choosing them appropriately, the SEA model reduces to adversarial OCO, SCO, or other intermediate settings. Additionally, for each  $t \in [T]$ , they define the (conditional) expected function as  $F_t(\mathbf{x}) = \mathbb{E}_{f_t \sim \mathfrak{D}_t}[f_t(\mathbf{x})]$ .

Due to the randomness in the online process, our goal in the SEA model is to bound the *expected regret* against any fixed comparator  $\mathbf{u} \in \mathcal{X}$ , defined as

$$\mathbb{E}[\mathbf{Reg}_T(\mathbf{u})] \triangleq \mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right]. \quad (1)$$

Furthermore, to capture the characteristics of the SEA model, Sachs et al. (2022) introduce the following quantities. For each  $t \in [T]$ , define the (conditional) variance of gradients as

$$\sigma_t^2 = \sup_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{f_t \sim \mathfrak{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2]. \quad (2)$$

Notice that both  $F_t(\mathbf{x})$  and  $\sigma_t^2$  can be random variables due to the randomness of distribution  $\mathfrak{D}_t$ . Then, the cumulative stochastic variance can be defined as

$$\sigma_{1:T}^2 = \mathbb{E} \left[ \sum_{t=1}^T \sigma_t^2 \right], \quad (3)$$

which reflects the stochastic aspect of the online process. Moreover, the cumulative adversarial variation is defined as

$$\Sigma_{1:T}^2 = \mathbb{E} \left[ \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right], \quad (4)$$

where  $\nabla F_0(\mathbf{x}) = \mathbf{0}$ , reflecting the adversarial difficulty.<sup>1</sup>

---

1. If the nature is oblivious, then both  $F_t(\mathbf{x})$  and  $\sigma_t^2$  will be deterministic and we can remove the expectation in (3) and (4).

## 1.2 Existing Results

With the smoothness of expected loss functions, [Sachs et al. \(2022\)](#) establish a series of results for the SEA model, including convex functions and strongly convex functions.

For convex and smooth functions, they prove an  $\mathcal{O}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$  regret bound of optimistic follow-the-regularized-leader (FTRL). Note that they require the individual functions  $\{f_t\}_{t=1}^T$  to be convex here, but we achieve the same result by only requiring the convexity of expected functions  $\{F_t\}_{t=1}^T$  in this paper. When facing the adversarial setting,  $\sigma_t^2 = 0$  for all  $t$  and  $\Sigma_{1:T}^2$  is equivalent to the gradient variation  $V_T \triangleq \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2$ , so their bound implies a regret bound in the form of  $\sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}) \leq \mathcal{O}(\sqrt{V_T})$ , matching the gradient-variation bound of [Chiang et al. \(2012\)](#) and also recovering the  $\mathcal{O}(\sqrt{T})$  bound in the worst case ([Zinkevich, 2003](#)). In the SCO setting,  $\Sigma_{1:T}^2 = 0$  since  $F_1 = \dots = F_T \triangleq F$ , and  $\sigma_t = \sigma$  for all  $t$ , where  $\sigma$  denotes the variance of stochastic gradients. Then their bound implies an  $\mathcal{O}(\sigma\sqrt{T})$  regret, leading to an excess risk bound in the form of  $F(\mathbf{x}_T) - \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}) \leq \mathcal{O}(\sigma/\sqrt{T})$  through the standard online-to-batch conversion ([Cesa-Bianchi et al., 2004](#)).

To investigate the strongly convex case, they assume that the *maximum value* of stochastic variance is  $\sigma_{\max}^2$  and the *maximum value* of adversarial variation is  $\Sigma_{\max}^2$ ; please refer to Assumption 3 for details. Then [Sachs et al. \(2022\)](#) prove an  $\mathcal{O}((\sigma_{\max}^2 + \Sigma_{\max}^2) \log T)$  expected regret bound of optimistic FTRL for  $\lambda$ -strongly convex and smooth functions. Considering the adversarial setting, we have  $\sigma_{\max}^2 = 0$  and  $\Sigma_{\max}^2 \leq 4G^2$  where  $G$  is the upper bound of individual function gradients, so their bound implies an  $\mathcal{O}(\log T)$  regret bound. We note that unlike in the convex and smooth case, their expected regret bound fails to recover the  $\mathcal{O}(\log V_T)$  gradient-variation bound ([Zhang et al., 2022](#)). In the SCO setting, we have  $\Sigma_{\max}^2 = 0$  and  $\sigma_{\max}^2 = \sigma^2$ . Therefore, their result brings an  $\mathcal{O}([\sigma^2 \log T]/T)$  excess risk bound through the online-to-batch conversion.

## 1.3 Our Contributions

Optimistic FTRL is an optimistic online learning algorithm ([Rakhlin and Sridharan, 2013](#)), which aims to exploit prior knowledge during the online process. Optimistic online mirror descent (OMD) is another popular optimistic online learning algorithm, from which the gradient-variation bound of [Chiang et al. \(2012\)](#) is (originally) derived. With the promising outcomes of optimistic FTRL ([Sachs et al., 2022](#)), it is natural to inquire about optimistic OMD's theoretical guarantees for the SEA model, and we address this below.

- For convex and smooth functions, optimistic OMD enjoys the same  $\mathcal{O}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$  expected regret bound as [Sachs et al. \(2022\)](#), but reduces their need for convexity of individual functions to a need for convexity of expected functions.
- For strongly convex and smooth functions, optimistic OMD attains an  $\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2)/(\sigma_{\max}^2 + \Sigma_{\max}^2)))$  bound, better than the  $\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T)$  bound of optimistic FTRL from [Sachs et al. \(2022\)](#).
- For exp-concave and smooth functions, our work establishes a new  $\mathcal{O}(d \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2))$  bound of optimistic OMD, where  $d$  denotes the dimensionality of decisions.

- Our better results stem from more careful analyses and do not imply that optimistic OMD has inherent superiority over optimistic FTRL for regret minimization. When encountering convex functions, we present a different analysis from the analysis of optimistic FTRL in [Sachs et al. \(2022\)](#), thereby similarly weakening the convexity-related assumption as in optimistic OMD while achieving the same regret bound. We also provide new analyses for strongly convex functions and exp-concave functions respectively, both obtaining the same expected regret bounds as optimistic OMD.

**Extension to Dynamic Regret.** The metric (1) is commonly referred to as expected static regret since the comparator is unchanged over time. We further extend the scope of the SEA model to optimize *expected dynamic regret* ([Zinkevich, 2003](#)), defined as

$$\mathbb{E}[\mathbf{Reg}_T^d(\mathbf{u}_1, \dots, \mathbf{u}_T)] \triangleq \mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \right], \quad (5)$$

where  $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X}$  is a sequence of (potentially) time-varying comparators. Note that the comparators can depend on the expected functions  $\{F_t\}_{t=1}^T$  and are required to be independent of the individual functions  $\{f_t\}_{t=1}^T$ . To optimize the dynamic regret, we introduce the path length  $P_T = \mathbb{E}[\sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2]$  to measure the non-stationarity level, where  $\mathbb{E}[\cdot]$  is taken over the potential randomness of the expected functions. Notably, the static regret (1) can be treated as a special case with  $\mathbf{u}_1 = \dots = \mathbf{u}_T = \mathbf{u}$ . For the SEA model with convex and smooth expected functions, we obtain an  $\mathcal{O}(P_T + \sqrt{1 + P_T}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}))$  expected dynamic regret. The bound is new and immediately recovers the  $\mathcal{O}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$  expected static regret given  $P_T = 0$ . It can also imply the  $\mathcal{O}(\sqrt{(1 + P_T + V_T)(1 + P_T)})$  gradient-variation dynamic regret bound of [Zhao et al. \(2020, 2024\)](#) in the adversarial setting and reduce to the  $\mathcal{O}(\sqrt{T(1 + P_T)})$  dynamic regret in the worst case ([Zhang et al., 2018](#)). We regard the support of dynamic regret as an advantage of optimistic OMD over optimistic FTRL. To the best of our knowledge, even  $\mathcal{O}(\sqrt{T(1 + P_T)})$  dynamic regret has not been established for FTRL-style methods in online convex optimization.

**Extension to Non-smooth Functions.** In addition, by combining optimistic OMD with *implicit update*, we extend our investigation to *non-smooth* loss functions. For the SEA model with convex and non-smooth functions, we first establish an  $\mathcal{O}(\sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$  static regret, based on which we further propose a two-layer algorithm equipped with an  $\mathcal{O}(\sqrt{1 + P_T}(\sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}))$  dynamic regret, where  $\tilde{\sigma}_{1:T}^2$  defined in (32) represents a slightly more relaxed measure than  $\sigma_{1:T}^2$ .

Based on all the above theoretical guarantees, we apply optimistic OMD to a variety of intermediate cases between adversarial OCO and SCO. This leads to *better* results for strongly convex functions and *new* results for exp-concave functions, thereby enriching our understanding of the intermediate scenarios. Furthermore, our emphasis on dynamic regret minimization enables us to derive novel corollaries for the online label shift problem ([Bai et al., 2022](#)), an interesting new problem setup with practical appeals.

Compared to our earlier conference version ([Chen et al., 2023](#)), this extended version provides significantly more results, along with refined presentations and more detailed anal-

ysis. Firstly, by revisiting and refining our analysis, we provide a better regret bound for strongly convex functions than our previous bound in [Chen et al. \(2023\)](#). Secondly, we incorporate a more detailed analysis of dynamic regret minimization within the SEA model, adding insights to explain the optimism design’s rationale and highlighting the disadvantages of alternative approaches. Thirdly, we investigate the SEA model with *non-smooth* functions, where we employ optimistic OMD with an implicit update and obtain favorable regret guarantees. Additionally, we explore dynamic regret minimization with non-smooth functions. Lastly, we apply our findings to address the online label shift problem, yielding results that further demonstrate the SEA model’s real-world applicability.

**Organization.** The remainder of the paper is structured as follows. Section 2 briefly reviews the related work. Our main results can be found in Section 3, in which we establish theoretical guarantees for convex, strongly convex, and exp-concave loss functions under the smoothness condition on loss functions respectively. In Section 4, we extend the investigations to dynamic regret minimization and non-smooth loss functions. In Section 5, we illustrate our results by giving some special implications, such as online learning with limited resources and online label shift. Section 6 concludes the paper and discusses future work. Some omitted details and proofs are provided in the appendix.

## 2 Related Work

This section reviews related works in adversarial OCO, SCO, and intermediate settings.

### 2.1 Adversarial Online Convex Optimization

Adversarial OCO can be seen as a repeated game between the online learner and the nature (or called the environment). In round  $t \in [T]$ , the online learner chooses a decision  $\mathbf{x}_t$  from the convex feasible set  $\mathcal{X} \subseteq \mathbb{R}^d$ , and suffers a convex loss  $f_t(\mathbf{x}_t)$  which the nature may adversarially select. The goal in adversarial OCO is to minimize the *regret*:

$$\mathbf{Reg}_T \triangleq \sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x}),$$

which measures the cumulative loss difference between the learner and the best decision in hindsight ([Orabona, 2019](#)). For convex functions, online gradient descent (OGD) achieves an  $\mathcal{O}(\sqrt{T})$  regret with a step size of  $\eta_t = \mathcal{O}(1/\sqrt{t})$  ([Zinkevich, 2003](#)). For  $\lambda$ -strongly convex functions, an  $\mathcal{O}(\frac{1}{\lambda} \log T)$  bound is attained by OGD with  $\eta_t = \mathcal{O}(1/[\lambda t])$  ([Shalev-Shwartz, 2007](#)). For  $\alpha$ -exp-concave functions, online Newton step (ONS) ([Hazan et al., 2007](#)) obtains an  $\mathcal{O}(\frac{d}{\alpha} \log T)$  bound. Those results are considered minimax optimal ([Ordentlich and Cover, 1998](#); [Abernethy et al., 2008](#)) and cannot be improved in general.

Furthermore, various algorithms have been proposed to achieve *problem-dependent* regret guarantees, which safeguard the minimax rates in the worst case and become better when problems satisfy benign properties such as smoothness ([Srebro et al., 2010](#); [Chiang et al., 2012](#); [Orabona et al., 2012](#); [Zhao et al., 2020, 2024](#)), sparsity ([Duchi et al., 2011](#); [McMahan and Streeter, 2010](#); [Gaillard and Wintenberger, 2018](#)), or other structural properties ([Kingma and Ba, 2015](#); [Joulani et al., 2020](#)). Among them, it is shown by [Chiang](#)

et al. (2012) that the regret for OCO with smooth functions can be upper bounded by the gradient-variation quantity, defined as

$$V_T = \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2. \quad (6)$$

Specifically, using the OMD framework with suitable configurations can attain an  $\mathcal{O}(\sqrt{V_T})$  regret for convex and smooth functions and attain an  $\mathcal{O}(\frac{d}{\alpha} \log V_T)$  regret for  $\alpha$ -exp-concave and smooth functions. Zhang et al. (2022) extend the result to  $\lambda$ -strongly convex and smooth functions, achieving an  $\mathcal{O}(\frac{1}{\lambda} \log V_T)$  bound. These bounds are notably tighter than previous problem-independent results when the loss functions change slowly such that the gradient variation  $V_T$  is small.

In addition, Rakhlin and Sridharan (2013) introduce the paradigm of optimistic online learning, designed to leverage prior knowledge about upcoming loss functions. In this approach, the learner receives a prediction of the next loss in each round, which is used to secure tighter bounds when the predictions prove accurate and still preserve the worst-case regret bound otherwise. Then two frameworks are developed: optimistic FTRL and optimistic OMD, where the latter generalized the algorithm of Chiang et al. (2012).

## 2.2 Stochastic Online Convex Optimization

SCO assumes i.i.d. loss functions and aims to minimize the convex objective in an expectation form:  $\min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x})$ , where  $F(\mathbf{x}) = \mathbb{E}_{f \sim \mathcal{D}}[f(\mathbf{x})]$ . The performance measure is the *excess risk* of the solution point over the optimum, that is,  $F(\mathbf{x}_T) - \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x})$ .

For Lipschitz and convex functions, stochastic gradient descent (SGD) achieves an  $\mathcal{O}(1/\sqrt{T})$  excess risk bound. Improved rates are achievable when functions have additional properties. For smooth functions, SGD reaches an  $\mathcal{O}(1/T + \sqrt{F_*/T})$  rate with  $F_* = \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x})$ , which will be tighter than  $\mathcal{O}(1/\sqrt{T})$  when  $F_*$  is small (Srebro et al., 2010). For  $\lambda$ -strongly convex functions, Hazan and Kale (2011) establish an  $\mathcal{O}(1/[\lambda T])$  excess risk bound through a variant of SGD. For  $\alpha$ -exp-concave functions, ONS provides an  $\mathcal{O}(d \log T/[\alpha T])$  rate (Hazan et al., 2007; Mahdavi et al., 2015). When functions satisfy strong convexity and smoothness simultaneously, accelerated stochastic approximation (AC-SA) achieves an  $\mathcal{O}(1/T)$  rate with a smaller constant (Ghadimi and Lan, 2012). Even faster results can be attained with strengthened conditions and advanced algorithms (Johnson and Zhang, 2013; Zhang et al., 2013; Neu and Rosasco, 2018; Zhang and Zhou, 2019).

## 2.3 Intermediate Setting

In recent years, intermediate settings between adversarial OCO and SCO have drawn attention in prediction with expert advice (PEA) problems (Amir et al., 2020) and bandit problems (Zimmert and Seldin, 2021). Amir et al. (2020) study the stochastic regime with adversarial corruptions in PEA problems, achieving an  $\mathcal{O}(\log N/\Delta + C_T)$  bound, where  $N$  is the number of experts,  $\Delta$  the suboptimality gap and  $C_T \geq 0$  the corruption level. In bandit problems, Zimmert and Seldin (2021) focus on the adversarial regime with a self-bounding constraint, establishing an  $\mathcal{O}(N \log T/\Delta + \sqrt{C_T N \log T/\Delta})$  bound. Ito (2021) further demonstrates an expected regret bound of  $\mathcal{O}(\log N/\Delta + \sqrt{C_T \log N/\Delta})$  in this con-



text. However, as mentioned by Ito (2021), we know very little about the intermediate setting in OCO, with recent contributions like Sachs et al. (2022) being exceptions.

### 3 Optimistic Mirror Descent for the SEA Model

In this section, we begin with listing the definitions and assumptions that will be used later. Then, we introduce OPTIMISTIC OMD, our main algorithmic framework. After that, we discuss its theoretical guarantees for the SEA model, along with new results of optimistic FTRL. The final subsection is dedicated to analyzing these results.

#### 3.1 Definitions and Assumptions

We first introduce the definitions of strongly convex functions and exp-concave functions (Boyd and Vandenberghe, 2004; Cesa-Bianchi and Lugosi, 2006).

**Definition 1** (strongly convex function). A function  $f : \mathcal{X} \mapsto \mathbb{R}$  is  $\lambda$ -strongly convex with respect to a norm  $\|\cdot\|$  if  $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\lambda}{2} \|\mathbf{y} - \mathbf{x}\|^2$  for all  $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ .

**Definition 2** (exp-concave function). A function  $f : \mathcal{X} \mapsto \mathbb{R}$  is  $\alpha$ -exp-convex if  $\exp(-\alpha f(\cdot))$  is concave over  $\mathcal{X}$ .

The assumptions listed below may be employed in our analysis. It is important to note that we will clearly specify the assumptions utilized in the theorem statements.

**Assumption 1** (gradient norms boundedness). The gradient norms of all the individual functions are bounded by  $G$ , i.e. for all  $t \in [T]$ , we have  $\max_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x})\|_2 \leq G$ .

**Assumption 2** (domain boundedness). The domain  $\mathcal{X}$  contains the origin  $\mathbf{0}$ , and the diameter of  $\mathcal{X}$  is bounded by  $D$ , i.e., for all  $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ , we have  $\|\mathbf{x} - \mathbf{y}\|_2 \leq D$ .

**Assumption 3** (maximal stochastic variance and adversarial variation). All the variances of realizable gradients are at most  $\sigma_{\max}^2$ , and all the adversarial variations are upper bounded by  $\Sigma_{\max}^2$ , i.e.,  $\forall t \in [T]$ , it holds that  $\sigma_t^2 \leq \sigma_{\max}^2$  and  $\sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \leq \Sigma_{\max}^2$ .

**Assumption 4** (smoothness of expected functions). For all  $t \in [T]$ , the expected function  $F_t(\cdot)$  is  $L$ -smooth over  $\mathcal{X}$ , i.e.,  $\|\nabla F_t(\mathbf{x}) - \nabla F_t(\mathbf{y})\|_2 \leq L\|\mathbf{x} - \mathbf{y}\|_2$ ,  $\forall \mathbf{x}, \mathbf{y} \in \mathcal{X}$ .

**Assumption 5** (convexity of expected functions). For all  $t \in [T]$ , the expected function  $F_t(\cdot)$  is convex over  $\mathcal{X}$ .

**Assumption 6** (strong convexity of expected functions). For  $t \in [T]$ , the expected function  $F_t(\cdot)$  is  $\lambda$ -strongly convex over  $\mathcal{X}$ .

**Assumption 7** (exponential concavity of individual functions). For  $t \in [T]$ , the individual function  $f_t(\cdot)$  is  $\alpha$ -exp-concave over  $\mathcal{X}$ .

**Assumption 8** (convexity of individual functions). For all  $t \in [T]$ , the individual function  $f_t(\cdot)$  is convex over  $\mathcal{X}$ .

Moreover, we will use the following property of exp-concave functions in our analysis.

**Lemma 1** (Lemma 3 of Hazan et al. (2007)). Suppose  $f : \mathcal{X} \mapsto \mathbb{R}$  is  $\alpha$ -exp-convex, under Assumption 1 and 2, we have  $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\beta}{2} (\mathbf{y} - \mathbf{x})^\top \nabla f(\mathbf{x}) \nabla f(\mathbf{x})^\top (\mathbf{y} - \mathbf{x})$  for all  $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ , where  $\beta = \frac{1}{2} \min \left\{ \frac{1}{4GD}, \alpha \right\}$ .

**Algorithm 1** Optimistic Online Mirror Descent (Optimistic OMD)

---

**Input:** Regularizer  $\psi_t : \mathcal{X} \mapsto \mathbb{R}$

- 1: Set  $\mathbf{x}_1 = \hat{\mathbf{x}}_1$  to be any point in  $\mathcal{X}$
- 2: **for**  $t = 1, \dots, T$  **do**
- 3:   Submit  $\mathbf{x}_t$  and the nature selects a distribution  $\mathfrak{D}_t$
- 4:   Receive  $f_t(\cdot)$ , which is sampled from  $\mathfrak{D}_t$
- 5:   Receive an optimistic vector  $M_{t+1}$  encoding certain prior knowledge of  $f_{t+1}(\cdot)$
- 6:   Update  $\hat{\mathbf{x}}_{t+1}$  and  $\mathbf{x}_{t+1}$  according to (7) and (8)
- 7: **end for**

---

**3.2 The Algorithm**

Optimistic OMD is a versatile and powerful framework for online learning (Rakhlin and Sridharan, 2013). During the learning process, it maintains two sequences  $\{\mathbf{x}_t\}_{t=1}^T$  and  $\{\hat{\mathbf{x}}_t\}_{t=1}^T$ . In round  $t \in [T]$ , the learner first submits the decision  $\mathbf{x}_t$  and observes the individual function  $f_t(\cdot)$ . Then, an optimistic vector  $M_{t+1} \in \mathbb{R}^d$  is received that encodes certain prior knowledge of the (unknown) function  $f_{t+1}(\cdot)$ , and the algorithm updates by

$$\hat{\mathbf{x}}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \mathcal{D}_{\psi_t}(\mathbf{x}, \hat{\mathbf{x}}_t), \quad (7)$$

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle M_{t+1}, \mathbf{x} \rangle + \mathcal{D}_{\psi_{t+1}}(\mathbf{x}, \hat{\mathbf{x}}_{t+1}), \quad (8)$$

where  $\mathcal{D}_{\psi}(\mathbf{x}, \mathbf{y}) = \psi(\mathbf{x}) - \psi(\mathbf{y}) - \langle \nabla \psi(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$  denotes the Bregman divergence induced by a differentiable convex function  $\psi : \mathcal{X} \mapsto \mathbb{R}$  (or usually called regularizer). In our work, we allow the regularizer to be time-varying. The specific choice of  $\psi_t(\cdot)$  depends on the type of online functions and will be determined later.

To leverage the possible smoothness of functions, we simply set the optimism as the last-round gradient, that is,  $M_{t+1} = \nabla f_t(\mathbf{x}_t)$  (Chiang et al., 2012). We initialize  $\mathbf{x}_1 = \hat{\mathbf{x}}_1$  as an arbitrary point in  $\mathcal{X}$ . The overall procedures are summarized in Algorithm 1.

**Remark 1.** If we drop the expectation operation, the measure (1) becomes the standard regret. Consequently, a straightforward way is to integrate existing regret bounds of optimistic OMD (Chiang et al., 2012; Rakhlin and Sridharan, 2013) and subsequently simplify the expectation. However, as elaborated in Sachs et al. (2022, Remark 4), this approach only yields very loose bounds. Therefore, it becomes necessary to dig into the analysis and scrutinize the influence of expectations during the intermediate steps.  $\triangleleft$

In the following, we consider three different instantiations of Algorithm 1, each corresponding to the SEA model with different types of functions: convex, strongly convex, and exp-concave functions, respectively. We also provide their respective theoretical guarantees.

**3.3 Convex and Smooth Functions**

In this part, we focus on the case that *expected functions* are convex and smooth. Sachs et al. (2022) require individual functions  $f_t(\cdot)$  ( $t \in [T]$ ) to be convex (see Assumption A1 of their paper), whereas we only require expected functions  $F_t(\cdot)$  ( $t \in [T]$ ) to be convex, which



is a much weaker condition. This relaxation, which has been studied in many stochastic optimization works (Shalev-Shwartz, 2016; Hu et al., 2017; Ahn et al., 2020), is due to the observation that the expectation in (1) eliminates the need for convexity in individual functions. Specifically, for any fixed  $\mathbf{u} \in \mathcal{X}$  we have

$$\mathbb{E}[f_t(\mathbf{x}_t) - f_t(\mathbf{u})] = \mathbb{E}[F_t(\mathbf{x}_t) - F_t(\mathbf{u})] \leq \mathbb{E}[\langle \nabla F_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle] = \mathbb{E}[\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle]. \quad (9)$$

The inequality arises from the convexity of  $F_t(\cdot)$  and the last step is due to the interchangeability of differentiation and integration by Leibniz integral rule. Note that the independence between  $\mathbf{u}$  and  $f_t$  is important for this derivation. We emphasize that if  $\mathbf{u}$  is chosen based on random functions, then the convexity of random functions will be necessary.<sup>2</sup>

Below, we focus on the optimization over bound the expected regret in terms of the linearized function, i.e.,  $\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle$ . For convex and smooth functions, we configure the algorithm with Euclidean regularizer

$$\psi_t(\mathbf{x}) = \frac{1}{2\eta_t} \|\mathbf{x}\|_2^2 \quad \text{and step size} \quad \eta_t = \frac{D}{\sqrt{\delta + 4G^2 + \bar{V}_{t-1}}}, \quad (10)$$

where  $\bar{V}_{t-1} = \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1})\|_2^2$  (assuming  $\nabla f_0(\mathbf{x}_0) = 0$ ) and  $\delta > 0$  is a parameter to be specified later. Then, the optimistic OMD updates in (7) and (8) become

$$\hat{\mathbf{x}}_{t+1} = \Pi_{\mathcal{X}}[\hat{\mathbf{x}}_t - \eta_t \nabla f_t(\mathbf{x}_t)], \quad \mathbf{x}_{t+1} = \Pi_{\mathcal{X}}[\hat{\mathbf{x}}_{t+1} - \eta_{t+1} \nabla f_t(\mathbf{x}_t)], \quad (11)$$

where  $\Pi_{\mathcal{X}}[\cdot]$  denotes the Euclidean projection onto the feasible domain  $\mathcal{X}$ . The algorithm executes gradient descent twice per round, using an adaptive step size akin to self-confident tuning (Auer et al., 2002). This approach obviates the need for the doubling trick used in prior works (Chiang et al., 2012; Rakhlin and Sridharan, 2013; Jadbabaie et al., 2015).

Below, we present the theoretical guarantee of optimistic OMD for the SEA model with convex and smooth functions. The proof is in Section 3.6.1.

**Theorem 1.** *Under Assumptions 1, 2, 4 and 5, optimistic OMD with regularizer (10) and updates (11) enjoys the following guarantee:*

$$\mathbb{E}[\mathbf{Reg}_T(\mathbf{u})] \leq 5\sqrt{10}D^2L + \frac{5\sqrt{5}DG}{2} + 5\sqrt{2}D\sqrt{\sigma_{1:T}^2} + 5D\sqrt{\Sigma_{1:T}^2} = \mathcal{O}\left(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}\right),$$

where we set  $\delta = 10D^2L^2$  in (10).

**Remark 2.** Theorem 1 demonstrates the same regret bound as the work of Sachs et al. (2022), but under weaker assumptions—we require only the convexity of expected functions, as opposed to individual functions in their work. The regret bound is optimal according to the lower bound of Sachs et al. (2022, Theorem 6).  $\triangleleft$

2. Fortunately, a favorable choice of  $\mathbf{u}$  is usually independent of random functions. For instance, if the nature is oblivious, we can choose  $\mathbf{u} = \mathbf{u}^* \in \arg \min_{\mathbf{u} \in \mathcal{X}} \sum_{t=1}^T F_t(\mathbf{u})$ , which only depends on expected functions. Additionally, fitting to  $\{F_1, \dots, F_T\}$  is preferable in practice as fitting to  $\{f_1, \dots, f_T\}$  might cause overfitting.

**Table 1:** Comparison of different theoretical guarantees of SEA for strongly convex functions.

Reference	Regret bound of SEA with $\lambda$ -strongly convex functions
<a href="#">Sachs et al. (2022)</a>	$\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T)$
<a href="#">Chen et al. (2023)</a>	$\mathcal{O}(\min\{\frac{G^2}{\lambda} \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2), \frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T\})$
This paper	$\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2)))$

In this subsection’s final part, we provide a new result of optimistic FTRL for the SEA model. Notably, we illustrate that even *without* the convexity of individual functions, optimistic FTRL can achieve the *same* guarantee as [Sachs et al. \(2022\)](#). This is achieved by using a linearized surrogate loss  $\{\langle \nabla f_t(\mathbf{x}_t), \cdot \rangle\}_{t=1}^T$  instead of the original loss  $\{f_t(\cdot)\}_{t=1}^T$ .

**Theorem 2.** *Under Assumptions 1, 2, 4, and 5 (without assuming convexity of individual functions), with an appropriate setup for the optimistic FTRL (see details in Appendix A.1), the expected regret is at most  $\mathcal{O}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$ .*

### 3.4 Strongly Convex and Smooth Functions

In this part, we examine the case when *expected functions* are strongly convex and smooth. We still employ optimistic OMD (Algorithm 1) and define the regularizer as

$$\psi_t(\mathbf{x}) = \frac{1}{2\eta_t} \|\mathbf{x}\|_2^2 \quad \text{with step size} \quad \eta_t = \frac{2}{\lambda t}. \quad (12)$$

It is worth mentioning that this step size configuration is *new* and much simpler than the self-confident step size used in earlier research on gradient-variation bounds for strongly convex and smooth functions ([Zhang et al., 2022](#)). Then the update rules maintain the same form as (11) in essence. We provide the following expected regret bound for the SEA model with strongly convex and smooth functions, the proof of which is in Section 3.6.2.

**Theorem 3.** *Under Assumptions 1, 2, 3, 4 and 6, optimistic OMD with regularizer (12) and updates (11) enjoys the following guarantee*

$$\begin{aligned} \mathbb{E}[\mathbf{Reg}_T(\mathbf{u})] &\leq \frac{32\sigma_{\max}^2 + 16\Sigma_{\max}^2}{\lambda} \ln \left( \frac{1}{2\sigma_{\max}^2 + \Sigma_{\max}^2} (2\sigma_{1:T}^2 + \Sigma_{1:T}^2) + 1 \right) + \frac{64\sigma_{\max}^2 + 32\Sigma_{\max}^2}{\lambda} \\ &\quad + \frac{16L^2D^2}{\lambda} \ln \left( 1 + 8\sqrt{2}\frac{L}{\lambda} \right) + \frac{16L^2D^2 + 4G^2}{\lambda} + \frac{\lambda D^2}{4} \\ &= \mathcal{O} \left( \frac{1}{\lambda} (\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2)) \right). \end{aligned}$$

Table 1 compares our result with those previously reported by [Sachs et al. \(2022\)](#) and our earlier conference version ([Chen et al., 2023](#)). Our result is strictly better than theirs, and we demonstrate the advantages in the following.

**Remark 3.** Compare to the  $\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T)$  bound of [Sachs et al. \(2022\)](#), our result shows advantages in benign problems with small cumulative quantities  $\sigma_{1:T}^2$  and  $\Sigma_{1:T}^2$ . Notably, even when  $\sigma_{1:T}^2$  and  $\Sigma_{1:T}^2$  are small,  $\sigma_{\max}^2$  and  $\Sigma_{\max}^2$  can be large, making their bound less effective. For instance, in an adversarial setting where  $\sigma_{1:T}^2 = \sigma_{\max}^2 = 0$  and online functions only change *once* such that  $\Sigma_{1:T}^2 = \Sigma_{\max}^2 = \mathcal{O}(1)$ , Theorem 3 yields an  $\mathcal{O}(1)$  bound, outperforming the  $\mathcal{O}(\log T)$  guarantee of [Sachs et al. \(2022\)](#). Furthermore, our bound can imply an  $\mathcal{O}(\frac{G^2}{\lambda} \log V_T)$  gradient-variation bound in adversarial OCO settings, whereas the bound of [Sachs et al. \(2022\)](#) cannot.  $\triangleleft$

**Remark 4.** Our new result surpasses the  $\mathcal{O}(\min\{\frac{G^2}{\lambda} \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2), \frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T\})$  bound from our earlier conference version ([Chen et al., 2023](#)). It exhibits greater adaptivity since  $\mathcal{O}(\sigma_{\max}^2 + \Sigma_{\max}^2)$  is always at most  $\mathcal{O}(G^2)$  and  $\mathcal{O}((\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2))$  is always at most  $\mathcal{O}(T)$ . This improvement is due to a refined analysis — we apply Lemma 6 to obtain a regret bound of the  $(\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2))$  form, which is inspired by Lemma 6 of [Chen et al. \(2023\)](#). See Section 3.6.2 for details.  $\triangleleft$

**Remark 5.** Our new upper bound in Theorem 3 does not contradict with the  $\Omega(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T)$  lower bound of [Sachs et al. \(2022, Theorem 8\)](#), because their lower bound focuses on the worst-case behavior while our result is better only in certain cases.  $\triangleleft$

Similar to Theorem 3, we demonstrate that for strongly convex and smooth functions, optimistic FTRL can also attain the *same* guarantee as optimistic OMD for the SEA model.

**Theorem 4.** *Under Assumptions 1, 2, 3, 4 and 6, with an appropriate setup for the optimistic FTRL (see details in Appendix A.2), the expected regret is at most  $\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2)))$ .*

### 3.5 Exp-concave and Smooth Functions

We further explore the SEA model for exp-concave and smooth functions. Notably, [Sachs et al. \(2022\)](#) only investigate convex and strongly convex functions, without studying exp-concave functions. Our results and analysis in this part is a *new* contribution.

Throughout this part, we will assume the individual functions are exp-concave rather than the expected functions, see Assumption 7. This is due to the need to use the exponential concavity of individual functions in our regret analysis. It is common in stochastic exp-concave optimization to assume exp-concavity of individual functions ([Mahdavi et al., 2015; Koren and Levy, 2015](#)). Importantly, we need to emphasize that the exponential concavity of individual functions *does not* imply the same for expected functions, which implies that the two assumptions are incomparable.

Following [Chiang et al. \(2012\)](#), we set the regularizer  $\psi_t(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_{H_t}^2$ , where  $H_t = I + \frac{\beta}{2} G^2 I + \frac{\beta}{2} \sum_{s=1}^{t-1} \nabla f_s(\mathbf{x}_s) \nabla f_s(\mathbf{x}_s)^\top$ ,  $I$  is the  $d$ -dimensional identity matrix, and  $\beta = \frac{1}{2} \min\{\frac{1}{4GD}, \alpha\}$ . Then, the updating rules of optimistic OMD in (7) and (8) become

$$\hat{\mathbf{x}}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}_t\|_{H_t}^2, \quad (13)$$

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}_{t+1}\|_{H_{t+1}}^2. \quad (14)$$

For exp-concave and smooth functions, we can realize the following bound of optimistic OMD for the SEA model with proof in Section 3.6.3.

**Theorem 5.** *Under Assumptions 1, 2, 4 and 7, optimistic OMD with updates (13) and (14) enjoys the following guarantee:*

$$\begin{aligned}\mathbb{E}[\text{Reg}_T(\mathbf{u})] &\leq \frac{16d}{\beta} \ln \left( \frac{\beta}{d} \sigma_{1:T}^2 + \frac{\beta}{2d} \Sigma_{1:T}^2 + \frac{\beta}{8d} G^2 + 1 \right) + \frac{16d}{\beta} \ln(32L^2 + 1) + D^2 \left( 1 + \frac{\beta}{2} G^2 \right) \\ &= \mathcal{O} \left( \frac{d}{\alpha} \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2) \right),\end{aligned}$$

where  $\beta = \frac{1}{2} \min \left\{ \frac{1}{4GD}, \alpha \right\}$ , and  $d$  is the dimensionality of decisions.

**Remark 6.** This is the *first* regret bound for the SEA model with exp-concave and smooth functions. Owing to analytical differences, we are unable to attain an  $\mathcal{O}(\frac{d}{\alpha}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T)$  regret bound, and further we can not get an  $\mathcal{O}(\frac{d}{\alpha}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2)/(\sigma_{\max}^2 + \Sigma_{\max}^2)))$  bound as in the strongly convex case (Theorem 3). We will investigate this possibility in the future.  $\triangleleft$

Similarly, we obtain the same guarantee by optimistic FTRL in the exp-concave case.

**Theorem 6.** *Under Assumptions 1, 2, 4 and 7 with an appropriate setup for the optimistic FTRL (see details in Appendix A.3), the expected regret is at most  $\mathcal{O}(\frac{d}{\alpha} \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2))$ .*

### 3.6 Analysis

In this section, we analyze the three theoretical guarantees based on optimistic OMD. Analyses of optimistic FTRL and proofs of all lemmas used are postponed to Appendix A.

#### 3.6.1 PROOF OF THEOREM 1

**Proof** According to (9), we can bound the expected regret for convex expected functions by a linearized function, that is

$$\begin{aligned}\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] &= \mathbb{E} \left[ \sum_{t=1}^T F_t(\mathbf{x}_t) - \sum_{t=1}^T F_t(\mathbf{u}) \right] \\ &\leq \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla F_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \right] = \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \right].\end{aligned}\quad (15)$$

Then we present a variant of the Bregman proximal inequality lemma (Nemirovski, 2005, Lemma 3.1), commonly used in optimistic OMD analysis, with proof in Appendix A.4.

**Lemma 2** (variant of Bregman proximal inequality). *Assume  $\psi_t(\cdot)$  is an  $\alpha$ -strongly convex function with respect to  $\|\cdot\|$ , and denote by  $\|\cdot\|_*$  the dual norm. Based on the updating rules of optimistic OMD in (7) and (8), for all  $\mathbf{x} \in \mathcal{X}$  and  $t \in [T]$ , we have*

$$\begin{aligned}\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x} \rangle &\leq \frac{1}{\alpha} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_*^2 \\ &\quad + \left( \mathcal{D}_{\psi_t}(\mathbf{x}, \hat{\mathbf{x}}_t) - \mathcal{D}_{\psi_t}(\mathbf{x}, \hat{\mathbf{x}}_{t+1}) \right) - \left( \mathcal{D}_{\psi_t}(\hat{\mathbf{x}}_{t+1}, \mathbf{x}_t) + \mathcal{D}_{\psi_t}(\mathbf{x}_t, \hat{\mathbf{x}}_t) \right),\end{aligned}$$

where we set  $\nabla f_0(\mathbf{x}_0) = 0$ .

Given that Theorem 1 performs optimistic OMD on individual functions  $\{f_1, \dots, f_T\}$ , we utilize Lemma 2 as  $\psi_t(\mathbf{x}) = \frac{1}{2\eta_t}\|\mathbf{x}\|_2^2$  is  $\frac{1}{\eta_t}$ -strongly convex with respect to  $\|\cdot\|_2$  and sum the inequality over  $t = 1, \dots, T$ :

$$\begin{aligned} \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle &\leq \underbrace{\sum_{t=1}^T \frac{1}{2\eta_t} (\|\mathbf{u} - \hat{\mathbf{x}}_t\|_2^2 - \|\mathbf{u} - \hat{\mathbf{x}}_{t+1}\|_2^2)}_{\text{term (a)}} \\ &+ \underbrace{\sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2}_{\text{term (b)}} - \underbrace{\sum_{t=1}^T \frac{1}{2\eta_t} (\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_2^2 + \|\hat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2)}_{\text{term (c)}}. \end{aligned} \quad (16)$$

In the following, we will bound the three terms on the right hand respectively.

First, given  $\eta_t = D/\sqrt{\delta + 4G^2 + \bar{V}_{t-1}}$  and  $\bar{V}_{t-1} = \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1})\|_2^2$ , we derive that  $\eta_t \leq D/\sqrt{\delta + \bar{V}_t}$  using Assumption 1 (boundedness of gradient norms). For term (a), by the fact  $\eta_t \leq \eta_{t-1}$  and Assumption 2 (domain boundedness), we have

$$\begin{aligned} \text{term (a)} &= \frac{1}{2\eta_1} \|\mathbf{u} - \hat{\mathbf{x}}_1\|_2^2 + \frac{1}{2} \sum_{t=2}^T \left( \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \|\mathbf{u} - \hat{\mathbf{x}}_t\|_2^2 - \frac{1}{2\eta_T} \|\mathbf{u} - \hat{\mathbf{x}}_{T+1}\|_2^2 \\ &\leq \frac{1}{2\eta_1} D^2 + \frac{1}{2} \sum_{t=2}^T \left( \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) D^2 = \frac{D^2}{2\eta_T} = \frac{D}{2} \sqrt{\delta + 4G^2 + \bar{V}_{T-1}}. \end{aligned}$$

For term (b), we utilize Lemma 11 to bound it as

$$\text{term (b)} \leq \sum_{t=1}^T \frac{D}{\sqrt{\delta + \bar{V}_t}} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \leq 2D\sqrt{\delta + \bar{V}_T}.$$

For term (c), we rely on the fact that  $\eta_t \leq \frac{D}{\sqrt{\delta}}$ :

$$\begin{aligned} \text{term (c)} &= \sum_{t=1}^T \frac{1}{2\eta_t} (\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_2^2 + \|\hat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2) \geq \frac{\sqrt{\delta}}{2D} \sum_{t=1}^T (\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_2^2 + \|\hat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2) \\ &\geq \frac{\sqrt{\delta}}{2D} \sum_{t=2}^T (\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_2^2 + \|\hat{\mathbf{x}}_t - \mathbf{x}_{t-1}\|_2^2) \geq \frac{\sqrt{\delta}}{4D} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2. \end{aligned}$$

Then we substitute the three bounds above into (16) and use Assumption 1 to get

$$\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \leq \frac{5D}{2} \sqrt{\delta + 4G^2 + \bar{V}_{T-1}} - \frac{\sqrt{\delta}}{4D} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2,$$

In order to bound the  $\bar{V}_{T-1}$  term, we incorporate a crucial lemma extracted from the analysis of Sachs et al. (2022). Refer to Appendix A.4 for the proof.

**Lemma 3** (boundedness of cumulative norm of gradient difference (Sachs et al., 2022, Analysis of Theorem 5)). *Under Assumptions 1 and 4, we have*

$$\begin{aligned} \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 &\leq G^2 + 4L^2 \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \\ &\quad + 8 \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + 4 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2. \end{aligned} \quad (17)$$

As a result, by applying Lemma 3, we have

$$\begin{aligned} &\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \\ &\leq \frac{5D}{2} \sqrt{\delta + 5G^2} + 5\sqrt{2}D \sqrt{\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2} + 5DL \sqrt{\sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2} \\ &\quad + 5D \sqrt{\sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2} - \frac{\sqrt{\delta}}{4D} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \\ &\leq \frac{5D}{2} \sqrt{\delta + 5G^2} + \frac{25D^3L^2}{\sqrt{\delta}} + 5\sqrt{2}D \sqrt{\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2} \\ &\quad + 5D \sqrt{\sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2} \end{aligned}$$

where the second step uses AM-GM inequality as  $5DL \sqrt{\sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2} \leq \frac{25D^3L^2}{\sqrt{\delta}} + \frac{\sqrt{\delta}}{4D} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2$ . Taking expectations and applying Jensen's inequality lead to

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \right] &\leq \frac{5D}{2} \sqrt{\delta} + \frac{25D^3L^2}{\sqrt{\delta}} + \frac{5\sqrt{5}DG}{2} + 5\sqrt{2}D \sqrt{\sigma_{1:T}^2} + 5D \sqrt{\Sigma_{1:T}^2} \\ &= 5\sqrt{10}D^2L + \frac{5\sqrt{5}DG}{2} + 5\sqrt{2}D \sqrt{\sigma_{1:T}^2} + 5D \sqrt{\Sigma_{1:T}^2} \\ &= \mathcal{O} \left( \sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right), \end{aligned}$$

where we set  $\delta = 10D^2L^2$  and recall definitions of  $\sigma_{1:T}^2$  in (3) and  $\Sigma_{1:T}^2$  in (4). We end the proof by noting the expectation upper-bounds the expected regret as in (15).  $\blacksquare$



## 3.6.2 PROOF OF THEOREM 3

**Proof** Since the expected functions are  $\lambda$ -strongly convex now, we have  $F_t(\mathbf{x}_t) - F_t(\mathbf{u}) \leq \langle \nabla F_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\lambda}{2} \|\mathbf{u} - \mathbf{x}_t\|_2^2$ . Then by the definition  $F_t(\mathbf{x}) = \mathbb{E}_{f_t \sim \mathfrak{D}_t}[f_t(\mathbf{x})]$ , we obtain

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] &= \mathbb{E} \left[ \sum_{t=1}^T F_t(\mathbf{x}_t) - \sum_{t=1}^T F_t(\mathbf{u}) \right] \\ &\leq \mathbb{E} \left[ \sum_{t=1}^T \left( \langle \nabla F_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\lambda}{2} \|\mathbf{u} - \mathbf{x}_t\|_2^2 \right) \right] = \mathbb{E} \left[ \sum_{t=1}^T \left( \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\lambda}{2} \|\mathbf{u} - \mathbf{x}_t\|_2^2 \right) \right]. \end{aligned} \quad (18)$$

Similar to the analysis of Theorem 1, we have the following regret upper bound,

$$\begin{aligned} &\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\lambda}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_2^2 \\ &\leq \underbrace{\sum_{t=1}^T \left( \frac{1}{2\eta_t} \|\mathbf{u} - \widehat{\mathbf{x}}_t\|_2^2 - \frac{1}{2\eta_t} \|\mathbf{u} - \widehat{\mathbf{x}}_{t+1}\|_2^2 \right)}_{\text{term (a)}} - \frac{\lambda}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_2^2 \\ &\quad + \underbrace{\sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2}_{\text{term (b)}} - \underbrace{\sum_{t=1}^T \frac{1}{2\eta_t} (\|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2 + \|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2)}_{\text{term (c)}}. \end{aligned} \quad (19)$$

We then provide the upper bounds of term (a), term (b), and term (c) respectively.

To bound term (a), we need the following classic lemma.

**Lemma 4** (stability lemma (Chiang et al., 2012, Proposition 7)). *Consider the following two updates: (i)  $\mathbf{x}_* = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{a}, \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{c})$ , and (ii)  $\mathbf{x}'_* = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{a}', \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{c})$ . When the regularizer  $\psi : \mathcal{X} \rightarrow \mathbb{R}$  is a 1-strongly convex function with respect to the norm  $\|\cdot\|$ , we have  $\|\mathbf{x}_* - \mathbf{x}'_*\| \leq \|(\nabla \psi(\mathbf{c}) - \mathbf{a}) - (\nabla \psi(\mathbf{c}) - \mathbf{a}')\|_* = \|\mathbf{a} - \mathbf{a}'\|_*$ .*

Using Lemma 4 with our algorithm yields  $\|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2 \leq \eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2$ . Considering Assumption 2 (domain boundedness) and the step size  $\eta_t = \frac{2}{\lambda t}$ , we obtain

$$\begin{aligned} \text{term (a)} &\leq \frac{1}{2\eta_1} D^2 + \frac{1}{2} \sum_{t=2}^T \left( \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \|\mathbf{u} - \widehat{\mathbf{x}}_t\|_2^2 - \frac{\lambda}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_2^2 \\ &\leq \frac{\lambda D^2}{4} + \frac{\lambda}{4} \sum_{t=1}^{T-1} (\|\mathbf{u} - \widehat{\mathbf{x}}_{t+1}\|_2^2 - 2\|\mathbf{u} - \mathbf{x}_t\|_2^2) \leq \frac{\lambda D^2}{4} + \frac{\lambda}{2} \sum_{t=1}^{T-1} \|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 \\ &\leq \frac{\lambda D^2}{4} + \frac{\lambda \eta_1}{2} \sum_{t=1}^{T-1} \eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \leq \frac{\lambda D^2}{4} + \text{term (b)}, \end{aligned}$$

where the last step is based on  $\eta_t$  being non-increasing. This shows that the upper bound of term (a) depends on term (b). For term (b), after inserting the definition of  $\eta_t$ , we get

$$\text{term (b)} = 2 \sum_{t=1}^T \frac{1}{\lambda t} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2.$$

Making use of the fact that  $\eta_t$  is non-increasing again, we bound term (c) by

$$\text{term (c)} \geq \sum_{t=2}^T \left( \frac{1}{2\eta_t} \|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2 + \frac{1}{2\eta_{t-1}} \|\widehat{\mathbf{x}}_t - \mathbf{x}_{t-1}\|_2^2 \right) \geq \sum_{t=2}^T \frac{1}{4\eta_{t-1}} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2.$$

Combining the upper bounds of term (a), term (b) and term (c) into (19) with  $\eta_t = \frac{2}{\lambda t}$  gives

$$\begin{aligned} & \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x} \rangle - \frac{\lambda}{2} \sum_{t=1}^T \|\mathbf{x} - \mathbf{x}_t\|_2^2 \\ & \leq \frac{\lambda D^2}{4} + 4 \sum_{t=1}^T \frac{1}{\lambda t} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 - \sum_{t=2}^T \frac{\lambda(t-1)}{8} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2. \end{aligned}$$

Then we need to use the following lemma with its proof in [Appendix A.4](#).

**Lemma 5** (boundedness of the norm of gradient difference ([Sachs et al., 2022](#), analysis of Theorem 5)). *Under Assumptions 4 and 1, we have*

$$\begin{aligned} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 & \leq 4\|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + 4\|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ & \quad + 4L^2\|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + 4\|\nabla F_{t-1}(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2, \end{aligned}$$

where  $\|\nabla f_1(\mathbf{x}_1) - \nabla f_0(\mathbf{x}_0)\|_2^2 = \|\nabla f_1(\mathbf{x}_1)\|_2^2 \leq G^2$ .

So applying Lemma 5 yields the following result,

$$\begin{aligned} & \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\lambda}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_2^2 \\ & \leq \frac{4G^2}{\lambda} + 4 \sum_{t=2}^T \frac{1}{\lambda t} (4\|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + 4\|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ & \quad + 4\|\nabla F_{t-1}(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2) + \sum_{t=2}^T \left( \frac{16L^2}{\lambda t} - \frac{\lambda(t-1)}{8} \right) \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + \frac{\lambda D^2}{4} \\ & \leq \frac{4G^2}{\lambda} + \sum_{t=2}^T \frac{16}{\lambda t} \|\nabla F_t(\mathbf{x}_t) - \nabla f_t(\mathbf{x}_t)\|_2^2 + \sum_{t=2}^T \frac{16}{\lambda t} \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ & \quad + \sum_{t=2}^T \frac{16}{\lambda(t-1)} \|\nabla F_{t-1}(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 + \sum_{t=1}^{T-1} \left( \frac{16L^2}{\lambda t} - \frac{\lambda t}{8} \right) \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2^2 + \frac{\lambda D^2}{4} \\ & \leq \frac{4G^2}{\lambda} + \sum_{t=1}^T \frac{32}{\lambda t} \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + \sum_{t=2}^T \frac{16}{\lambda t} \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ & \quad + \sum_{t=1}^{T-1} \left( \frac{16L^2}{\lambda t} - \frac{\lambda t}{8} \right) \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2^2 + \frac{\lambda D^2}{4}. \end{aligned} \tag{20}$$

Following [Sachs et al. \(2022\)](#), we define  $\kappa = \frac{L}{\lambda}$ . Then for  $t \geq 8\sqrt{2}\kappa$ , we have  $\frac{16L^2}{\lambda t} - \frac{\lambda t}{8} \leq 0$ . Using Assumption 2 (domain boundedness), the fourth term above is bounded as

$$\begin{aligned} \sum_{t=1}^{T-1} \left( \frac{16L^2}{\lambda t} - \frac{\lambda t}{8} \right) \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2^2 &\leq \sum_{t=1}^{\lceil 8\sqrt{2}\kappa \rceil} \left( \frac{16L^2}{\lambda t} - \frac{\lambda t}{8} \right) D^2 \leq \frac{16L^2 D^2}{\lambda} \sum_{t=1}^{\lceil 8\sqrt{2}\kappa \rceil} \frac{1}{t} \\ &\leq \frac{16L^2 D^2}{\lambda} \left( 1 + \int_{t=1}^{\lceil 8\sqrt{2}\kappa \rceil} \frac{1}{t} dt \right) = \frac{16L^2 D^2}{\lambda} \ln \left( 1 + 8\sqrt{2} \frac{L}{\lambda} \right) + \frac{16L^2 D^2}{\lambda}. \end{aligned}$$

Combining the above two formulas and taking the expectation, we can get that

$$\begin{aligned} &\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\lambda}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_2^2 \right] \\ &\leq \mathbb{E} \left[ \sum_{t=1}^T \frac{32}{\lambda t} \sigma_t^2 + \sum_{t=2}^T \frac{16}{\lambda t} \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right] + \frac{16L^2 D^2}{\lambda} \ln \left( 1 + 8\sqrt{2} \frac{L}{\lambda} \right) \\ &\quad + \frac{16L^2 D^2 + 4G^2}{\lambda} + \frac{\lambda D^2}{4}, \end{aligned}$$

where  $\sigma_t^2 = \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{f_t \sim \mathcal{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2]$  as defined in (2). To deal with the first term, we introduce a new lemma below, with its proof in [Appendix A.4](#).

**Lemma 6.** *Under Assumption 3, we have*

$$\begin{aligned} &\sum_{t=1}^T \frac{1}{\lambda t} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \\ &\leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \ln \left( \sum_{t=1}^T \frac{1}{2\sigma_{\max}^2 + \Sigma_{\max}^2} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) + 1 \right) + \frac{4\sigma_{\max}^2 + 2\Sigma_{\max}^2}{\lambda}. \end{aligned}$$

Then, we can arrive at

$$\begin{aligned} &\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\lambda}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_2^2 \right] \\ &\leq \frac{32\sigma_{\max}^2 + 16\Sigma_{\max}^2}{\lambda} \ln \left( \frac{1}{2\sigma_{\max}^2 + \Sigma_{\max}^2} (2\sigma_{1:T}^2 + \Sigma_{1:T}^2) + 1 \right) + \frac{64\sigma_{\max}^2 + 32\Sigma_{\max}^2}{\lambda} \\ &\quad + \frac{16L^2 D^2}{\lambda} \ln \left( 1 + 8\sqrt{2} \frac{L}{\lambda} \right) + \frac{16L^2 D^2 + 4G^2}{\lambda} + \frac{\lambda D^2}{4} \\ &= \mathcal{O} \left( \frac{1}{\lambda} (\sigma_{\max}^2 + \Sigma_{\max}^2) \log \left( (\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2) \right) \right). \end{aligned}$$

This ends the proof. ■

## 3.6.3 PROOF OF THEOREM 5

**Proof** Due to the exp-concavity assumption and Lemma 1, we have  $f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \leq \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\beta}{2} \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2$ , where  $\beta = \frac{1}{2} \min \left\{ \frac{1}{4GD}, \alpha \right\}$ , and  $h_t = \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\top$ . Therefore, we can get tighter regret bounds for exp-concave functions as follows

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2 \right]. \quad (21)$$

Clearly,  $\psi_t(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_{H_t}^2$  is a 1-strongly convex function with respect to  $\|\cdot\|_{H_t}$ , and  $\|\cdot\|_{H_t}^{-1}$  is the dual norm. Thus, from Lemma 2 (Variant of Bregman proximal inequality), we have

$$\begin{aligned} & \sum_{t=1}^T \langle \mathbf{x}_t - \mathbf{u}, \nabla f_t(\mathbf{x}_t) \rangle - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2 \\ & \leq \underbrace{\sum_{t=1}^T \left( \frac{1}{2} \|\mathbf{u} - \hat{\mathbf{x}}_t\|_{H_t}^2 - \frac{1}{2} \|\mathbf{u} - \hat{\mathbf{x}}_{t+1}\|_{H_t}^2 \right)}_{\text{term (a)}} - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2 \\ & \quad + \underbrace{\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_{H_t^{-1}}^2}_{\text{term (b)}} - \underbrace{\sum_{t=1}^T \frac{1}{2} (\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_{H_t}^2 + \|\hat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_{H_t}^2)}_{\text{term (c)}}. \end{aligned} \quad (22)$$

Then, we discuss the upper bounds of term (a), term (b) and term (c), respectively. According to Chiang et al. (2012, Proof of Lemma 14), we write term (a) as

$$\frac{1}{2} \left( \|\mathbf{u} - \hat{\mathbf{x}}_1\|_{H_1}^2 - \|\mathbf{u} - \hat{\mathbf{x}}_{T+1}\|_{H_{T+1}}^2 + \sum_{t=1}^T (\|\mathbf{u} - \hat{\mathbf{x}}_{t+1}\|_{H_{t+1}}^2 - \|\mathbf{u} - \hat{\mathbf{x}}_{t+1}\|_{H_t}^2) \right) - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2.$$

Based on Assumption 2 (domain boundedness) and Assumption 1 (boundedness of gradient norms), with the definition that  $H_t = I + \frac{\beta}{2} G^2 I + \frac{\beta}{2} \sum_{\tau=1}^{t-1} \nabla f_\tau(\mathbf{x}_\tau) \nabla f_\tau(\mathbf{x}_\tau)^\top$  and  $h_t = \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\top$ , we have  $\|\mathbf{u} - \hat{\mathbf{x}}_1\|_{H_1}^2 \leq D^2 (1 + \frac{\beta}{2} G^2)$  and  $H_{t+1} - H_t = \frac{\beta}{2} h_t$  for every  $t$ . So we can simplify term (a) to

$$\begin{aligned} \text{term (a)} & \leq \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) + \frac{\beta}{4} \sum_{t=1}^T \|\mathbf{u} - \hat{\mathbf{x}}_{t+1}\|_{h_t}^2 - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2 \\ & \leq \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) + \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{x}_t - \hat{\mathbf{x}}_{t+1}\|_{h_t}^2 \leq \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) + \sum_{t=1}^T \|\mathbf{x}_t - \hat{\mathbf{x}}_{t+1}\|_{H_t}^2 \\ & \leq \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_{H_t^{-1}}^2 = \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) + \text{term (b)}, \end{aligned}$$

where we use  $H_t \succeq \frac{\beta}{2} G^2 I \succeq \frac{\beta}{2} h_t$  for the third inequality and Lemma 4 (Stability lemma) in the fourth inequality. Notably, the upper bound of term (b) determines that of term (a).

Hence we move to bound term (b). By definition of  $H_t$ , there is  $G^2 I \succeq \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\top$  for every  $t$ . In addition, we know  $\nabla f_0(\mathbf{x}_0) = 0$ , so

$$H_t \succeq I + \frac{\beta}{4} \sum_{\tau=1}^t \left( \nabla f_\tau(\mathbf{x}_\tau) \nabla f_\tau(\mathbf{x}_\tau)^\top + \nabla f_{\tau-1}(\mathbf{x}_{\tau-1}) \nabla f_{\tau-1}(\mathbf{x}_{\tau-1})^\top \right). \quad (23)$$

Similar to [Chiang et al. \(2012\)](#), we claim that

$$\begin{aligned} & \nabla f_\tau(\mathbf{x}_\tau) \nabla f_\tau(\mathbf{x}_\tau)^\top + \nabla f_{\tau-1}(\mathbf{x}_{\tau-1}) \nabla f_{\tau-1}(\mathbf{x}_{\tau-1})^\top \\ & \succeq \frac{1}{2} (\nabla f_\tau(\mathbf{x}_\tau) - \nabla f_{\tau-1}(\mathbf{x}_{\tau-1})) (\nabla f_\tau(\mathbf{x}_\tau) - \nabla f_{\tau-1}(\mathbf{x}_{\tau-1}))^\top. \end{aligned} \quad (24)$$

The above inequality comes from subtracting the RHS of it from the left and getting that  $\frac{1}{2} (\nabla f_\tau(\mathbf{x}_\tau) + \nabla f_{\tau-1}(\mathbf{x}_{\tau-1})) (\nabla f_\tau(\mathbf{x}_\tau) + \nabla f_{\tau-1}(\mathbf{x}_{\tau-1}))^\top \succeq 0$ . Based on this, we obtain

$$H_t \stackrel{(24)}{\succeq} I + \frac{\beta}{8} \sum_{\tau=1}^t (\nabla f_\tau(\mathbf{x}_\tau) - \nabla f_{\tau-1}(\mathbf{x}_{\tau-1})) (\nabla f_\tau(\mathbf{x}_\tau) - \nabla f_{\tau-1}(\mathbf{x}_{\tau-1}))^\top.$$

Let  $P_t = I + \frac{\beta}{8} \sum_{\tau=1}^t (\nabla f_\tau(\mathbf{x}_\tau) - \nabla f_{\tau-1}(\mathbf{x}_{\tau-1})) (\nabla f_\tau(\mathbf{x}_\tau) - \nabla f_{\tau-1}(\mathbf{x}_{\tau-1}))^\top$ , we have

$$\begin{aligned} \text{term (b)} & \leq \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_{P_t^{-1}}^2 = \frac{8}{\beta} \sum_{t=1}^T \left\| \sqrt{\frac{\beta}{8}} (\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})) \right\|_{P_t^{-1}}^2 \\ & \leq \frac{8d}{\beta} \ln \left( \frac{\beta}{8d} \bar{V}_T + 1 \right), \end{aligned}$$

where we apply Lemma 14 with  $\mathbf{u}_t = \sqrt{\frac{\beta}{8}} (\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1}))$  and  $\varepsilon = 1$ .

Then, derived from the fact that  $H_t \succeq H_{t-1} \succeq I$ , we can bound term (c) as

$$\begin{aligned} \text{term (c)} & = \frac{1}{2} \sum_{t=1}^T \|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_{H_t}^2 + \frac{1}{2} \sum_{t=2}^{T+1} \|\mathbf{x}_{t-1} - \hat{\mathbf{x}}_t\|_{H_{t-1}}^2 \\ & \geq \frac{1}{2} \sum_{t=2}^T \|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_{H_{t-1}}^2 + \frac{1}{2} \sum_{t=2}^T \|\mathbf{x}_{t-1} - \hat{\mathbf{x}}_t\|_{H_{t-1}}^2 \geq \frac{1}{4} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2. \end{aligned}$$

Combining the above bounds of term (a), term (b) and term (c), we can get

$$\begin{aligned} & \sum_{t=1}^T \langle \mathbf{x}_t - \mathbf{u}, \nabla f_t(\mathbf{x}_t) \rangle - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2 \\ & \leq \frac{16d}{\beta} \ln \left( \frac{\beta}{8d} \bar{V}_T + 1 \right) + \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) - \frac{1}{4} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2. \end{aligned}$$

Further exploiting Lemma 3 (Boundedness of cumulative norm of gradient difference) with the inequality  $\ln(1+u+v) \leq \ln(1+u) + \ln(1+v)$  ( $u, v > 0$ ), we have

$$\sum_{t=1}^T \langle \mathbf{x}_t - \mathbf{u}, \nabla f_t(\mathbf{x}_t) \rangle - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2$$

$$\begin{aligned}
&\leq \frac{16d}{\beta} \ln \left( \frac{\beta}{d} \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + \frac{\beta}{2d} \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 + \frac{\beta}{8d} G^2 + 1 \right) \\
&\quad + \frac{16d}{\beta} \ln \left( \frac{\beta L^2}{2d} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + 1 \right) + \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) - \frac{1}{4} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \\
&\leq \frac{16d}{\beta} \ln \left( \frac{\beta}{d} \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + \frac{\beta}{2d} \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 + \frac{\beta}{8d} G^2 + 1 \right) \\
&\quad + \frac{16d}{\beta} \ln (32L^2 + 1) + \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right),
\end{aligned}$$

where the last step is due to Lemma 9.

Taking the expectation, and making use of Jensen's inequality, the above bound becomes

$$\begin{aligned}
&\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle - \frac{\beta}{2} \sum_{t=1}^T \|\mathbf{u} - \mathbf{x}_t\|_{h_t}^2 \right] \\
&\leq \frac{16d}{\beta} \ln \left( \frac{\beta}{d} \sigma_{1:T}^2 + \frac{\beta}{2d} \Sigma_{1:T}^2 + \frac{\beta}{8d} G^2 + 1 \right) + \frac{16d}{\beta} \ln (32L^2 + 1) + \frac{D^2}{2} \left( 1 + \frac{\beta}{2} G^2 \right) \\
&= \mathcal{O} \left( \frac{d}{\alpha} \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2) \right)
\end{aligned}$$

We finish the proof by integrating the above inequality to (21). ■

## 4 Extensions: Dynamic Regret Minimization and Non-smooth Functions

In this section, we investigate a new measure for the SEA model—dynamic regret, a more suitable metric for non-stationary environments. Subsequently, we explore the SEA model for non-smooth loss functions, proposing algorithms for minimizing static regret and dynamic regret respectively. Detailed analysis and proofs are placed in Section 4.4.

### 4.1 Dynamic Regret Minimization

To optimize the expected dynamic regret in (5), following the recent study of non-stationary online learning (Zhao et al., 2024), we develop a two-layer online ensemble method building upon the optimistic OMD framework. Our method consists of a meta-learner running over a group of base-learners, with full procedures summarized in Algorithm 2. Specifically, we maintain a pool for candidate step sizes  $\mathcal{H} = \{\eta_i = c \cdot 2^i \mid i \in [N]\}$ , where  $N$  is the number of base-learners and is of order  $\mathcal{O}(\log T)$ ,  $c$  is some small constant that will be given later. We denote by  $\mathcal{B}_i$  the  $i$ -th base-learner for  $i \in [N]$ . At round  $t \in [T]$ , the online learner obtains the decision  $\mathbf{x}_t$  by aggregating local base decisions via the meta-learner, namely,  $\mathbf{x}_t = \sum_{i=1}^N p_{t,i} \mathbf{x}_{t,i}$ , where  $\mathbf{x}_{t,i}$  is the decision returned by the base-learner  $\mathcal{B}_i$  for  $i \in [N]$  and  $\mathbf{p}_t \in \Delta_N$  is the weight vector returned by the meta-algorithm. The nature then chooses a distribution  $\mathfrak{D}_t$  and the individual function  $f_t(\cdot)$  is sampled from  $\mathfrak{D}_t$ . Subsequently, the online learner suffers the loss  $f_t(\mathbf{x}_t)$  and observes the gradient  $\nabla f_t(\mathbf{x}_t)$ .



For the base-learner  $\mathcal{B}_i$ , in each round  $t$ , she obtains her local decision  $\mathbf{x}_{t+1,i}$  by instantiating the optimistic OMD algorithm (see Algorithm 1) with  $\psi(\mathbf{x}) = \frac{1}{2\eta_i}\|\mathbf{x}\|_2^2$  and  $M_{t+1} = \nabla f_t(\mathbf{x}_t)$  over the linearized surrogate loss  $g_t(\mathbf{x}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle$ , where  $\eta_i \in \mathcal{H}$  is the step size associated with the  $i$ -th base-learner. Since  $\nabla g_t(\mathbf{x}_{t,i}) = \nabla f_t(\mathbf{x}_t)$ , the updating rules of  $\mathcal{B}_i$  are demonstrated as

$$\hat{\mathbf{x}}_{t+1,i} = \Pi_{\mathcal{X}}[\hat{\mathbf{x}}_{t,i} - \eta_i \nabla f_t(\mathbf{x}_t)], \quad \mathbf{x}_{t+1,i} = \Pi_{\mathcal{X}}[\hat{\mathbf{x}}_{t+1,i} - \eta_i \nabla f_t(\mathbf{x}_t)]. \quad (25)$$

The meta-learner updates the weight vector  $\mathbf{p}_{t+1} \in \Delta_N$  by optimistic hedge (Syrgkanis et al., 2015) with a time-varying learning rate  $\varepsilon_t$ , that is,

$$p_{t+1,i} \propto \exp\left(-\varepsilon_t \left(\sum_{s=1}^t \ell_{s,i} + m_{t+1,i}\right)\right), \quad (26)$$

where the feedback loss  $\ell_t \in \mathbb{R}^N$  is constructed by

$$\ell_{t,i} = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_{t,i} \rangle + \lambda \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 \quad (27)$$

for  $t \geq 2$  and  $\ell_{1,i} = \langle \nabla f_1(\mathbf{x}_1), \mathbf{x}_{1,i} \rangle$ ; and the optimism  $\mathbf{m}_{t+1} \in \mathbb{R}^N$  is constructed as

$$m_{t+1,i} = \langle M_{t+1}, \mathbf{x}_{t+1,i} \rangle + \lambda \|\mathbf{x}_{t+1,i} - \mathbf{x}_{t,i}\|_2^2 \quad (28)$$

with  $M_{t+1} = \nabla f_t(\mathbf{x}_t)$  for  $t \geq 2$  and  $M_1 = \mathbf{0}$ ;  $\lambda \geq 0$  being the coefficient of the correction terms; and we set  $\mathbf{x}_{0,i} = \mathbf{0}$  for  $i \in [N]$ . Note that the correction term  $\lambda \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2$  in the meta-algorithm (both feedback loss and optimism) plays an important role. Indeed, our algorithm design and regret analysis follow the collaborative online ensemble framework proposed by Zhao et al. (2024) for optimizing the gradient-variation dynamic regret. Technically, for such a two-layer structure, to cancel the additional positive term  $\sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2$  appearing in the derivation of  $\sigma_{1:T}^2$  and  $\Sigma_{1:T}^2$ , one needs to ensure an effective collaboration between the meta and base layers. This involves simultaneously exploiting negative terms of the regret upper bounds in both the base and meta layers as well as leveraging additional negative terms introduced by the above correction term.

**Remark 7.** After the submission of our conference paper, Sachs et al. (2022) released an updated version (Sachs et al., 2023), where they also utilized optimistic OMD to achieve the same dynamic regret as our approach. However, there is a significant difference between their method and ours. They employed an optimism design with  $m_{t,i} = \langle \nabla f_{t-1}(\bar{\mathbf{x}}_t), \mathbf{x}_{t,i} \rangle$ , based on another solution for gradient-variation dynamic regret of online convex optimization (Zhao et al., 2020), where  $\bar{\mathbf{x}}_t = \sum_{i=1}^N p_{t-1,i} \mathbf{x}_{t,i}$ . This design actually introduces a dependence issue in the SEA model because  $\bar{\mathbf{x}}_t$  depends on  $f_{t-1}(\cdot)$ . We provide more elaborations in Appendix B.1 and technical discussions in Remark 10.  $\triangleleft$

Below, we provide the dynamic regret upper bound of Algorithm 2 for the SEA model, and we will give the proof in Section 4.4.1.

**Theorem 7.** Under Assumptions 1, 2, 4 and 5, setting the step size pool  $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$  with  $\eta_i = \min\{1/(8L), \sqrt{(D^2/(8G^2T)) \cdot 2^{i-1}}\}$  and  $N = \lceil 2^{-1} \log_2(G^2T/(8L^2D^2)) \rceil + 1$ , and

**Algorithm 2** Dynamic Regret Minimization of the SEA Model

---

**Input:** step size pool  $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$ , learning rate of meta-algorithm  $\varepsilon_t > 0$ , correction coefficient  $\lambda > 0$

- 1: Initialization:  $\mathbf{x}_1 = \hat{\mathbf{x}}_1 \in \mathcal{X}$ ,  $\mathbf{p}_1 = \frac{1}{N} \cdot \mathbf{1}_N$
- 2: **for**  $t = 1$  **to**  $T$  **do**
- 3:   Submit the decision  $\mathbf{x}_t = \sum_{i=1}^N p_{t,i} \mathbf{x}_{t,i}$
- 4:   Observe the online function  $f_t : \mathcal{X} \mapsto \mathbb{R}$  sampled from the underlying distribution  $\mathfrak{D}_t$  and suffer the loss  $f_t(\mathbf{x}_t)$
- 5:   Base-learner  $\mathcal{B}_i$  updates the local decision by optimistic OMD, see (25),  $\forall i \in [N]$
- 6:   Receive  $\mathbf{x}_{t+1,i}$  from base-learner  $\mathcal{B}_i$  for  $i \in [N]$
- 7:   Construct the feedback loss  $\ell_t \in \mathbb{R}^N$  and optimism  $\mathbf{m}_{t+1} \in \mathbb{R}^N$  by (27) and (28)
- 8:   Update the weight  $\mathbf{p}_{t+1} \in \Delta_N$  by optimistic Hedge in (26)
- 9: **end for**

---

setting the learning rate of meta-algorithm as  $\varepsilon_t = \min\{1/(8D^2L), \sqrt{(\ln N)/(D^2\bar{V}_t)}\}$  for all  $t \in [T]$ , Algorithm 2 ensures

$$\mathbb{E}[\mathbf{Reg}_T^d(\mathbf{u}_1, \dots, \mathbf{u}_T)] \leq \mathcal{O}\left(P_T + \sqrt{1 + P_T}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})\right)$$

for any comparator sequence  $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X}$ , where  $\bar{V}_t = \sum_{s=2}^t \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1})\|_2^2$  with  $\nabla f_0(\mathbf{x}_0)$  defined as  $\mathbf{0}$ , and  $P_T = \mathbb{E}[\sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2]$  is the path length of comparators.

**Remark 8.** As mentioned, the static regret studied in earlier sections is a special case of dynamic regret with a fixed comparator. As a consequence, Theorem 7 directly implies an  $\mathcal{O}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$  static regret bound by noticing that  $P_T = 0$  when comparing to a fixed benchmark, which recovers the result in Theorem 1. Moreover, Theorem 7 also recovers the  $\mathcal{O}(\sqrt{(1 + P_T + V_T)(1 + P_T)})$  gradient-variation bound of Zhao et al. (2020, 2024) for the adversarial setting and the minimax optimal  $\mathcal{O}(\sqrt{T(1 + P_T)})$  bound of Zhang et al. (2018) since  $\sigma_{1:T}^2 = 0$  and  $\Sigma_{1:T}^2 = V_T \leq 4G^2T$  in this case.  $\triangleleft$

We focus on the convex and smooth case, while for the strongly convex and exp-concave cases, current understandings of their dynamic regret are still far from complete (Baby and Wang, 2022). In particular, how to realize optimistic online learning in strongly convex/exp-concave dynamic regret minimization remains open. Lastly, we note that to the best of our knowledge, FTRL has not yet achieved the worst-case  $\mathcal{O}(\sqrt{T(1 + P_T)})$  dynamic regret (Zhang et al., 2018), let alone the gradient-variation bound. In fact, FTRL is more like a lazy update (Hazan, 2016), which seems unable to track a sequence of changing comparators. We found that Jacobsen and Cutkosky (2022) have given preliminary results (in Theorem 2 and Theorem 3 of their work): *all the parameter-free FTRL-based algorithms we are aware of cannot achieve a dynamic regret bound better than  $\mathcal{O}(P_T\sqrt{T})$* . Although this cannot cover all the cases of FTRL-based algorithms on dynamic regret, it has at least shown that FTRL-based algorithms do have certain limitations in dynamic regret minimization.

## 4.2 SEA with Non-smooth Functions

The analysis in the previous section depends on the smoothness assumptions of expected functions (see Assumption 4). In this part, we further generalize the scope of the SEA model to the *non-smooth* functions. This is facilitated by the optimistic OMD framework again, but we replace gradient-descent updates with *implicit updates* in the optimistic step.

We consider the static regret minimization for the SEA model with convex and non-smooth functions. Assuming that all individual functions  $f_t(\cdot)$ 's are convex on  $\mathcal{X}$ , we update the decision  $\mathbf{x}_t$  by deploying optimistic OMD with  $\psi_t(\mathbf{x}) = \frac{1}{2\eta_t}\|\mathbf{x}\|_2^2$ , i.e.,

$$\widehat{\mathbf{x}}_{t+1} = \Pi_{\mathcal{X}} [\widehat{\mathbf{x}}_t - \eta_t \nabla f_t(\mathbf{x}_t)], \quad (29)$$

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} f_t(\mathbf{x}) + \frac{1}{2\eta_{t+1}} \|\mathbf{x} - \widehat{\mathbf{x}}_{t+1}\|_2^2, \quad (30)$$

where the update (30) is an implicit update and the step size is set as

$$\eta_t = \frac{D}{\sqrt{1 + 4G^2 + \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_s)\|_2^2}} \quad (31)$$

for  $t \in [T]$  (we define  $\eta_{T+1} = \eta_T$ ). Note that the second step (30) is crucial to remove the dependence on the smoothness of loss functions. Unlike the gradient-based update  $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}} [\widehat{\mathbf{x}}_{t+1} - \eta_{t+1} \nabla f_t(\mathbf{x}_t)]$  used in previous sections, it directly updates over the original function  $f_t(\mathbf{x})$  without linearization, so this is often referred to as “implicit update” (Campolongo and Orabona, 2020; Chen and Orabona, 2023; Bai et al., 2022).

Our algorithm can achieve a similar regret form as the smooth case scaling with the quantities  $\Sigma_{1:T}^2$  to reflect the adversarial difficulty and  $\tilde{\sigma}_{1:T}^2$  to indicate the stochastic aspect, where the variance quantity  $\tilde{\sigma}_{1:T}^2$  is defined as

$$\tilde{\sigma}_{1:T}^2 = \mathbb{E} \left[ \sum_{t=1}^T \tilde{\sigma}_t^2 \right], \text{ with } \tilde{\sigma}_t^2 = \mathbb{E}_{f_t \sim \mathcal{D}_t} \left[ \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2 \right]. \quad (32)$$

**Remark 9.** Note that  $\tilde{\sigma}_{1:T}^2$  also captures the stochastic difficulty of the SEA model due to the sample randomness. However, admittedly it is larger than  $\sigma_{1:T}^2$  because of the convex nature of the supremum operator. Despite this, we are unable to obtain any  $\sigma_{1:T}^2$ -type bound for the non-smooth case, and the technical discussions are deferred to Remark 10. It is crucial to highlight that later implications will demonstrate significant relevance of this quantity, particularly in real-world problems like online label shift (Section 5.7).  $\triangleleft$

Below we present the regret guarantee for SEA with *non-smooth* and convex functions. Refer to Section 4.4.2 for the proof.

**Theorem 8.** *Under Assumptions 1, 2 and 8, optimistic OMD with updates (29)–(32) enjoys the following guarantee:*

$$\mathbb{E}[\mathbf{Reg}_T(\mathbf{u})] \leq 5D\sqrt{1+G^2} + 10\sqrt{2}D\sqrt{\tilde{\sigma}_{1:T}^2} + 10D\sqrt{\Sigma_{1:T}^2} = \mathcal{O} \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right).$$

This bound is similar in form to the bound for the smooth case (Theorem 1), albeit with a slight loss in terms of the variance definition. However, in specific cases, this bound can be as good as the smooth case. For example, for fully adversarial OCO, we have  $\tilde{\sigma}_{1:T}^2 = \sigma_{1:T}^2 = 0$  since  $f_t(\cdot) = F_t(\cdot)$  for each  $t \in [T]$ . Moreover, when applying the result to the online label shift problem (see Section 5.7), using no matter  $\sigma_{1:T}^2$  or  $\tilde{\sigma}_{1:T}^2$  will deliver the same regret guarantee that scales with meaningful quantities for online label shift, the detailed analysis of which will be provided in Section 5.7 and Remark 13.

### 4.3 SEA with Non-smooth Functions: Dynamic Regret Minimization

We further investigate the dynamic regret of SEA with non-smooth and convex functions. To minimize the dynamic regret, we still employ a two-layer online ensemble structure based on the optimistic OMD framework as in Section 4.1, but with *implicit updates* in the base learners and additional ingredients for the design of meta learner.

Specifically, we construct a step size pool  $\mathcal{H} = \{\eta_i = c \cdot 2^i \mid i \in [N]\}$  to cover the (approximate) optimal step size, where  $N = \mathcal{O}(\log T)$  is the number of candidate step sizes and  $c$  is a constant given later. Then we maintain a meta-learner running over a group of base-learners  $\{\mathcal{B}_i\}_{i \in [N]}$ , each associated with a candidate step size  $\eta_i$  from the pool  $\mathcal{H}$ . The main procedure is summarized in Algorithm 3.

Consistent with the learner for static regret, each base-learner  $\mathcal{B}_i$  here performs the optimistic OMD algorithm parallelly with  $\psi(\mathbf{x}) = \frac{1}{2\eta_i} \|\mathbf{x}\|_2^2$  and an implicit update in the optimistic step. That means the updating rules of base-learner  $\mathcal{B}_i$  are

$$\hat{\mathbf{x}}_{t+1,i} = \Pi_{\mathcal{X}} [\hat{\mathbf{x}}_{t,i} - \eta_i \nabla f_t(\mathbf{x}_{t,i})], \quad \mathbf{x}_{t+1,i} = \arg \min_{\mathbf{x} \in \mathcal{X}} f_t(\mathbf{x}) + \frac{1}{2\eta_i} \|\mathbf{x} - \hat{\mathbf{x}}_{t+1,i}\|_2^2, \quad (33)$$

where  $\eta_i \in \mathcal{H}$  is the corresponding candidate step size and  $\mathbf{x}_{t+1,i}$  is the local decision.

Then the meta-learner collects local decisions and updates the weight  $\mathbf{p}_{t+1} \in \Delta_N$  by

$$p_{t+1,i} \propto \exp \left( -\varepsilon_t \left( \sum_{s=1}^t f_s(\mathbf{x}_{s,i}) + f_t(\mathbf{x}_{t+1,i}) \right) \right), \quad (34)$$

where  $p_{t+1,i}$  denotes the weight of the  $i$ -th base-learner and  $\varepsilon_t$  is the learning rate to be set later. After that, the online learner submits the decision  $\mathbf{x}_{t+1} = \sum_{i=1}^N p_{t+1,i} \mathbf{x}_{t+1,i}$  to the nature, and consequently suffers the loss  $f_{t+1}(\mathbf{x}_{t+1})$ , where  $f_{t+1}$  is sampled from the distribution  $\mathfrak{D}_{t+1}$  selected by the nature. Compared to Algorithm 2 in the smooth case, we no longer use surrogate losses and correction terms because we apply the technique of converting function variation into gradient variation without any negative term cancellations.

We have the following theoretical guarantee for Algorithm 3 with proof in Section 4.4.3.

**Theorem 9.** *Under Assumptions 1, 2 and 8, setting the step size pool  $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$  with  $\eta_i = (D/\sqrt{1+4TG^2}) \cdot 2^{i-1}$  and  $N = \lceil \frac{1}{2} \log((1+2T)(1+4TG^2)) \rceil + 1$ , and setting the learning rate of meta-algorithm as  $\varepsilon_t = 1/\sqrt{1 + \sum_{s=1}^t (\max_{i \in [N]} \{|\tilde{f}_s(\mathbf{x}_{s,i}) - \tilde{f}_{s-1}(\mathbf{x}_{s,i})|\})^2}$ , Algorithm 3 ensures*

$$\mathbb{E}[\mathbf{Reg}_T^d(\mathbf{u}_1, \dots, \mathbf{u}_T)] \leq \mathcal{O} \left( \sqrt{1 + P_T} \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right) \right),$$

which holds for any comparator sequence  $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X}$ .

**Algorithm 3** Dynamic Regret Minimization of SEA Model with Non-smooth Functions

---

**Input:** step size pool  $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$ , learning rate of meta-algorithm  $\varepsilon_t > 0$

- 1: Initialization:  $\mathbf{x}_1 = \hat{\mathbf{x}}_1 \in \mathcal{X}$ ,  $\mathbf{p}_1 = \frac{1}{N} \cdot \mathbf{1}_N$
- 2: **for**  $t = 1$  **to**  $T$  **do**
- 3:   Submit the decision  $\mathbf{x}_t = \sum_{i=1}^N p_{t,i} \mathbf{x}_{t,i}$
- 4:   Observe the online function  $f_t : \mathcal{X} \mapsto \mathbb{R}$  sampled from the underlying distribution  $\mathfrak{D}_t$  and suffer the loss  $f_t(\mathbf{x}_t)$
- 5:   Base-learner  $\mathcal{B}_i$  updates by optimistic OMD with implicit updates (33) for  $i \in [N]$
- 6:   Receive  $\mathbf{x}_{t+1,i}$  from base-learner  $\mathcal{B}_i$  for  $i \in [N]$
- 7:   Update the weight  $\mathbf{p}_{t+1} \in \Delta_N$  by optimistic Hedge in (34)
- 8: **end for**

---

**Remark 10.** Theorem 9 is not dependent on the smoothness of expected functions but is applicable to the smooth scenario as well. The  $\mathcal{O}(\sqrt{1 + P_T}(\sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}))$  bound detailed here and the  $\mathcal{O}(P_T + \sqrt{1 + P_T}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}))$  bound obtained under smoothness in Theorem 7 exhibit similar scaling in their corresponding variance quantities— $\tilde{\sigma}_{1:T}^2$  and  $\sigma_{1:T}^2$ , respectively. We emphasize that this work focuses on demonstrating that using implicit updates can theoretically achieve good regret bounds for non-smooth functions, and it is acknowledged that using implicit updates in the non-smooth case instead of using the first-order method in the smooth case may be more costly. Moreover, we argue that methods employing information about the function value  $f_t(\hat{\mathbf{x}})$  (or the gradient  $\nabla f_t(\hat{\mathbf{x}})$ ) where  $\hat{\mathbf{x}}$  is generated *afterward* the decision  $\mathbf{x}_t$  can hardly achieve regret bounds scaling with  $\sigma_{1:T}^2$ . This holds for the non-smooth part, as optimistic update steps of base-learners demand the full function information, and the meta-learner requires the value of  $f_t(\mathbf{x}_{t+1,i})$ . It also applies to the case of Sachs et al. (2023), who use the optimism design of Zhao et al. (2020) to optimize the dynamic regret of SEA with smooth functions. This would require the gradient  $\nabla f_{t-1}(\bar{\mathbf{x}}_t)$  with  $\bar{\mathbf{x}}_t = \sum_{i=1}^N p_{t-1,i} \mathbf{x}_{t,i}$  as mentioned in Remark 7, and can only obtain a weaker bound scaling with  $\tilde{\sigma}_{1:T}^2$ . We provide the details in Appendix B.1.  $\triangleleft$

#### 4.4 Analysis

In this section, we give the analysis of Theorem 7, Theorem 8 and Theorem 9 respectively, with some supplementary analysis and useful lemmas provided in Appendix B.

##### 4.4.1 PROOF OF THEOREM 7

This part presents the proof of Theorem 7. Since our algorithmic design is based on the collaborative online ensemble framework of Zhao et al. (2024), we first introduce the general theorem (Zhao et al., 2024, Theorem 9) and provide the proof for our theorem based on it.

**Theorem 10** (adaptation of Theorem 9 of Zhao et al. (2024)). *Under Assumption 1 (boundedness of gradient norms) and Assumption 2 (domain boundedness), setting the step*

size pool  $\mathcal{H}$  as

$$\mathcal{H} = \left\{ \eta_i = \min \left\{ \bar{\eta}, \sqrt{\frac{D^2}{8G^2T}} \cdot 2^{i-1} \right\} \mid i \in [N] \right\}, \quad (35)$$

where  $N = \lceil 2^{-1} \log_2((8G^2T\bar{\eta}^2)/D^2) \rceil + 1$ , and setting meta-algorithm's learning rate as

$$\varepsilon_t = \min \left\{ \bar{\varepsilon}, \sqrt{\frac{\ln N}{D^2 \sum_{s=1}^t \|\nabla f_s(\mathbf{x}_s) - f_{s-1}(\mathbf{x}_{s-1})\|_2^2}} \right\},$$

Algorithm 2 enjoys the following dynamic regret guarantee:

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u}_t \rangle \right] \\ & \leq 5\sqrt{D^2 \ln N \mathbb{E}[\bar{V}_T]} + 2\sqrt{(D^2 + 2DP_T) \mathbb{E}[\bar{V}_T]} + \mathbb{E} \left[ \frac{\ln N}{\bar{\varepsilon}} + 8\bar{\varepsilon}D^2G^2 + \frac{D^2 + 2DP_T}{\bar{\eta}} \right. \\ & \quad \left. + \left( \lambda - \frac{1}{4\bar{\eta}} \right) \sum_{t=2}^T \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 - \frac{1}{4\bar{\varepsilon}} \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2 - \lambda \sum_{t=2}^T \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 \right]. \end{aligned}$$

In the above,  $\bar{V}_T = \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2$  is the adaptivity term measuring the quality of optimistic gradient vectors  $\{M_t = \nabla f_{t-1}(\mathbf{x}_{t-1})\}_{t=1}^T$ , and  $P_T = \mathbb{E}[\sum_{t=2}^T \|\mathbf{u}_{t-1} - \mathbf{u}_t\|_2]$  is the path length of comparators.

**Remark 11.** Note that  $\mathbf{u}_1, \dots, \mathbf{u}_T$  may exhibit randomness in the SEA model, so the path length  $P_T$  we define is in the expected form. Consequently, we have introduced a subtle modification to Theorem 5 of Zhao et al. (2024), in which the expectation is taken before tuning the step size in its analysis.  $\triangleleft$

In the following, we prove Theorem 7 based on Theorem 10.

**Proof** [of Theorem 7] In Theorem 10, where  $\bar{V}_T = \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2$ , applying Lemma 3 (boundedness of cumulative norm of gradient difference) allows us to bound the first and second term as

$$\begin{aligned} & 5\sqrt{D^2 \ln N \mathbb{E}[\bar{V}_T]} + 2\sqrt{(D^2 + 2DP_T) \mathbb{E}[\bar{V}_T]} \\ & \leq G \left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) + \left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) \sqrt{4L^2 \mathbb{E} \left[ \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \right]} \\ & \quad + \left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) \left( 2\sqrt{2} \sqrt{\sigma_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right). \quad (36) \end{aligned}$$

To eliminate the relevant terms of  $\|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2$ , we first notice that

$$\|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 = \left\| \sum_{i=1}^N p_{t,i} \mathbf{x}_{t,i} - \sum_{i=1}^N p_{t-1,i} \mathbf{x}_{t-1,i} \right\|_2^2$$



$$\begin{aligned}
 &\leq 2 \left\| \sum_{i=1}^N p_{t,i} \mathbf{x}_{t,i} - \sum_{i=1}^N p_{t,i} \mathbf{x}_{t-1,i} \right\|_2^2 + 2 \left\| \sum_{i=1}^N p_{t,i} \mathbf{x}_{t-1,i} - \sum_{i=1}^N p_{t-1,i} \mathbf{x}_{t-1,i} \right\|_2^2 \\
 &\leq 2 \left( \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2 \right)^2 + 2 \left( \sum_{i=1}^N |p_{t,i} - p_{t-1,i}| \|\mathbf{x}_{t-1,i}\|_2 \right)^2 \\
 &\leq 2 \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 + 2D^2 \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2.
 \end{aligned}$$

Thus we get  $\sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \leq 2 \sum_{t=2}^T \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 + 2D^2 \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2$ . Then we can use it and the AM-GM inequality to bound the second term in (36):

$$\begin{aligned}
 &\left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) \sqrt{4L^2 \mathbb{E} \left[ \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \right]} \\
 &\leq 5 \sqrt{D^2 \ln N \left( 8L^2 \mathbb{E} \left[ \sum_{t=2}^T \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 \right] + 8L^2 D^2 \mathbb{E} \left[ \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2 \right] \right)} \\
 &\quad + 2 \sqrt{(D^2 + 2DP_T) \left( 8L^2 \mathbb{E} \left[ \sum_{t=2}^T \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 \right] + 8L^2 D^2 \mathbb{E} \left[ \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2 \right] \right)} \\
 &\leq \frac{25 \ln N}{4\bar{\varepsilon}} + \frac{D^2 + 2DP_T}{\bar{\eta}} + (8\bar{\varepsilon}D^2L^2 + 8\bar{\eta}L^2) \mathbb{E} \left[ \sum_{t=2}^T \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 \right] \\
 &\quad + (8\bar{\varepsilon}L^2D^4 + 8\bar{\eta}L^2D^2) \mathbb{E} \left[ \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2 \right].
 \end{aligned}$$

Combining (36) and the above formula with the regret in Theorem 10, we have

$$\begin{aligned}
 &\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u}_t \rangle \right] \\
 &\leq G \left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) \\
 &\quad + \left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) \left( 2\sqrt{2}\sqrt{\sigma_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) + \frac{2D^2 + 4DP_T}{\bar{\eta}} \\
 &\quad + \left( \lambda - \frac{1}{4\bar{\eta}} \right) \mathbb{E} \left[ \sum_{t=2}^T \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 \right] + \left( 8\bar{\varepsilon}L^2D^4 + 8\bar{\eta}L^2D^2 - \frac{1}{4\bar{\varepsilon}} \right) \mathbb{E} \left[ \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2 \right] \\
 &\quad + (8\bar{\varepsilon}D^2L^2 + 8\bar{\eta}L^2 - \lambda) \mathbb{E} \left[ \sum_{t=2}^T \sum_{i=1}^N p_{t,i} \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 \right] + \frac{29 \ln N}{4\bar{\varepsilon}} + 8\bar{\varepsilon}D^2G^2.
 \end{aligned}$$

Setting  $\lambda = 2L$ ,  $\bar{\eta} = \frac{1}{8L}$  and  $\bar{\varepsilon} = \frac{1}{8D^2L}$ , we can drop the last three non-positive terms to get

$$\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u}_t \rangle \right]$$

$$\begin{aligned}
&\leq G \left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) + \left( 5\sqrt{D^2 \ln N} + 2\sqrt{(D^2 + 2DP_T)} \right) \left( 2\sqrt{2}\sqrt{\sigma_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) \\
&\quad + (58 \ln N + 16)D^2L + 32DLP_T + \frac{1}{L}G^2 = \mathcal{O} \left( P_T + \sqrt{(1 + P_T)} \left( \sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right) \right),
\end{aligned} \tag{37}$$

which completes the proof.  $\blacksquare$

#### 4.4.2 PROOF OF THEOREM 8

Before giving proofs of the non-smooth case, for the sake of simplicity of the presentation, we first introduce the following notation:

$$\tilde{V}_T = \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2, \tag{38}$$

which adds a supremum operation before summing compared with  $V_T$ .

**Proof** Referring to (9) from the previous article, for convex random functions, we have:

$$\mathbb{E}[f_t(\mathbf{x}_t) - f_t(\mathbf{u})] \leq \mathbb{E}[\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle]. \tag{39}$$

We decompose the instantaneous loss above as

$$\begin{aligned}
&\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \\
&\leq \underbrace{\langle \nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_t), \mathbf{x}_t - \hat{\mathbf{x}}_{t+1} \rangle}_{\text{term (a)}} + \underbrace{\langle \nabla f_{t-1}(\mathbf{x}_t), \mathbf{x}_t - \hat{\mathbf{x}}_{t+1} \rangle}_{\text{term (b)}} + \underbrace{\langle \nabla f_t(\mathbf{x}_t), \hat{\mathbf{x}}_{t+1} - \mathbf{u} \rangle}_{\text{term (c)}}.
\end{aligned} \tag{40}$$

So we give the upper bounds of these three terms respectively in the following. For term (a), by Fenchel's inequality for the squared  $L_2$  norm, we have

$$\text{term (a)} \leq 2\eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_t)\|_2^2 + \frac{1}{2\eta_t} \|\mathbf{x}_t - \hat{\mathbf{x}}_{t+1}\|_2^2. \tag{41}$$

We introduce the following lemma to bound term (b), which is related to the implicit update procedure with the proof presented in [Appendix B.2](#). Note that to make the following lemma hold, we need the convexity of individual functions.

**Lemma 7.** *Let  $\hat{\mathbf{x}}_{t+1}$  and  $\mathbf{x}_{t+1}$  be defined as in (29) and (30). Then, for any  $\mathbf{x} \in \mathcal{X}$ ,*

$$\langle \nabla f_t(\mathbf{x}_{t+1}), \mathbf{x}_{t+1} - \mathbf{x} \rangle \leq \frac{1}{2\eta_{t+1}} \left( \|\mathbf{x} - \hat{\mathbf{x}}_{t+1}\|_2^2 - \|\mathbf{x} - \mathbf{x}_{t+1}\|_2^2 - \|\mathbf{x}_{t+1} - \hat{\mathbf{x}}_{t+1}\|_2^2 \right).$$

According to Lemma 7, we set  $\mathbf{x} = \hat{\mathbf{x}}_{t+1}$  and obtain

$$\text{term (b)} \leq \frac{1}{2\eta_t} \left( \|\hat{\mathbf{x}}_{t+1} - \hat{\mathbf{x}}_t\|_2^2 - \|\hat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 - \|\hat{\mathbf{x}}_t - \mathbf{x}_t\|_2^2 \right). \tag{42}$$

For term (c), we leverage Lemma 7 of [Zhao et al. \(2020\)](#) to get

$$\text{term (c)} \leq \frac{1}{2\eta_t} (\|\mathbf{u} - \hat{\mathbf{x}}_t\|_2^2 - \|\mathbf{u} - \hat{\mathbf{x}}_{t+1}\|_2^2 - \|\hat{\mathbf{x}}_t - \hat{\mathbf{x}}_{t+1}\|_2^2). \quad (43)$$

Combining the three upper bounds above and summing over  $t = 1, \dots, T$ , we have

$$\begin{aligned} & \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \\ & \leq \sum_{t=1}^T 2\eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_t)\|_2^2 + \sum_{t=1}^T \frac{1}{2\eta_t} (\|\mathbf{u} - \hat{\mathbf{x}}_t\|_2^2 - \|\mathbf{u} - \hat{\mathbf{x}}_{t+1}\|_2^2) + \frac{D^2}{2\eta_T} \\ & \leq \sum_{t=1}^T 2\eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_t)\|_2^2 + \frac{D^2}{2\eta_1} + \frac{D^2}{2} \sum_{t=2}^T \left( \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) + \frac{D^2}{2\eta_T} \\ & \leq \sum_{t=1}^T 2\eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_t)\|_2^2 + \frac{D^2}{\eta_T}, \end{aligned} \quad (44)$$

where we drop the negative term  $-\frac{1}{2\eta_t} \|\hat{\mathbf{x}}_t - \mathbf{x}_t\|_2^2$  to get the first inequality. Then we apply the inequality  $\eta_t \leq D/\sqrt{1 + \sum_{s=1}^t \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_s)\|_2^2}$  and Lemma 11 to obtain

$$\begin{aligned} \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle & \leq 2 \sum_{t=1}^T \frac{D}{\sqrt{1 + \sum_{s=1}^t \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_s)\|_2^2}} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_t)\|_2^2 \\ & \quad + D \sqrt{1 + 4G^2 + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_t)\|_2^2} \\ & \leq 5D \sqrt{1 + 4G^2 + \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2} \end{aligned}$$

Moreover, we develop a lemma to bound the  $\sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2$  term with its proof in [Appendix B.2](#).

**Lemma 8.** *Under Assumption 1, we have*

$$\begin{aligned} & \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2 \\ & \leq G^2 + 6 \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2 + 4 \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2. \end{aligned}$$

According to this lemma, we get that

$$\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \leq 5D\sqrt{1 + 5G^2} + 10\sqrt{2}D \sqrt{\sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2}$$

$$+ 10D \sqrt{\sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2}.$$

Taking expectations with Jensen's inequality and combining with (39), we arrive at

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] &\leq 5D\sqrt{1+5G^2} + 10\sqrt{2}D\sqrt{\tilde{\sigma}_{1:T}^2} + 10D\sqrt{\Sigma_{1:T}^2} \\ &= \mathcal{O} \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right), \end{aligned}$$

which ends the proof.  $\blacksquare$

#### 4.4.3 PROOF OF THEOREM 9

**Proof** For dynamic regret minimization based on Algorithm 3, we can decompose the expected dynamic regret into the *meta-regret* and *base-regret*:

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \right] = \underbrace{\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_{t,i}) \right]}_{\text{meta-regret}} + \underbrace{\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_{t,i}) - \sum_{t=1}^T f_t(\mathbf{u}_t) \right]}_{\text{base-regret}}. \quad (45)$$

The first part quantifies the cumulative loss difference between overall and base decisions, while the second part measures the dynamic regret of base-learner  $\mathcal{B}_i$ . This decomposition applies to any base-learner's index  $i \in [N]$ . We then present upper bounds for both terms.

**Bounding the meta-regret.** For the meta-regret, due to Jensen's inequality, we have

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_{t,i}) \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^N p_{t,j} f_t(\mathbf{x}_{t,j}) - \sum_{t=1}^T f_t(\mathbf{x}_{t,i}) \right].$$

By introducing the reference losses  $\tilde{f}_t(\mathbf{x}_{t,i}) = f_t(\mathbf{x}_{t,i}) - f_t(\mathbf{x}_{\text{ref}})$  and  $\tilde{f}_{t-1}(\mathbf{x}_{t,i}) = f_{t-1}(\mathbf{x}_{t,i}) - f_{t-1}(\mathbf{x}_{\text{ref}})$ , where  $\mathbf{x}_{\text{ref}}$  is an arbitrary reference point in  $\mathcal{X}$ , we can easily verify that

$$p_{t,i} = \frac{\exp \left( \varepsilon_t \left( \sum_{s=1}^{t-1} f_s(\mathbf{x}_{s,i}) + f_{t-1}(\mathbf{x}_{t,i}) \right) \right)}{\sum_{j=1}^N \exp \left( \varepsilon_t \left( \sum_{s=1}^{t-1} f_s(\mathbf{x}_{s,j}) + f_{t-1}(\mathbf{x}_{t,j}) \right) \right)} = \frac{\exp \left( \varepsilon_t \left( \sum_{s=1}^{t-1} \tilde{f}_s(\mathbf{x}_{s,i}) + \tilde{f}_{t-1}(\mathbf{x}_{t,i}) \right) \right)}{\sum_{j=1}^N \exp \left( \varepsilon_t \left( \sum_{s=1}^{t-1} \tilde{f}_s(\mathbf{x}_{s,j}) + \tilde{f}_{t-1}(\mathbf{x}_{t,j}) \right) \right)}.$$

That means the updating rule of  $\mathbf{p}_{t+1}$  for meta-learner in (34) can also be written as

$$p_{t+1,i} \propto \exp \left( -\varepsilon_t \left( \sum_{s=1}^t \tilde{f}_s(\mathbf{x}_{s,i}) + \tilde{f}_t(\mathbf{x}_{t+1,i}) \right) \right). \quad (46)$$

According to [Zhao et al. \(2024\)](#), the updating rule (46) which uses adaptive learning rate  $\varepsilon_t$  is identical to the optimistic FTRL algorithm which updates by

$$\mathbf{p}_{t+1} = \arg \min_{\mathbf{p} \in \Delta_N} \left\langle \mathbf{p}, \sum_{s=1}^t \boldsymbol{\ell}_s + \mathbf{m}_{t+1} \right\rangle + \psi_{t+1}(\mathbf{p})$$

with  $\psi_{t+1}(\mathbf{p}) = \frac{1}{\varepsilon_t} (\sum_{i=1}^N p_i \ln p_i + \ln N)$ , where the  $i$ -th component of  $\boldsymbol{\ell}_s$  is  $\ell_{s,i} = \tilde{f}_s(\mathbf{x}_{s,i})$  for  $i \in [N]$  and the  $i$ -th component of  $\mathbf{m}_{t+1}$  is  $m_{t+1,i} = \tilde{f}_t(\mathbf{x}_{t+1,i})$  for  $i \in [N]$  (this is easily proved by computing the closed-form solution). Then we can apply Lemma 15 and obtain

$$\begin{aligned} & \sum_{t=1}^T \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle - \sum_{t=1}^T \ell_{t,i} \\ & \leq \max_{\mathbf{p} \in \Delta} \psi_{T+1}(\mathbf{p}) + \sum_{t=1}^T \left( \langle \boldsymbol{\ell}_t - \mathbf{m}_t, \mathbf{p}_t - \mathbf{p}_{t+1} \rangle - \frac{1}{2\varepsilon_{t-1}} \|\mathbf{p}_t - \mathbf{p}_{t+1}\|_1^2 \right) \\ & \leq \frac{\ln N}{\varepsilon_T} + \sum_{t=1}^T \varepsilon_{t-1} \|\boldsymbol{\ell}_t - \mathbf{m}_t\|_\infty^2 + \frac{1}{4\varepsilon_{t-1}} \|\mathbf{p}_t - \mathbf{p}_{t+1}\|_1^2 - \frac{1}{2\varepsilon_{t-1}} \|\mathbf{p}_t - \mathbf{p}_{t+1}\|_1^2 \\ & \leq \frac{\ln N}{\varepsilon_T} + \sum_{t=1}^T \varepsilon_{t-1} \|\boldsymbol{\ell}_t - \mathbf{m}_t\|_\infty^2 = \frac{\ln N}{\varepsilon_T} + \sum_{t=1}^T \varepsilon_{t-1} \left( \max_{i \in [N]} \left\{ \left| \tilde{f}_t(\mathbf{x}_{t,i}) - \tilde{f}_{t-1}(\mathbf{x}_{t,i}) \right| \right\} \right)^2. \end{aligned}$$

Since  $\varepsilon_t = 1/\sqrt{1 + \sum_{s=1}^t \left( \max_{i \in [N]} \left\{ \left| \tilde{f}_s(\mathbf{x}_{s,i}) - \tilde{f}_{s-1}(\mathbf{x}_{s,i}) \right| \right\} \right)^2}$ , we have

$$\begin{aligned} & \sum_{t=1}^T \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle - \sum_{t=1}^T \ell_{t,i} \\ & \leq \frac{\ln N}{\varepsilon_T} + \sum_{t=1}^T \frac{\left( \max_{i \in [N]} \left\{ \left| \tilde{f}_t(\mathbf{x}_{t,i}) - \tilde{f}_{t-1}(\mathbf{x}_{t,i}) \right| \right\} \right)^2}{\sqrt{1 + \sum_{s=1}^{t-1} \left( \max_{i \in [N]} \left\{ \left| \tilde{f}_s(\mathbf{x}_{s,i}) - \tilde{f}_{s-1}(\mathbf{x}_{s,i}) \right| \right\} \right)^2}} \\ & \leq (\ln N + 4) \sqrt{1 + \sum_{t=1}^T \left( \max_{i \in [N]} \left\{ \left| \tilde{f}_t(\mathbf{x}_{t,i}) - \tilde{f}_{t-1}(\mathbf{x}_{t,i}) \right| \right\} \right)^2} + \max_{t \in [T]} \left( \max_{i \in [N]} \left\{ \left| \tilde{f}_t(\mathbf{x}_{t,i}) - \tilde{f}_{t-1}(\mathbf{x}_{t,i}) \right| \right\} \right)^2, \end{aligned}$$

where we exploit Lemma 12 in the second inequality. Next, to convert the function variation to the gradient variation, we define  $H_t(\mathbf{x}_{t,i}) = f_t(\mathbf{x}_{t,i}) - f_{t-1}(\mathbf{x}_{t,i})$  and get

$$\begin{aligned} \left| \tilde{f}_t(\mathbf{x}_{t,i}) - \tilde{f}_{t-1}(\mathbf{x}_{t,i}) \right| &= |H_t(\mathbf{x}_{t,i}) - H_t(\mathbf{x}_{\text{ref}})| = |\langle \nabla H_t(\boldsymbol{\xi}_{t,i}), \mathbf{x}_{t,i} - \mathbf{x}_{\text{ref}} \rangle| \\ &\leq D \|\nabla f_t(\boldsymbol{\xi}_{t,i}) - \nabla f_{t-1}(\boldsymbol{\xi}_{t,i})\|_2 \leq D \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2, \end{aligned}$$

where the second equality is due to the mean value theorem and  $\boldsymbol{\xi}_{t,i} = c_{t,i} \mathbf{x}_{t,i} + (1 - c_{t,i}) \mathbf{x}_{\text{ref}}$  with  $c_{t,i} \in [0, 1]$ . So by Assumption 1 (boundedness of gradient norms), we have

$$\sum_{t=1}^T \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle - \sum_{t=1}^T \ell_{t,i} \leq (\ln N + 4) \sqrt{1 + D^2 \tilde{V}_T} + 4G^4 \leq (\ln N + 4) D \sqrt{\tilde{V}_T} + 4G^4 + \ln N + 4.$$

Further combining the definitions of  $\ell_t$  and  $\ell_{t,i}$ , we finally get

$$\begin{aligned} \sum_{t=1}^T \sum_{j=1}^N p_{t,j} f_t(\mathbf{x}_{t,j}) - \sum_{t=1}^T f_t(\mathbf{x}_{t,i}) &= \sum_{t=1}^T \sum_{j=1}^N p_{t,j} \tilde{f}_t(\mathbf{x}_{t,j}) - \sum_{t=1}^T \tilde{f}_t(\mathbf{x}_{t,i}) = \sum_{t=1}^T \langle \mathbf{p}_t, \ell_t \rangle - \sum_{t=1}^T \ell_{t,i} \\ &\leq (\ln N + 4)D\sqrt{\tilde{V}_T} + 4G^4 + \ln N + 4. \end{aligned} \quad (47)$$

**Bounding the base-regret.** Owing to the convexity of individual functions, we have

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_{t,i}) - \sum_{t=1}^T f_t(\mathbf{u}_t) \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_{t,i}), \mathbf{x}_{t,i} - \mathbf{u}_t \rangle \right].$$

Similar to the non-smooth case of static regret, we can get the upper bound of the above instantaneous loss following the same arguments in obtaining (44):

$$\begin{aligned} &\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_{t,i}), \mathbf{x}_{t,i} - \mathbf{u}_t \rangle \right] \\ &\leq \mathbb{E} \left[ 2\eta_i \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla f_{t-1}(\mathbf{x}_{t,i})\|_2^2 + \frac{1}{2\eta_i} \sum_{t=2}^T \left( \|\mathbf{u}_t - \hat{\mathbf{x}}_{t,i}\|_2^2 - \|\mathbf{u}_{t-1} - \hat{\mathbf{x}}_{t,i}\|_2^2 \right) + \frac{D^2}{2\eta_i} \right] \\ &\leq 2\eta_i \mathbb{E} [\tilde{V}_T] + \mathbb{E} \left[ \frac{1}{2\eta_i} \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2 \|\mathbf{u}_t - \hat{\mathbf{x}}_{t,i} + \mathbf{u}_{t-1} - \hat{\mathbf{x}}_{t,i}\|_2 \right] + \frac{D^2}{2\eta_i} \\ &\leq 2\eta_i \mathbb{E} [\tilde{V}_T] + \frac{D^2 + 2DP_T}{2\eta_i}, \end{aligned}$$

where the second inequality comes from that  $\|\mathbf{u}_t - \hat{\mathbf{x}}_{t,i}\|_2^2 - \|\mathbf{u}_{t-1} - \hat{\mathbf{x}}_{t,i}\|_2^2 = \langle \mathbf{u}_t - \hat{\mathbf{x}}_{t,i} - (\mathbf{u}_{t-1} - \hat{\mathbf{x}}_{t,i}), \mathbf{u}_t - \hat{\mathbf{x}}_{t,i} + (\mathbf{u}_{t-1} - \hat{\mathbf{x}}_{t,i}) \rangle \leq \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2 \|\mathbf{u}_t - \hat{\mathbf{x}}_{t,i} + \mathbf{u}_{t-1} - \hat{\mathbf{x}}_{t,i}\|_2$ . Since we have

$$\tilde{V}_T \leq \sum_{t=1}^T \left( 2 \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x})\|_2^2 + 2 \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_{t-1}(\mathbf{x})\|_2^2 \right) \leq 4TG^2,$$

the optimal step size  $\eta^* = \frac{1}{2} \sqrt{\frac{D^2 + 2DP_T}{1 + \mathbb{E}[\tilde{V}_T]}}$  should lie in the range  $[\frac{1}{2} \sqrt{\frac{D^2}{1 + 4TG^2}}, \frac{1}{2} \sqrt{D^2 + 2D^2T}]$ .

Our designed step size pool is  $\mathcal{H} = \left\{ \frac{D}{\sqrt{1 + 4TG^2}} \cdot 2^{i-1} \mid i \in [N] \right\}$  with  $N = \lceil \frac{1}{2} \log((1 + 2T)(1 + 4TG^2)) \rceil + 1$ . There must be an  $\eta_{i^*} \in \mathcal{H}$  satisfying  $\eta_{i^*} \leq \eta^* \leq 2\eta_{i^*}$  and we can obtain that

$$\begin{aligned} &\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_{t,i}), \mathbf{x}_{t,i} - \mathbf{u}_t \rangle \right] \\ &\leq 2\eta_{i^*} \mathbb{E} [\tilde{V}_T] + \frac{D^2 + 2DP_T}{2\eta_{i^*}} \leq 2\eta^* \mathbb{E} [\tilde{V}_T] + \frac{D^2 + 2DP_T}{\eta^*} \leq 2\sqrt{2(D^2 + 2DP_T)\mathbb{E} [\tilde{V}_T]}. \end{aligned} \quad (48)$$

**Bounding the overall dynamic regret.** Combining the meta-regret (47) and the base-regret (48), we further obtain that

$$\begin{aligned}
& \mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \right] \\
& \leq \left( D(\ln N + 4) + 2\sqrt{2(D^2 + 2DP_T)} \right) \sqrt{\mathbb{E} [\tilde{V}_T]} + 4G^4 + \ln N + 4 \\
& \leq \left( D(\ln N + 4) + 2\sqrt{2(D^2 + 2DP_T)} \right) \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) + 4G^4 + \ln N + 4 \\
& = \mathcal{O} \left( \sqrt{1 + P_T} \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right) \right), \tag{49}
\end{aligned}$$

where we make use of Lemma 8 in the last inequality and finish the proof.  $\blacksquare$

## 5 Implications

In this section, first, we demonstrate how our results can be applied to recover the regret bound for adversarial data and the excess risk bound for stochastic data. Then, we discuss the implications for other intermediate examples.

We begin by listing two points followed by all the examples. First, for convex and smooth functions, we obtain the same  $\mathcal{O}(\sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2})$  bound as [Sachs et al. \(2022\)](#), so we will not repeat the analysis below unless necessary. But we emphasize that our result eliminates the assumption for convexity of individual functions, which is required in their work. Second, for strongly convex and smooth functions, we will omit the  $(\sigma_{\max}^2 + \Sigma_{\max}^2)$  denominator in the logarithmic term of our  $\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log((\sigma_{1:T}^2 + \Sigma_{1:T}^2)/(\sigma_{\max}^2 + \Sigma_{\max}^2)))$  bound below for simplicity.

### 5.1 Fully Adversarial Data

For fully adversarial data, we have  $\sigma_{1:T}^2 = 0$  as  $\sigma_t^2 = 0$  for  $t \in [T]$ , and  $\Sigma_{1:T}^2$  is equivalent to  $V_T$ . In this case, our bound in Theorem 3 guarantees an  $\mathcal{O}(\frac{1}{\lambda} \log V_T)$  regret bound for  $\lambda$ -strongly convex and smooth functions, recovering the gradient-variation bound of [Zhang et al. \(2022\)](#). By contrast, the result of [Sachs et al. \(2022\)](#) can only recover the  $\mathcal{O}(\frac{1}{\lambda} \log T)$  worst-case bound. Furthermore, for  $\alpha$ -exp-concave functions, our new result (Theorem 5) implies an  $\mathcal{O}(\frac{d}{\alpha} \log V_T)$  regret bound for OCO, recovering the result of [Chiang et al. \(2012\)](#).

### 5.2 Fully Stochastic Data

For fully stochastic data, the loss functions are i.i.d., so we have  $\Sigma_{1:T}^2 = 0$  and  $\sigma_t = \sigma$ ,  $\forall t \in [T]$ . Then for  $\lambda$ -strongly convex functions, Theorem 3 implies the same  $\mathcal{O}(\log T/[\lambda T])$  excess risk bound as [Sachs et al. \(2022\)](#). Besides, Theorem 5 further delivers a new  $\mathcal{O}(d \log T/[\alpha T])$  bound for  $\alpha$ -exp-concave functions. These results match the well-known bounds in SCO ([Hazan et al., 2007](#)) through online-to-batch conversion.



### 5.3 Adversarially Corrupted Stochastic Data

In the adversarially corrupted stochastic model, the loss function consists of two parts:  $f_t(\cdot) = h_t(\cdot) + c_t(\cdot)$ , where  $h_t(\cdot)$  is the loss of i.i.d. data sampled from a fixed distribution  $\mathfrak{D}$ , and  $c_t(\cdot)$  is a smooth adversarial perturbation satisfying that  $\sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{X}} \|\nabla c_t(\mathbf{x})\| \leq C_T$ , where  $C_T \geq 0$  is a parameter called the *corruption level*. Ito (2021) studies this model in expert and bandit problems, proposing a bound consisting of regret of i.i.d. data and an  $\sqrt{C_T}$  term measuring the corrupted performance. Sachs et al. (2022) achieve a similar  $\mathcal{O}(\sigma\sqrt{T} + \sqrt{C_T})$  bound in OCO problems under convexity and smoothness conditions, and raise an open question about how to extend the results to strongly convex losses. We resolve the problem by applying Theorem 3 of optimistic OMD to this model.

**Corollary 1.** *In the adversarially corrupted stochastic model, Our Theorem 3 implies an  $\mathcal{O}(\frac{1}{\lambda} \log(\sigma^2 T + C_T))$  bound for  $\lambda$ -strongly convex expected functions; and Theorem 5 implies an  $\mathcal{O}(\frac{d}{\alpha} \log(\sigma^2 T + C_T))$  bound for  $\alpha$ -exp-concave individual functions.*

The proof of Corollary 1 is in Appendix C.1. We successfully extend results of Ito (2021) not only to strongly convex functions, but also to exp-concave functions.

### 5.4 Random Order Model

Random order model (ROM) (Garber et al., 2020; Sherman et al., 2021) relaxes the adversarial setting in standard adversarial OCO, where the nature is allowed to choose the set of loss functions even with complete knowledge of the algorithm. However, nature cannot choose the order of loss functions, which will be arranged in uniformly random order.

Same as Sachs et al. (2022), let  $\bar{\nabla}_T(\mathbf{x}) \triangleq \frac{1}{T} \sum_{s=1}^T \nabla f_s(\mathbf{x})$ . Then we have  $\sigma_1^2 = \max_{\mathbf{x} \in \mathcal{X}} \frac{1}{T} \sum_{t=1}^T \|\nabla f_t(\mathbf{x}) - \bar{\nabla}_T(\mathbf{x})\|_2^2$  and we define  $\Lambda = \frac{1}{T} \sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \bar{\nabla}_T(\mathbf{x})\|_2^2$ . Note that  $\Lambda$  is a relaxation of  $\sigma_1^2$  and the logarithm of  $\Lambda/\sigma_1^2$  will not be large in reasonable scenarios. Sachs et al. (2022) establish an  $\mathcal{O}(\sigma_1 \sqrt{\log(\Lambda/\sigma_1)T})$  bound but require the convexity of individual functions, and they ask whether  $\sigma$ -dependent regret bounds can be realized under weaker assumptions on convexity of expected functions like Sherman et al. (2021). In Corollary 2, we give an affirmative answer based on Theorem 1 and obtain the results with weak assumptions. The proof is in Appendix C.2.

**Corollary 2.** *For convex expected functions, ROM enjoys an  $\mathcal{O}(\sigma_1 \sqrt{\log(\Lambda/\sigma_1)T})$  bound.*

For  $\lambda$ -strongly convex expected functions, Theorem 3 leads to an  $\mathcal{O}(\frac{1}{\lambda} \log(T\sigma_1^2 \log(\Lambda/\sigma_1^2)))$  bound, which is more stronger than the  $\mathcal{O}(\frac{1}{\lambda} \sigma_1^2 \log T)$  bound of Sachs et al. (2022) when  $\sigma_1^2$  is not too small. Meanwhile, the best-of-both-worlds guarantee in Theorem 3 safeguards that our final bound is never worse than theirs. Besides, for  $\alpha$ -exp-concave functions, we establish a new  $\mathcal{O}(\frac{d}{\alpha} \log(T\sigma_1^2 \log(\Lambda/\sigma_1^2)))$  bound from Theorem 5, but the curvature assumption is imposed over individual functions. Thus an open question is whether a similar  $\sigma$ -dependent bound can be obtained under the convexity of expected functions.

### 5.5 Slow Distribution Shift

We consider a simple problem instance of online learning with slow distribution shifts, in which the underlying distributions selected by the nature in every two adjacent rounds are

close on average. Formally, we suppose that  $(1/T) \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \leq \varepsilon$ , where  $\varepsilon$  is a constant. So we can get that  $\Sigma_{1:T}^2 \leq T\varepsilon$ . For  $\lambda$ -strongly convex functions, our Theorem 3 realizes an  $\mathcal{O}(\frac{1}{\lambda} \log(\sigma_{1:T}^2 + \varepsilon T))$  regret bound, which is tighter than the  $\mathcal{O}(\frac{1}{\lambda}(\sigma_{\max}^2 \log T + \varepsilon T))$  bound of Sachs et al. (2022) for a large range of  $\varepsilon$ . Extending the analysis to  $\alpha$ -exp-concave functions yields an  $\mathcal{O}(\frac{d}{\alpha} \log(\sigma_{1:T}^2 + \varepsilon T))$  regret from Theorem 5.

## 5.6 Online Learning with Limited Resources

In real-world online learning applications, functions often arrive not individually but rather in groups. Let  $K_t$  denote the number of functions coming in round  $t$  and  $f_t(\cdot, i)$  denote the  $i$ -th function. Denote by  $F_t(\cdot) \triangleq \frac{1}{K_t} \sum_{i=1}^{K_t} f_t(\cdot, i)$  the average of all functions.

We consider the scenarios with limited computing resources such that gradient estimation can only be achieved by sampling a portion of the functions, leading to gradient variance. Assume that at each time  $t$  we sample  $1 \leq B_t \leq K_t$  functions, where the  $i$ -th function is expressed as  $\hat{f}_t(\cdot, i)$ . We can then estimate  $F_t(\cdot)$  by  $f_t(\cdot) \triangleq \frac{1}{B_t} \sum_{i=1}^{B_t} \hat{f}_t(\cdot, i)$ , and further we have an upper bound for  $\sigma_t^2$  as follows.

$$\begin{aligned} \sigma_t^2 &= \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E} \left[ \left\| \frac{1}{B_t} \sum_{i=1}^{B_t} \nabla \hat{f}_t(\mathbf{x}, i) - \nabla F_t(\mathbf{x}) \right\|_2^2 \right] \\ &= \frac{1}{B_t^2} \max_{\mathbf{x} \in \mathcal{X}} \left( \sum_{i=1}^{B_t} \mathbb{E} \left[ \left\| \nabla \hat{f}_t(\mathbf{x}, i) - \nabla F_t(\mathbf{x}) \right\|_2^2 \right] \right. \\ &\quad \left. + \mathbb{E} \left[ \sum_{i \neq j} \left\langle \mathbb{E} \left[ \nabla \hat{f}_t(\mathbf{x}, i) - \nabla F_t(\mathbf{x}) \right], \mathbb{E} \left[ \nabla \hat{f}_t(\mathbf{x}, j) - \nabla F_t(\mathbf{x}) \right] \right\rangle \right] \right) \\ &= \frac{1}{B_t^2} \max_{\mathbf{x} \in \mathcal{X}} \left( \sum_{i=1}^{B_t} \mathbb{E} \left[ \left\| \nabla \hat{f}_t(\mathbf{x}, i) - \nabla F_t(\mathbf{x}) \right\|_2^2 \right] \right) \leq \frac{4G^2}{B_t}, \end{aligned}$$

where we use the fact that  $\nabla \hat{f}_t(\mathbf{x}, i)$  and  $\nabla \hat{f}_t(\mathbf{x}, j)$  are independent when  $i \neq j$ , and the fact that  $\mathbb{E}[\nabla \hat{f}_t(\mathbf{x}, i) - \nabla F_t(\mathbf{x})] = 0$ . The last inequality is due to Assumption 1. As a result, we have  $\sigma_{1:T}^2 = \mathbb{E}[\sum_{t=1}^T \sigma_t^2] \leq 4G^2 \sum_{t=1}^T \frac{1}{B_t}$  and obtain the following corollary by substituting it into Theorem 1, Theorem 3, and Theorem 5, respectively.

**Corollary 3.** *In online learning with limited resources, we can obtain an  $\mathcal{O}(2G\sqrt{\sum_{t=1}^T \frac{1}{B_t}} + \sqrt{\Sigma_{1:T}^2})$  bound for convex functions by Theorem 1; and Theorem 3 implies an  $\mathcal{O}(\frac{1}{\lambda} \log(4G^2 \sum_{t=1}^T \frac{1}{B_t} + \Sigma_{1:T}^2))$  bound for  $\lambda$ -strongly convex functions; and Theorem 5 leads to an  $\mathcal{O}(\frac{d}{\alpha} \log(4G^2 \sum_{t=1}^T \frac{1}{B_t} + \Sigma_{1:T}^2))$  bound for  $\alpha$ -exp-concave functions.*

When the number of sampled functions increases, the estimated gradient will gradually approach the real gradient and the variance will be close to 0. Note that the ratio  $B_t/K_t$  can be viewed as the *data throughput* determined by the available computing resources (Zhou, 2024; Wang et al., 2024). Corollary 3 demonstrates the impact of data throughput on the performance of learning algorithms.

### 5.7 Online Label Shift

In this part, we demonstrate the close connection of the SEA model to the problem of online label shift (OLS), which is independently studied in the recent literature (Bai et al., 2022; Qian et al., 2023). We show that our dynamic regret of the SEA model can be translated to OLS and obtain favorable guarantees.

**Problem setup.** OLS considers the multi-class classification in non-stationary environments, where the label distribution changes over time while the class-conditional is fixed. Denote by  $\mathcal{Z} \subseteq \mathbb{R}^{d'}$  the feature space and  $\mathcal{Y} = [K] \triangleq \{1, \dots, K\}$  the label space. OLS consists of two stages: (i) during the *offline initialization* stage, the learner trains a well-performed initial model  $h_0(\cdot) = h(\mathbf{x}_0, \cdot) : \mathcal{Z} \rightarrow \mathcal{Y}$  based on a labeled sample set  $S_0 = \{(\mathbf{z}_n, y_n)\}_{n=1}^{N_0}$  drawn from the distribution  $\mathfrak{D}_0(\mathbf{z}, y)$ ; (ii) during the *online adaptation* stage, at each round  $t \in [T]$ , the learner needs to make predictions of a small number of unlabeled data  $S_t = \{\mathbf{z}_n\}_{n=1}^{N_t}$  drawn from the distribution  $\mathfrak{D}_t(\mathbf{z})$ . The distributions  $\mathfrak{D}_t(\mathbf{z})$  are continuously shifting over time and thereby the learner should update the model  $\mathbf{x}_t \in \mathcal{X}$  adaptively. As mentioned, a *label shift* assumption is imposed:  $\mathfrak{D}_t(y)$  can vary over time while  $\mathfrak{D}_t(\mathbf{z} | y)$  is identical throughout the process.

**Performance measure and Key quantities.** In OLS, the model's quality is evaluated by its risk  $F_t(\mathbf{x}) = \mathbb{E}_{(\mathbf{z}, y) \sim \mathfrak{D}_t}[\ell(h(\mathbf{x}, \mathbf{z}), y)]$  at round  $t$ , where  $\ell(\cdot, \cdot)$  can be any convex surrogate loss for classification and  $h : \mathcal{Z} \times \mathcal{W} \rightarrow \mathbb{R}^K$  is the predictive function parametrized by  $\mathbf{x}$ . We use *dynamic regret* to measure the performance of online algorithms. However, we cannot directly use  $F_t$  for updating since it is unknown due to the lack of supervision. To address this problem, Bai et al. (2022) conduct the risk rewriting as follows:

$$F_t(\mathbf{x}) \triangleq \sum_{k=1}^K [\boldsymbol{\mu}_{y_t}]_k \cdot F_0^k(\mathbf{x}), \quad \text{with } F_t^k(\mathbf{x}) \triangleq \mathbb{E}_{\mathbf{z} \sim \mathfrak{D}_t(\mathbf{z} | y=k)}[\ell(h(\mathbf{x}, \mathbf{z}), k)] \quad (50)$$

where  $\boldsymbol{\mu}_{y_t} \in \Delta_K$  denotes the label distribution vector with the  $k$ -th entry  $[\boldsymbol{\mu}_{y_t}]_k \triangleq \mathfrak{D}_t(y = k)$  and  $F_t^k(\mathbf{x}) \triangleq \mathbb{E}_{\mathbf{z} \sim \mathfrak{D}_t(\mathbf{z} | y=k)}[\ell(h(\mathbf{x}, \mathbf{z}), k)]$  is the risk evaluated over the  $k$ -th label at round  $t$ . Note that we use  $F_0^k(\mathbf{x}) = F_t^k(\mathbf{x})$  here, owing to the label shift assumption that  $\mathfrak{D}_t(\mathbf{z} | y)$  is fixed over the process. Further, an estimator of  $F_t(\mathbf{x})$  can be established as

$$f_t(\mathbf{x}) \triangleq \sum_{k=1}^K [\hat{\boldsymbol{\mu}}_{y_t}]_k \cdot f_0^k(\mathbf{x}), \quad \text{with } f_0^k(\mathbf{x}) = \frac{1}{|S_0^k|} \sum_{\mathbf{z}_n \in S_0^k} \ell(h(\mathbf{x}, \mathbf{z}_n), k), \quad (51)$$

where  $S_0^k$  denotes a subset of  $S_0$  containing all samples with label  $k$  and  $\hat{\boldsymbol{\mu}}_{y_t}$  is an estimator of  $\boldsymbol{\mu}_{y_t}$  that can be constructed by the Black Box Shift Estimation (BBSE) method (Lipton et al., 2018). Specifically, they first obtain the predictive labels  $\hat{y}_t$  by using the initial model  $h_0$  to predict over the unlabeled data  $S_t$ , and then compute the label distribution  $\boldsymbol{\mu}_{y_t}$  via solving the crucial equation  $\boldsymbol{\mu}_{y_t} = C_{h_0}^{-1} \boldsymbol{\mu}_{\hat{y}_t}$ , where  $\boldsymbol{\mu}_{\hat{y}_t} \in \Delta_K$  is the distribution vector of the predictive labels  $\hat{y}_t$  and  $C_{h_0} \in \mathbb{R}^{K \times K}$  is the confusion matrix with  $[C_{h_0}]_{ij} = \mathbb{E}_{\mathbf{z} \sim \mathfrak{D}_0(\mathbf{z} | y=j)}[\mathbb{1}\{h_0(\mathbf{z}) = i\}]$ . Then  $C_{h_0}$  can be estimated empirically by  $[\hat{C}_{h_0}]_{ij} = \sum_{(\mathbf{z}, y) \in S_0} \mathbb{1}\{h_0(\mathbf{z}) = i \text{ and } y = j\} / \mathbb{1}\{y = j\}$ , using the offline labeled data  $S_0$ . And  $\boldsymbol{\mu}_{\hat{y}_t}$  can be estimated empirically with online data  $S_t$ , which is given by

$[\hat{\boldsymbol{\mu}}_{\hat{y}_t}]_j = \frac{1}{|S_t|} \sum_{\mathbf{z} \in S_t} \mathbb{1}\{h_0(\mathbf{z}) = j\}$ . With the above estimation, the final estimator for the label distribution vector is constructed as  $\hat{\boldsymbol{\mu}}_{y_t} = \hat{C}_{h_0}^{-1} \hat{\boldsymbol{\mu}}_{\hat{y}_t}$ . Following Bai et al. (2022), we assume that  $S_0$  has a sufficient sample size such that  $\hat{C}_{h_0} = C_{h_0}$  and  $f_0^k(\mathbf{x}) = F_0^k(\mathbf{x})$ . While this assumption may appear strong at first glance, it is primarily introduced to simplify the presentation without trivializing the online adaptation and can be removed by incorporating a concentration analysis as discussed in Bai et al. (2022). As a result,  $f_t(\mathbf{x})$  is an unbiased estimator with respect to  $F_t(\mathbf{x})$ .

**Regret guarantee.** Under such a setup, the SEA model can be applied to the OLS problem by considering the expected function (50) and the randomized function (51). By Theorem 9, we can obtain the following guarantee, whose proof is in Appendix C.3.

**Corollary 4.** *Modeling the online label shift problem as the SEA model with the expected function defined as (50) and the randomized function (51), and further applying Algorithm 3, we can obtain that for  $\mathbf{x}_t^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} F_t(\mathbf{x})$ ,*

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \right] \\ & \leq \mathcal{O} \left( L_T^{\frac{1}{3}} T^{\frac{1}{3}} \left( \sqrt{\sum_{t=1}^T \mathbb{E} [\|\hat{\boldsymbol{\mu}}_{y_t} - \boldsymbol{\mu}_{y_t}\|_2^2]} + \sqrt{\sum_{t=1}^T \mathbb{E} [\|\boldsymbol{\mu}_{y_t} - \boldsymbol{\mu}_{y_{t-1}}\|_2^2]} \right)^{\frac{2}{3}} \right), \end{aligned}$$

where  $L_T = \sum_{t=2}^T \|\boldsymbol{\mu}_{y_t} - \boldsymbol{\mu}_{y_{t-1}}\|_1$  measures the label distributions changes.

**Remark 12.** For the OLS problem, Bai et al. (2022) provide a general optimistic regret bound in the form of  $\mathcal{O}(L_T^{\frac{1}{3}} G_T^{\frac{1}{3}} T^{\frac{1}{3}})$ , where  $G_T \triangleq \sum_{t=1}^T \mathbb{E} [\sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla H_t(\mathbf{x})\|_2^2]$  with the hint function  $H_t(\mathbf{x}) = \sum_{k=1}^K [\mathbf{h}_{y_t}]_k \cdot f_0^k(\mathbf{x})$ . They further provide a bias-variance decomposition for  $G_T$ , as restated below:

$$G_T \leq KG^2 \sum_{t=1}^T \mathbb{E} [\|\mathbf{h}_{y_t} - \hat{\boldsymbol{\mu}}_{y_t}\|_2^2] \leq 2KG^2 \sum_{t=1}^T \left( \underbrace{\mathbb{E} [\|\mathbf{h}_{y_t} - \boldsymbol{\mu}_{y_t}\|_2^2]}_{\text{bias}} + \underbrace{\mathbb{E} [\|\boldsymbol{\mu}_{y_t} - \hat{\boldsymbol{\mu}}_{y_t}\|_2^2]}_{\text{variance}} \right). \quad (52)$$

Therefore, when setting the hint as last-round empirical class prior, their  $\mathcal{O}(L_T^{\frac{1}{3}} G_T^{\frac{1}{3}} T^{\frac{1}{3}})$  bound implies the same result as ours in Corollary 4, as we have the following upper bound for  $G_T$  when setting  $\mathbf{h}_{y_t} = \hat{\boldsymbol{\mu}}_{y_{t-1}}$ ,

$$G_T \leq KG^2 + 4KG^2 \sum_{t=2}^T \left( \mathbb{E} [\|\boldsymbol{\mu}_{y_t} - \boldsymbol{\mu}_{y_{t-1}}\|_2^2] \right) + 6KG^2 \sum_{t=1}^T \left( \mathbb{E} [\|\hat{\boldsymbol{\mu}}_{y_t} - \boldsymbol{\mu}_{y_t}\|_2^2] \right).$$

This demonstrates the close connection between the SEA model and the OLS problem. In fact, the bias-variance decomposition presented in (52) of Bai et al. (2022) is similar in spirit to the key quantities (cumulative stochastic variance  $\sigma_{1:T}^2$  and cumulative adversarial variation  $\Sigma_{1:T}^2$ ) introduced in the SEA model (Sachs et al., 2022).  $\triangleleft$

**Remark 13.** Besides, in the OLS problem, using the bound with  $\sigma_{1:T}^2$  (Theorem 7) or  $\tilde{\sigma}_{1:T}^2$  (Theorem 9) can actually give the *same* upper bound that scales with meaningful quantities. Specifically, we can respectively bound  $\sigma_{1:T}^2$  and  $\tilde{\sigma}_{1:T}^2$  by

$$\begin{aligned}\sigma_{1:T}^2 &= \mathbb{E} \left[ \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \mathbb{E} \left[ \left\| \sum_{k=1}^K ([\hat{\boldsymbol{\mu}}_{y_t}]_k - [\boldsymbol{\mu}_{y_t}]_k) \cdot \nabla F_0^k(\mathbf{x}) \right\|_2^2 \right] \right] \leq KG^2 \sum_{t=1}^T \mathbb{E} \left[ \|\hat{\boldsymbol{\mu}}_{y_t} - \boldsymbol{\mu}_{y_t}\|_2^2 \right], \\ \tilde{\sigma}_{1:T}^2 &= \mathbb{E} \left[ \sum_{t=1}^T \mathbb{E} \left[ \sup_{\mathbf{x} \in \mathcal{X}} \left\| \sum_{k=1}^K ([\hat{\boldsymbol{\mu}}_{y_t}]_k - [\boldsymbol{\mu}_{y_t}]_k) \cdot \nabla F_0^k(\mathbf{x}) \right\|_2^2 \right] \right] \leq KG^2 \sum_{t=1}^T \mathbb{E} \left[ \|\hat{\boldsymbol{\mu}}_{y_t} - \boldsymbol{\mu}_{y_t}\|_2^2 \right].\end{aligned}$$

Both quantities share the same upper bound in the form of label distribution variances.  $\triangleleft$

## 6 Conclusion and Future Work

In this paper, we investigate the stochastically extended adversarial (SEA) model of [Sachs et al. \(2022\)](#) and propose a different solution via the optimistic OMD framework. Our results yield the *same* regret bound for convex and smooth functions under weaker assumptions and a *better* regret bound for strongly convex and smooth functions; moreover, we establish the *first* regret bound for exp-concave and smooth functions. For all three cases, we enhance analyses of optimistic FTRL, proving equal regret bounds with optimistic OMD for the SEA model. Furthermore, we study the SEA model under *dynamic regret* and propose a new two-layer algorithm based on optimistic OMD, which obtains the *first* dynamic regret guarantee for the SEA model. Additionally, we explore the SEA model under *non-smooth* scenarios, in which we propose to use optimistic OMD with an implicit update to achieve static and dynamic regret guarantees. Lastly, we discuss implications for intermediate learning scenarios, leading to various new results.

Although our algorithms for various functions can be unified using the optimistic OMD framework, they still necessitate distinct configurations for parameters such as step sizes and regularizers. Consequently, it becomes crucial to conceive and develop more adaptive online algorithms that eliminate the need for pre-set parameters. Exploring this area of research and designing such algorithms will be an important focus in future studies.

## Acknowledgments and Disclosure of Funding

This work was partially supported by the National Science and Technology Major Project (2022ZD0114801) and NSFC (U23A20382, 62122037). Peng Zhao was supported in part by the Xiaomi Foundation.

## Appendix A. Omitted Proofs for Section 3

This section contains the omitted proofs of optimistic FTRL for Section 3, including Theorem 2, 4, 6 in Appendix A.1–Appendix A.3, followed by useful lemmas in Appendix A.4.

### A.1 Proof of Theorem 2

**Proof** For convex and smooth functions, we start by outlining the optimistic FTRL procedure. At each step  $t$ , a surrogate loss is defined:  $\ell_t(\mathbf{x}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} - \mathbf{x}_t \rangle$ . Unlike Sachs et al. (2022), we use this surrogate loss instead of the original function  $f_t(\cdot)$  to update  $\mathbf{x}_t$ , avoiding the need for convexity in individual functions (which is required by Sachs et al. (2022)). The decision  $\mathbf{x}_t$  is updated by deploying optimistic FTRL over the linearized loss:

$$\mathbf{x}_t = \arg \min_{\mathbf{x} \in \mathcal{X}} \sum_{s=1}^{t-1} \left\{ \ell_s(\mathbf{x}) + \langle M_t, \mathbf{x} \rangle + \frac{1}{\eta_t} \|\mathbf{x}\|_2^2 \right\},$$

where  $\mathbf{x}_0$  can be an arbitrary point in  $\mathcal{X}$ , and the optimistic vector  $M_t = \nabla f_{t-1}(\mathbf{x}_{t-1})$  (we set  $M_1 = \nabla f_0(\mathbf{x}_0) = 0$ ). The step size  $\eta_t$  is designed as  $\eta_t = D^2 / (\delta + \sum_{s=1}^{t-1} \eta_s \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1})\|_2^2)$  with  $\delta$  to be defined latter, which is non-increasing for  $t \in [T]$ .

We can easily obtain that

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u} \rangle \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \right].$$

As a result, we only need to consider the regret of the surrogate loss  $\ell_t(\cdot)$ . The following proof is similar to Sachs et al. (2022). To exploit Lemma 15 (regret analysis of optimistic FTRL), we map the  $G_t$  term in Lemma 15 to  $\frac{1}{\eta_t} \|\mathbf{x}\|_2^2 + \sum_{s=1}^{t-1} \ell_s(\mathbf{x})$  and map the  $\tilde{\mathbf{g}}_t$  term to  $M_t$ . Note that  $G_t$  is  $\frac{2}{\eta_t}$ -strongly convex and  $\ell_t$  is convex, we have

$$\begin{aligned} \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) &\leq \frac{D^2}{\eta_T} + \sum_{t=1}^T \left( \langle \nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}_t - \mathbf{x}_{t+1} \rangle - \frac{1}{\eta_t} \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 \right) \\ &\leq \frac{D^2}{\eta_T} + \sum_{t=1}^T \left( \frac{\eta_t}{2} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 - \frac{1}{2\eta_t} \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 \right) \\ &\leq \delta + \frac{3}{2} \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 - \frac{\delta}{2D^2} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 \\ &\leq \frac{3\sqrt{2}}{2} D \sqrt{\bar{V}_T} + \frac{6D^2 G^2}{\delta} + \delta - \frac{\delta}{2D^2} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2. \end{aligned}$$

where we use the fact that  $\langle a, b \rangle \leq \|a\|_* \|b\| \leq \frac{1}{2c} \|a\|_*^2 + \frac{c}{2} \|b\|^2$  in the second inequality ( $\|\cdot\|_*$  denotes the dual norm of  $\|\cdot\|$ ), based on the Hölder's inequality. The third step is due to the fact  $\eta_t \leq \frac{D^2}{\delta}$  ( $t \in [T]$ ) and the last step use the inequality  $\sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \leq D \sqrt{2\bar{V}_T} + \frac{4D^2 G^2}{\delta}$  from Sachs et al. (2022, proof of Theorem 5).

Using Lemma 3 (Boundedness of cumulative norm of the gradient difference), we have

$$\begin{aligned}
& \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \\
& \leq 6D \sqrt{\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2} + 3\sqrt{2}D \sqrt{\sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2} \\
& \quad + 3\sqrt{2}DL \sqrt{\sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2} - \frac{\delta}{2D^2} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 + \frac{6D^2G^2}{\delta} + \delta + \frac{3\sqrt{2}}{2}DG \\
& \leq 6D \sqrt{\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2} + 3\sqrt{2}D \sqrt{\sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2} \\
& \quad + \frac{9D^4L^2}{\delta} + \frac{6D^2G^2}{\delta} + \delta + \frac{3\sqrt{2}}{2}DG, \tag{53}
\end{aligned}$$

where we use the following inequality in the last step  $3\sqrt{2}DL \sqrt{\sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2} \leq \frac{9D^4L^2}{\delta} + \frac{\delta}{2D^2} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2$ , canceling out the the negative term in (53) with the second term.

Then, we take expectations over (53) with the help of definitions of  $\sigma_{1:T}^2$  and  $\Sigma_{1:T}^2$ , and use Jensen's inequality. Given that the expected regret of surrogate loss functions upper bounds the expected regret of original functions, we get the final result:

$$\begin{aligned}
\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] & \leq \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \right] \\
& \leq 6D \sqrt{\sigma_{1:T}^2} + 3\sqrt{2}D \sqrt{\Sigma_{1:T}^2} + \frac{9D^4L + 6D^2G^2}{\delta} + \delta + \frac{3\sqrt{2}}{2}DG \\
& = 6D \sqrt{\sigma_{1:T}^2} + 3\sqrt{2}D \sqrt{\Sigma_{1:T}^2} + 2\sqrt{9D^4L + 6D^2G^2} + \frac{3\sqrt{2}}{2}DG \\
& = \mathcal{O} \left( \sqrt{\sigma_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right),
\end{aligned}$$

where we set  $\delta = \sqrt{9D^4L + 6D^2G^2}$ . Hence, we complete the proof.  $\blacksquare$

## A.2 Proof of Theorem 4

**Proof** We first present the procedure of optimistic FTRL for  $\lambda$ -strongly convex and smooth functions (Sachs et al., 2022). In each round  $t$ , we define a new surrogate loss:  $\ell_t(\mathbf{x}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} - \mathbf{x}_t \rangle + \frac{\lambda}{2} \|\mathbf{x} - \mathbf{x}_t\|_2^2$ . And the decision  $\mathbf{x}_{t+1}$  is determined by

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{\lambda}{2} \|\mathbf{x} - \mathbf{x}_0\|_2^2 + \sum_{s=1}^t \ell_s(\mathbf{x}) + \langle M_{t+1}, \mathbf{x} \rangle \right\},$$

where  $\mathbf{x}_0$  is an arbitrary point in  $\mathcal{X}$ , and the optimistic vector  $M_{t+1} = \nabla f_t(\mathbf{x}_t)$ . In the beginning, we set  $M_1 = \nabla f_0(\mathbf{x}_0) = 0$  and thus  $\mathbf{x}_1 = \mathbf{x}_0$ . Compared with the original

algorithm of [Sachs et al. \(2022\)](#), we insert an additional  $\frac{\lambda}{2}\|\mathbf{x} - \mathbf{x}_0\|_2^2$  term in the updating rule above, and in this way, the objective function in the  $t$ -th round is  $\lambda t$ -strongly convex, which facilitates the subsequent analysis.

According to (18), it is easy to verify that

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \right]. \quad (54)$$

Thus, we can focus on the regret of surrogate loss  $\ell_t(\cdot)$ . From Lemma 15 (regret analysis of optimistic FTRL), since  $\frac{\lambda}{2}\|\mathbf{x} - \mathbf{x}_0\|_2^2 + \sum_{s=1}^{t-1} \ell_s(\mathbf{x})$  is  $\lambda t$ -strongly convex, we obtain

$$\begin{aligned} & \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \\ & \leq \frac{\lambda}{2} \|\mathbf{u} - \mathbf{x}_0\|_2^2 + \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}_t - \mathbf{x}_{t+1} \rangle - \sum_{t=1}^T \frac{\lambda t}{2} \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 \\ & \leq \frac{\lambda D^2}{2} + \sum_{t=1}^T \frac{1}{\lambda t} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 - \frac{\lambda}{4} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2. \end{aligned} \quad (55)$$

Then we directly use Lemma 5 (Boundedness of the norm of gradient difference) to obtain

$$\begin{aligned} & \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \\ & \leq \frac{G^2}{\lambda} + \sum_{t=2}^T \frac{1}{\lambda t} (4\|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + 4\|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ & \quad + 4\|\nabla F_{t-1}(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2) + \sum_{t=1}^T \left( \frac{4L^2}{\lambda(t+1)} - \frac{\lambda t}{4} \right) \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + \frac{\lambda D^2}{2} \\ & \leq \frac{G^2}{\lambda} + \sum_{t=1}^T \frac{8}{\lambda t} \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + \sum_{t=2}^T \frac{4}{\lambda t} \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ & \quad + \sum_{t=1}^T \left( \frac{4L^2}{\lambda t} - \frac{\lambda t}{4} \right) \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + \frac{\lambda D^2}{2}. \end{aligned} \quad (56)$$

The above formula reuses the simplification techniques in (20). Still defining  $\kappa = \frac{L}{\lambda}$ , then for  $t \geq 16\kappa$ , there is  $\frac{4L^2}{\lambda t} - \frac{\lambda t}{4} \leq 0$ . For this reason, it turns out that

$$\begin{aligned} & \sum_{t=1}^T \left( \frac{4L^2}{\lambda t} - \frac{\lambda t}{4} \right) \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \leq \sum_{t=1}^{\lceil 16\kappa \rceil} \left( \frac{4L^2}{\lambda t} - \frac{\lambda t}{4} \right) D^2 \leq \frac{4L^2 D^2}{\lambda} \sum_{t=1}^{\lceil 16\kappa \rceil} \frac{1}{t} \\ & \leq \frac{4L^2 D^2}{\lambda} \left( 1 + \int_{t=1}^{\lceil 16\kappa \rceil} \frac{1}{t} \right) = \frac{4L^2 D^2}{\lambda} \ln \left( 1 + 16\frac{L}{\lambda} \right) + \frac{4L^2 D^2}{\lambda}. \end{aligned}$$



By substituting the above inequality into (56) and taking the expectation, we can obtain

$$\begin{aligned}
& \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \right] \\
& \leq \mathbb{E} \left[ \sum_{t=1}^T \frac{8}{\lambda t} \sigma_t^2 + \sum_{t=2}^T \frac{4}{\lambda t} \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right] + \frac{4L^2 D^2}{\lambda} \ln \left( 1 + 16 \frac{L}{\lambda} \right) \\
& \quad + \frac{4L^2 D^2 + G^2}{\lambda} + \frac{\lambda D^2}{2}.
\end{aligned}$$

Similar to the derivation using optimistic OMD, we take advantage of Lemma 6 to get

$$\begin{aligned}
& \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \right] \\
& \leq \frac{8\sigma_{\max}^2 + 4\Sigma_{\max}^2}{\lambda} \ln \left( \frac{1}{2\sigma_{\max}^2 + \Sigma_{\max}^2} (2\sigma_{1:T}^2 + \Sigma_{1:T}^2) + 1 \right) + \frac{8\sigma_{\max}^2 + 4\Sigma_{\max}^2 + 4}{\lambda} \\
& \quad + \frac{4L^2 D^2}{\lambda} \ln \left( 1 + 16 \frac{L}{\lambda} \right) + \frac{4L^2 D^2 + G^2}{\lambda} + \frac{\lambda D^2}{2} \\
& = \mathcal{O} \left( \frac{1}{\lambda} (\sigma_{\max}^2 + \Sigma_{\max}^2) \log \left( (\sigma_{1:T}^2 + \Sigma_{1:T}^2) / (\sigma_{\max}^2 + \Sigma_{\max}^2) \right) \right),
\end{aligned}$$

which completes the proof.  $\blacksquare$

### A.3 Proof of Theorem 6

**Proof** We use the following optimistic FTRL for  $\alpha$ -exp-concave and smooth functions,

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{1}{2} (1 + \beta G^2) \|\mathbf{x}\|_2^2 + \sum_{s=1}^t \ell_s(\mathbf{x}) + \langle M_{t+1}, \mathbf{x} \rangle \right\},$$

where  $\mathbf{x}_0$  is an arbitrary point in  $\mathcal{X}$ ,  $M_{t+1} = \nabla f_t(\mathbf{x}_t)$ , and the surrogate loss  $\ell_t(\mathbf{x}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} - \mathbf{x}_t \rangle + \frac{\beta}{2} \|\mathbf{x} - \mathbf{x}_t\|_{h_t}^2$  with  $\beta = \frac{1}{2} \min \left\{ \frac{1}{4GD}, \alpha \right\}$ , and  $h_t = \nabla f_t(\mathbf{x}_t) \nabla f_t(\mathbf{x}_t)^\top$ . Furthermore, we set  $M_1 = \nabla f_0(\mathbf{x}_0) = 0$ . From (21), we can easily derive that

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \right]. \quad (57)$$

So in the following, we focus on the regret of surrogate losses. Denoting by  $H_t = I + \beta G^2 I + \beta \sum_{s=1}^{t-1} h_s$  (where  $I$  is the  $d \times d$  identity matrix) and  $G_t(\mathbf{x}) = \frac{1}{2} (1 + \beta G^2) \|\mathbf{x}\|_2^2 + \sum_{s=1}^{t-1} \ell_s(\mathbf{x})$ , we have that  $G_t(\mathbf{x})$  is 1-strongly convex w.r.t.  $\|\cdot\|_{H_t}$ . Hence, using Lemma 15 (regret analysis of optimistic FTRL), we immediately get the following guarantee

$$\sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u})$$

$$\begin{aligned}
 &\leq \frac{1 + \beta G^2}{2} \|\mathbf{u}\|_2^2 + \sum_{t=1}^T \left( \langle \nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}_t - \mathbf{x}_{t+1} \rangle - \frac{1}{2} \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_{H_t}^2 \right) \\
 &\leq \frac{(1 + \beta G^2) D^2}{2} + \underbrace{\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_{H_t^{-1}}^2}_{\text{term (a)}} - \underbrace{\frac{1}{4} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_{H_t}^2}_{\text{term (b)}}, \quad (58)
 \end{aligned}$$

where we denote the dual norm of  $\|\cdot\|_{H_t}$  by  $\|\cdot\|_{H_t^{-1}}$ , and use Assumption 2 (domain boundedness) and  $\langle a, b \rangle \leq \|a\|_* \|b\| \leq \frac{1}{2c} \|a\|_*^2 + \frac{c}{2} \|b\|^2$  in the second inequality.

To bound term (a) in (58), we begin with the fact that

$$\begin{aligned}
 H_t &\succeq I + \beta \sum_{s=1}^t \nabla f_s(\mathbf{x}_s) \nabla f_s(\mathbf{x}_s)^\top \\
 &\succeq I + \frac{\beta}{2} \sum_{s=1}^t \left( \nabla f_s(\mathbf{x}_s) \nabla f_s(\mathbf{x}_s)^\top + \nabla f_{s-1}(\mathbf{x}_{s-1}) \nabla f_{s-1}(\mathbf{x}_{s-1})^\top \right), \quad (59)
 \end{aligned}$$

where the first inequality is due to Assumption 1 (boundedness of gradient norms) and the second inequality comes from the definition that  $\nabla f_0(\mathbf{x}_0) = 0$ . We substitute (24) in the proof of Theorem 5 into (59) and obtain that

$$H_t \succeq I + \frac{\beta}{4} \sum_{s=1}^t (\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1})) (\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1}))^\top.$$

Let  $P_t = I + \frac{\beta}{4} \sum_{s=1}^t (\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1})) (\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\mathbf{x}_{s-1}))^\top$  so that  $H_t \succeq P_t$ , then we can bound term (a) in (58) as

$$\text{term (a)} \leq \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_{P_t^{-1}}^2 = \frac{4}{\beta} \sum_{t=1}^T \left\| \sqrt{\frac{\beta}{4}} (\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})) \right\|_{P_t^{-1}}^2.$$

By applying Lemma 14 with  $\mathbf{u}_t = \sqrt{\frac{\beta}{4}} (\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1}))$  and  $\varepsilon = 1$ , we get that

$$\text{term (a)} \leq \frac{4d}{\beta} \ln \left( \frac{\beta}{4d} \bar{V}_T + 1 \right).$$

Then we move to term (b). Since  $H_t = I + \beta G^2 I + \beta \sum_{s=1}^{t-1} h_s \succeq I$ , we can derive that

$$\text{term (b)} = \frac{1}{4} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_{H_t}^2 \geq \frac{1}{4} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_I^2 = \frac{1}{4} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2.$$

So we bound the guarantee in (58) by substituting the bounds of term (a) and term (b):

$$\sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \leq \frac{(1 + \beta G^2) D^2}{2} + \frac{4d}{\beta} \ln \left( \frac{\beta}{4d} \bar{V}_T + 1 \right) - \frac{1}{4} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2.$$

Through Lemma 3 (Boundedness of cumulative norm of gradient difference) together with the inequality of  $\ln(1+u+v) \leq \ln(1+u) + \ln(1+v)$  ( $u, v > 0$ ), we get that

$$\begin{aligned}
& \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \\
& \leq \frac{4d}{\beta} \ln \left( \frac{2\beta}{d} \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + \frac{\beta}{d} \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 + \frac{\beta}{4d} G^2 + 1 \right) \\
& \quad + \frac{(1+\beta G^2)D^2}{2} + \frac{4d}{\beta} \ln \left( \frac{\beta L^2}{d} \sum_{t=2}^T \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + 1 \right) - \frac{1}{4} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 \\
& \leq \frac{4d}{\beta} \ln \left( \frac{2\beta}{d} \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + \frac{\beta}{d} \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 + \frac{\beta}{4d} G^2 + 1 \right) \\
& \quad + \frac{(1+\beta G^2)D^2}{2} + \frac{4d}{\beta} \ln(16L^2 + 1).
\end{aligned}$$

where the last step comes from Lemma 9.

Then we compute the expected regret by taking the expectation over the above regret with the help of Jensen's inequality and the derived result in (57):

$$\begin{aligned}
& \mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \right] \leq \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{u}) \right] \\
& \leq \frac{4d}{\beta} \ln \left( \frac{2\beta}{d} \sigma_{1:T}^2 + \frac{\beta}{d} \Sigma_{1:T}^2 + \frac{\beta}{4d} G^2 + 1 \right) + \frac{(1+\beta G^2)D^2}{2} + \frac{4d}{\beta} \ln(16L^2 + 1) \\
& = \mathcal{O} \left( \frac{d}{\alpha} \log(\sigma_{1:T}^2 + \Sigma_{1:T}^2) \right).
\end{aligned}$$

■

#### A.4 Useful Lemmas

We first provide the proof of Lemma 2, 3, 5 and 6, and then present other lemmas useful for the proofs.

**Proof** [of Lemma 2] We can decompose the instantaneous loss as

$$\begin{aligned}
& \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x} \rangle \\
& = \underbrace{\langle \nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle}_{\text{term (a)}} + \underbrace{\langle \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle}_{\text{term (b)}} + \underbrace{\langle \nabla f_t(\mathbf{x}_t), \widehat{\mathbf{x}}_{t+1} - \mathbf{x} \rangle}_{\text{term (c)}}.
\end{aligned}$$

For term (a), we use Lemma 4 (stability lemma) and get that

$$\text{term (a)} = \langle \nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle$$

$$\leq \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_* \|\mathbf{x}_t - \hat{\mathbf{x}}_{t+1}\| \leq \frac{1}{\alpha} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_*^2.$$

For term (b) and term (c), due to the updating rules of optimistic OMD in (7) and (8), we can apply Lemma 10 (Bregman proximal inequality) and obtain that

$$\begin{aligned} \text{term (b)} &= \langle \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}_t - \hat{\mathbf{x}}_{t+1} \rangle \leq \mathcal{D}_{\psi_t}(\hat{\mathbf{x}}_{t+1}, \hat{\mathbf{x}}_t) - \mathcal{D}_{\psi_t}(\hat{\mathbf{x}}_{t+1}, \mathbf{x}_t) - \mathcal{D}_{\psi_t}(\mathbf{x}_t, \hat{\mathbf{x}}_t), \\ \text{term (c)} &= \langle \nabla f_t(\mathbf{x}_t), \hat{\mathbf{x}}_{t+1} - \mathbf{x} \rangle \leq \mathcal{D}_{\psi_t}(\mathbf{x}, \hat{\mathbf{x}}_t) - \mathcal{D}_{\psi_t}(\mathbf{x}, \hat{\mathbf{x}}_{t+1}) - \mathcal{D}_{\psi_t}(\hat{\mathbf{x}}_{t+1}, \hat{\mathbf{x}}_t). \end{aligned}$$

We complete the proof by combining the three upper bounds.  $\blacksquare$

**Proof** [of Lemma 3] It is easy to verify the above lemma by substituting (60) and (61) in Lemma 5 into  $\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2$  and simplifying the result.  $\blacksquare$

**Proof** [of Lemma 5] For  $t \geq 2$ , from Jensen's inequality and Assumption 4 (smoothness of expected function), we have

$$\begin{aligned} & \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ &= 16 \left\| \frac{1}{4} [\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)] + \frac{1}{4} [\nabla F_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_{t-1})] \right. \\ & \quad \left. + \frac{1}{4} [\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})] + \frac{1}{4} [\nabla F_{t-1}(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})] \right\|_2^2 \\ &\leq 4 \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + 4 \|\nabla F_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_{t-1})\|_2^2 \\ & \quad + 4 \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 + 4 \|\nabla F_{t-1}(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \\ &\leq 4 \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + 4L^2 \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \\ & \quad + 4 \|\nabla F_t(\mathbf{x}_{t-1}) - \nabla F_{t-1}(\mathbf{x}_{t-1})\|_2^2 + 4 \|\nabla F_{t-1}(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2. \end{aligned} \tag{60}$$

For  $t = 1$ , from Assumption 1 (boundedness of the gradient norm), we have

$$\|\nabla f_1(\mathbf{x}_1) - \nabla f_0(\mathbf{x}_0)\|_2^2 = \|\nabla f_1(\mathbf{x}_1)\|_2^2 \leq G^2. \tag{61}$$

Combining both cases finishes the proof.  $\blacksquare$

**Proof** [of Lemma 6] We first define the following quantity:

$$\alpha = \left\lceil \sum_{t=1}^T \frac{1}{2\sigma_{\max}^2 + \Sigma_{\max}^2} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \right\rceil.$$

If  $1 \leq \alpha < T$ , we bound  $\sum_{t=1}^T \frac{1}{\lambda t} (2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2)$  as follows.

$$\begin{aligned} & \sum_{t=1}^T \frac{1}{\lambda t} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \\ &= \sum_{t=1}^{\alpha} \frac{1}{\lambda t} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) + \sum_{t=\alpha+1}^T \frac{1}{\lambda t} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \end{aligned}$$

$$\begin{aligned}
&\leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \sum_{t=1}^{\alpha} \frac{1}{t} + \frac{1}{\lambda(\alpha+1)} \sum_{t=\alpha+1}^T \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \\
&\leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \left( 1 + \int_{t=1}^{\alpha} \frac{1}{t} dt \right) + \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} (\ln \alpha + 1) + \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \\
&\leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \ln \left( \sum_{t=1}^T \frac{1}{2\sigma_{\max}^2 + \Sigma_{\max}^2} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) + 1 \right) + \frac{4\sigma_{\max}^2 + 2\Sigma_{\max}^2}{\lambda}.
\end{aligned}$$

Else if  $\alpha = T$ , we have

$$\begin{aligned}
&\sum_{t=1}^T \frac{1}{\lambda t} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \\
&= \sum_{t=1}^{\alpha} \frac{1}{\lambda t} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \sum_{t=1}^{\alpha} \frac{1}{t} \\
&\leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \left( 1 + \int_{t=1}^{\alpha} \frac{1}{t} dt \right) \leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} (\ln \alpha + 1) \\
&\leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda} \ln \left( \sum_{t=1}^T \frac{1}{2\sigma_{\max}^2 + \Sigma_{\max}^2} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) + 1 \right) + \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda}.
\end{aligned}$$

Else if  $\alpha < 1$ , we have

$$\begin{aligned}
&\sum_{t=1}^T \frac{1}{\lambda t} \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \\
&\leq \frac{1}{\lambda \alpha} \sum_{t=1}^T \left( 2\sigma_t^2 + \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right) \leq \frac{2\sigma_{\max}^2 + \Sigma_{\max}^2}{\lambda}.
\end{aligned}$$

■

**Lemma 9.** *Let  $A_T$  be a non-negative term,  $a, b$  be non-negative constants and  $c$  be a positive constant, then we have*

$$a \ln(bA_T + 1) - cA_T \leq a \ln \left( \frac{ab}{c} + 1 \right). \quad (62)$$

**Proof** We use the following inequality:  $\ln p \leq \frac{p}{q} + \ln q - 1$  holds for all  $p > 0, q > 0$ . By setting  $p = bA_T + 1$  and  $q = \frac{ab}{c} + 1$ , we obtain

$$\begin{aligned}
a \ln(bA_T + 1) - cA_T &\leq a \left( \frac{bA_T + 1}{ab/c + 1} + \ln \left( \frac{ab}{c} + 1 \right) - 1 \right) - cA_T \\
&= c \left( \frac{ab}{ab + c} - 1 \right) A_T + \left( \frac{1}{ab/c + 1} - 1 \right) a + a \ln \left( \frac{ab}{c} + 1 \right) \leq a \ln \left( \frac{ab}{c} + 1 \right).
\end{aligned}$$

■

## Appendix B. Omitted Proofs for Section 4

In this section, we present the omitted details for Section 4, including a discussion of the method using alternative optimism design and a useful lemma.

### B.1 Elaborations on an alternative method

In this part, we demonstrate that when employing an alternative optimism design with  $M_{t+1} = \nabla f_t(\bar{\mathbf{x}}_{t+1})$  where  $\bar{\mathbf{x}}_{t+1} = \sum_{i=1}^N p_{t,i} \mathbf{x}_{t+1,i}$ , we can only obtain a slightly worse regret scaling with the quantity  $\tilde{\sigma}_{1:T}^2$ .

We first briefly describe the algorithm, which is a variant of Algorithm 2. With the same two-layer structure as Algorithm 2, each base-learner  $\mathcal{B}_i$  updates its local decision by

$$\hat{\mathbf{x}}_{t+1,i} = \Pi_{\mathcal{X}}[\hat{\mathbf{x}}_{t,i} - \eta_i \nabla f_t(\mathbf{x}_{t,i})], \quad \mathbf{x}_{t+1,i} = \Pi_{\mathcal{X}}[\hat{\mathbf{x}}_{t+1,i} - \eta_i \nabla f_t(\mathbf{x}_{t,i})],$$

which requires its own gradient direction; the meta-learner omits correction terms and updates the weight vector  $\mathbf{p}_{t+1} \in \Delta_N$  by  $p_{t+1,i} \propto \exp(-\varepsilon_t (\sum_{s=1}^t \ell_{s,i} + m_{t+1,i}))$  where the feedback loss  $\ell_t \in \mathbb{R}^N$  is constructed by  $\ell_{t,i} = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_{t,i} \rangle$  and the optimism  $\mathbf{m}_{t+1} \in \mathbb{R}^N$  is constructed as  $m_{t+1,i} = \langle M_{t+1}, \mathbf{x}_{t+1,i} \rangle$  with  $M_{t+1} = \nabla f_t(\bar{\mathbf{x}}_{t+1})$ . The step size  $\eta_i$  of base-learners and the learning rate  $\varepsilon_t$  of meta-learner will be given later. Then for the above alternative algorithm, we can obtain the following theoretical guarantee.

**Theorem 11.** *Under Assumptions 1, 2, 4 and 5, setting the step size pool  $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$  with  $\eta_i = \min\{1/(4L), 2^{i-1} \sqrt{D^2/(98G^2T)}\}$  and  $N = \lceil 2^{-1} \log_2(8G^2T/(L^2D^2)) \rceil + 1$ , and setting the learning rate as  $\varepsilon_t = 1/\sqrt{\delta + 4G^2 + \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\bar{\mathbf{x}}_s)\|_2^2}$  with  $\delta = 4D^2L^2(\ln N + 2D^2)$  for all  $t \in [T]$ , this variant of Algorithm 2 using the above optimism design and with no correction terms (more specifically, setting  $\lambda = 0$  and  $M_{t+1} = \nabla f_t(\bar{\mathbf{x}}_{t+1})$  with  $\bar{\mathbf{x}}_{t+1} = \sum_{i=1}^N p_{t,i} \mathbf{x}_{t+1,i}$ ) can obtain the following bound*

$$\mathbb{E}[\mathbf{Reg}_T^d] \leq \mathcal{O}\left(P_T + \sqrt{1 + P_T} \left(\sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}\right)\right).$$

**Proof** [of Theorem 11] Notice that the dynamic regret can be decomposed into two parts:

$$\mathbb{E}\left[\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t^*)\right] = \underbrace{\mathbb{E}\left[\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_{t,i})\right]}_{\text{meta-regret}} + \underbrace{\mathbb{E}\left[\sum_{t=1}^T f_t(\mathbf{x}_{t,i}) - \sum_{t=1}^T f_t(\mathbf{u}_t^*)\right]}_{\text{base-regret}}.$$

Then we provide the upper bounds for the two terms respectively.

**Bounding the meta-regret.** Before giving the analysis, we first define  $\hat{V}_t = \sum_{s=1}^t \|\nabla f_s(\mathbf{x}_s) - \nabla f_{s-1}(\bar{\mathbf{x}}_s)\|_2^2$  for the brevity of subsequent analysis. Similar to the proof in the previous section, we can easily get

$$\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t,i} \rangle \leq \sum_{t=1}^T \varepsilon_t \|\ell_t - \mathbf{m}_t\|_\infty^2 + \frac{\ln N}{\varepsilon_{T+1}} - \sum_{t=2}^T \frac{1}{4\varepsilon_t} \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2$$

$$\begin{aligned}
&\leq D^2 \sum_{t=1}^T \varepsilon_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\bar{\mathbf{x}}_t)\|_2^2 + \frac{\ln N}{\varepsilon_{T+1}} - \sum_{t=2}^T \frac{1}{4\varepsilon_t} \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2 \\
&\leq (\ln N + 2D^2) \sqrt{\delta + 4G^2 + \hat{V}_T} - \frac{\sqrt{\delta}}{4} \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2, \tag{63}
\end{aligned}$$

where we bound the adaptivity term in the second inequality by

$$\|\ell_t - \mathbf{m}_t\|_\infty^2 = \max_{i \in [N]} \langle \nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\bar{\mathbf{x}}_t), \mathbf{x}_{t,i} \rangle^2 \leq D^2 \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\bar{\mathbf{x}}_t)\|_2^2,$$

and the last inequality comes from Lemma 12 and the fact that  $\varepsilon_t \leq \frac{1}{\sqrt{\delta + \hat{V}_t}} \leq \frac{1}{\sqrt{\delta}}$ . Based on Assumption 4 (smoothness of expected function),  $\hat{V}_T$  can be bounded by

$$\begin{aligned}
\hat{V}_T &= \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\bar{\mathbf{x}}_t)\|_2^2 \\
&\leq G^2 + 4 \sum_{t=2}^T \left( \|\nabla f_t(\mathbf{x}_t) - \nabla F_t(\mathbf{x}_t)\|_2^2 + \|\nabla F_t(\mathbf{x}_t) - \nabla F_{t-1}(\mathbf{x}_t)\|_2^2 + \|\nabla F_{t-1}(\mathbf{x}_t) - \nabla F_{t-1}(\bar{\mathbf{x}}_t)\|_2^2 \right. \\
&\quad \left. + \|\nabla F_{t-1}(\bar{\mathbf{x}}_t) - \nabla f_{t-1}(\bar{\mathbf{x}}_t)\|_2^2 \right) \tag{64} \\
&\leq G^2 + 8 \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2 + 4 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_t) - \nabla F_{t-1}(\mathbf{x}_t)\|_2^2 + 4L^2 \sum_{t=2}^T \|\mathbf{x}_t - \bar{\mathbf{x}}_t\|_2^2 \\
&\leq G^2 + 8 \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2 + 4 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_t) - \nabla F_{t-1}(\mathbf{x}_t)\|_2^2 + 4D^2L^2 \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2,
\end{aligned}$$

where the last inequality is due to the fact

$$\|\mathbf{x}_t - \bar{\mathbf{x}}_t\|_2^2 = \left\| \sum_{i=1}^N (p_{t,i} - p_{t-1,i}) \mathbf{x}_{t,i} \right\|_2^2 \leq \left( \sum_{i=1}^N |p_{t,i} - p_{t-1,i}| \|\mathbf{x}_{t,i}\|_2 \right)^2 \leq D^2 \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2.$$

As a result, substitute the above upper bound into (63), we arrive at

$$\begin{aligned}
&\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t,i} \rangle \\
&\leq (\ln N + 2D^2) \sqrt{\delta + 5G^2} + \left( 2\sqrt{2} \ln N + 4\sqrt{2}D^2 \right) \sqrt{\sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2} \\
&\quad + (2 \ln N + 4D^2) \sqrt{\sum_{t=2}^T \|\nabla F_t(\mathbf{x}_t) - \nabla F_{t-1}(\mathbf{x}_t)\|_2^2} + (2DL \ln N + 4D^3L) \sqrt{\sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2} \\
&\quad - \frac{\sqrt{\delta}}{4} \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2
\end{aligned}$$

$$\begin{aligned} &\leq (\ln N + 2D^2) \sqrt{\delta + 5G^2} + \left(2\sqrt{2} \ln N + 4\sqrt{2}D^2\right) \sqrt{\sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2} \\ &\quad + (2 \ln N + 4D^2) \sqrt{\sum_{t=2}^T \|\nabla F_t(\mathbf{x}_t) - \nabla F_{t-1}(\mathbf{x}_t)\|_2^2} + \frac{(2DL \ln N + 4D^3L)^2}{\sqrt{\delta}}, \end{aligned}$$

where the last inequality is due to

$$(2DL \ln N + 4D^3L) \sqrt{\sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2} \leq \frac{(2DL \ln N + 4D^3L)^2}{\sqrt{\delta}} + \frac{\sqrt{\delta}}{4} \sum_{t=2}^T \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1^2.$$

Finally, we take expectations with the help of Jensen's inequality and obtain

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t,i} \rangle \right] &\leq \left(2\sqrt{2} \ln N + 4\sqrt{2}D^2\right) \sqrt{\tilde{\sigma}_{1:T}^2} + (2 \ln N + 4D^2) \sqrt{\Sigma_{1:T}^2} \\ &\quad + (\ln N + 2D^2) \left( \sqrt{5}G + 4DL \sqrt{\ln N + 2D^2} \right), \end{aligned}$$

where we set  $\delta = 4D^2L^2 (\ln N + 2D^2)$ .

**Bounding the base-regret.** Notice that the base-learner actually performs optimistic OMD with the regularizer  $\psi_t(\mathbf{x}) = \frac{1}{2\eta_i} \|\mathbf{x}\|_2^2$ , we can apply Lemma 2 to obtain the base-regret for any index  $i \in [N]$  as:

$$\begin{aligned} &\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_{t,i}), \mathbf{x}_{t,i} - \mathbf{u}_t^* \rangle \\ &\leq \eta_i \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla f_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 + \frac{1}{2\eta_i} \sum_{t=1}^T \left( \|\mathbf{u}_t^* - \hat{\mathbf{x}}_{t,i}\|_2^2 - \|\mathbf{u}_t^* - \hat{\mathbf{x}}_{t+1,i}\|_2^2 \right) \\ &\quad - \frac{1}{2\eta_i} \sum_{t=1}^T \left( \|\hat{\mathbf{x}}_{t+1,i} - \mathbf{x}_{t,i}\|_2^2 - \|\hat{\mathbf{x}}_{t,i} - \mathbf{x}_{t,i}\|_2^2 \right) \\ &\leq \eta_i \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla f_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 + \frac{D^2 + 2DP_T}{2\eta_i} - \frac{1}{4\eta_i} \sum_{t=2}^T \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2, \end{aligned}$$

where the derivation of the last inequality is similar to the previous proof and will not be repeated here. By exploiting Lemma 3, the above formula can be further bounded by

$$\begin{aligned} &\sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_{t,i}), \mathbf{x}_{t,i} - \mathbf{u}_t^* \rangle \\ &\leq \eta_i G^2 + 8\eta_i \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla F_t(\mathbf{x}_{t,i})\|_2^2 + 4\eta_i \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1,i}) - \nabla F_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 \\ &\quad + \left( 4\eta_i L^2 - \frac{1}{4\eta_i} \right) \sum_{t=2}^T \|\mathbf{x}_{t,i} - \mathbf{x}_{t-1,i}\|_2^2 + \frac{D^2 + 2DP_T}{2\eta_i} \end{aligned}$$



$$\begin{aligned}
&\leq 8\eta_i \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla F_t(\mathbf{x}_{t,i})\|_2^2 + 4\eta_i \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1,i}) - \nabla F_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 \\
&\quad + \frac{D^2 + 2DP_T}{2\eta_i} + \eta_i G^2.
\end{aligned}$$

The last inequality holds by ensuring the step size satisfies  $\eta_i \leq 1/(4L)$  for any  $i \in [N]$ . Moreover, the above formula shows that the best step size is  $\eta^\dagger = \min\{1/(4L), \eta^*\}$ , where

$$\eta^* = \sqrt{\frac{D^2 + 2DP_T}{16 \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla F_t(\mathbf{x}_{t,i})\|_2^2 + 8 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1,i}) - \nabla F_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 + 2G^2}}.$$

We set the step size pool  $\mathcal{H} = \left\{ \eta_i = \min \left\{ \frac{1}{4L}, 2^{i-1} \sqrt{\frac{D^2}{98G^2T}} \right\} \mid i \in [N] \right\}$ , which ensures that  $\eta^\dagger$  is included. Then, if  $\eta^* \leq 1/(4L)$ , there must be an  $\eta_{i^*} \in \mathcal{H}$  satisfying that  $\eta_{i^*} \leq \eta^* \leq 2\eta_{i^*}$ , and we can obtain that

$$\begin{aligned}
&\eta_{i^*} \left( 8 \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla F_t(\mathbf{x}_{t,i})\|_2^2 + 4 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1,i}) - \nabla F_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 + G^2 \right) + \frac{D^2 + 2DP_T}{2\eta_{i^*}} \\
&\leq \eta^* \left( 8 \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla F_t(\mathbf{x}_{t,i})\|_2^2 + 4 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1,i}) - \nabla F_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 + G^2 \right) + \frac{D^2 + 2DP_T}{\eta^*} \\
&= 2\sqrt{(D^2 + 2DP_T) \left( 8 \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla F_t(\mathbf{x}_{t,i})\|_2^2 + 4 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1,i}) - \nabla F_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 + G^2 \right)}.
\end{aligned}$$

Otherwise, if  $\eta^* > 1/(4L)$ , we will choose  $\eta_i = 1/(4L)$  and obtain that

$$\begin{aligned}
&\frac{1}{4L} \left( 8 \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_{t,i}) - \nabla F_t(\mathbf{x}_{t,i})\|_2^2 + 4 \sum_{t=2}^T \|\nabla F_t(\mathbf{x}_{t-1,i}) - \nabla F_{t-1}(\mathbf{x}_{t-1,i})\|_2^2 + G^2 \right) \\
&\quad + 2L(D^2 + 2DP_T) \leq 6L(D^2 + 2DP_T).
\end{aligned}$$

Hence, we get the final regret bound of the base-learner by taking both cases into account and taking expectations with Jensen's inequality

$$\mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_{t,i}), \mathbf{x}_{t,i} - \mathbf{u}_t^* \rangle \right] \leq 2\sqrt{D^2 + 2DP_T} \left( 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} + G \right) + 6L(D^2 + 2DP_T).$$

**Bounding the overall dynamic regret.** Combining the meta-regret and the base-regret, and using the convexity of expected functions, we have

$$\begin{aligned}
&\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t^*) \right] \\
&\leq \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t,i} \rangle \right] + \mathbb{E} \left[ \sum_{t=1}^T \langle \nabla f_t(\mathbf{x}_{t,i}), \mathbf{x}_{t,i} - \mathbf{u}_t^* \rangle \right]
\end{aligned}$$

$$\begin{aligned}
&\leq \left(2 \ln N + 4D^2 + 4\sqrt{D^2 + 2DP_T}\right) \left(\sqrt{2\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}\right) \\
&\quad + (\ln N + 2D^2) \left(\sqrt{5}G + 4DL\sqrt{\ln N + 2D^2}\right) + 2G\sqrt{D^2 + 2DP_T} + 6L(D^2 + 2DP_T) \\
&= \mathcal{O}\left(P_T + \sqrt{1 + P_T} \left(\sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2}\right)\right),
\end{aligned}$$

which completes the proof.  $\blacksquare$

**Remark 14.** In fact, we can also prove Theorem 11 simply by taking expectations over the  $\mathcal{O}(\sqrt{(1 + P_T + V_T)(1 + P_T)})$  bound in Theorem 3 of Zhao et al. (2020), but we give the above specific proof to illustrate the *dependence issue* of this alternative optimism design. Specifically, according to the definition of  $\bar{\mathbf{x}}_t$ , it has dependency on  $f_{t-1}$ . So we cannot directly obtain the expectation of the gray item in (64), but can only perform the supremum operation first, which results in the inability to get a bound scaling with  $\sigma_{1:T}^2$ .  $\triangleleft$

## B.2 Useful Lemma

**Proof** [of Lemma 7] According to the proof of the second inequality in Proposition 4.1 of Campolongo and Orabona (2020), using the first-order optimality condition for  $\mathbf{x}_{t+1}$  yields

$$\langle \nabla f_t(\mathbf{x}_{t+1}) + \nabla \psi_{t+1}(\mathbf{x}_{t+1}) - \nabla \psi_{t+1}(\hat{\mathbf{x}}_{t+1}), \mathbf{x} - \mathbf{x}_{t+1} \rangle \geq 0, \quad \forall \mathbf{x} \in \mathcal{X}.$$

By moving terms, the above equation can be rewritten as

$$\begin{aligned}
\langle \nabla f_t(\mathbf{x}_{t+1}), \mathbf{x}_{t+1} - \mathbf{x} \rangle &\leq \langle \nabla \psi_{t+1}(\hat{\mathbf{x}}_{t+1}) - \nabla \psi_{t+1}(\mathbf{x}_{t+1}), \mathbf{x}_{t+1} - \mathbf{x} \rangle \\
&= B_{\psi_{t+1}}(\mathbf{x}, \hat{\mathbf{x}}_{t+1}) - B_{\psi_{t+1}}(\mathbf{x}, \mathbf{x}_{t+1}) - B_{\psi_{t+1}}(\mathbf{x}_{t+1}, \hat{\mathbf{x}}_{t+1}), \quad \forall \mathbf{x} \in \mathcal{X}.
\end{aligned}$$

Since  $\psi_{t+1}(\mathbf{x}) = \frac{1}{2\eta_{t+1}} \|\mathbf{x}\|_2^2$ , we can transform the above equation into

$$\langle \nabla f_t(\mathbf{x}_{t+1}), \mathbf{x}_{t+1} - \mathbf{x} \rangle \leq \frac{1}{2\eta_{t+1}} (\|\mathbf{x} - \hat{\mathbf{x}}_{t+1}\|_2^2 - \|\mathbf{x} - \mathbf{x}_{t+1}\|_2^2 - \|\mathbf{x}_{t+1} - \hat{\mathbf{x}}_{t+1}\|_2^2), \quad \forall \mathbf{x} \in \mathcal{X},$$

which completes the proof.  $\blacksquare$

**Proof** [of Lemma 8] For  $t \geq 2$ , using Jensen's inequality, we have

$$\begin{aligned}
&\sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2 \\
&\leq \sup_{\mathbf{x} \in \mathcal{X}} (2\|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2 + 2\|\nabla F_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2) \\
&\leq \sup_{\mathbf{x} \in \mathcal{X}} (2\|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2 + 4\|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 + 4\|\nabla F_{t-1}(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2).
\end{aligned}$$

For  $t = 1$ , from Assumption 1 (boundedness of the gradient norm), we have

$$\sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_1(\mathbf{x}) - \nabla f_0(\mathbf{x})\|_2^2 = \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_1(\mathbf{x})\|_2^2 \leq G^2.$$

As a result, we can complete the proof by adding the terms from  $t = 1$  to  $t = T$ .  $\blacksquare$

## Appendix C. Omitted Proofs for Section 5

This section presents omitted proofs of corollaries in Section 5, including proof of Corollary 1 in Appendix C.1, proof of Corollary 2 in Appendix C.2 and proof of Corollary 4 in Appendix C.3.

### C.1 Proof of Corollary 1

**Proof** Recall that in Section 5.3 the loss functions in adversarially corrupted stochastic model satisfy  $f_t(\mathbf{x}) = h_t(\mathbf{x}) + c_t(\mathbf{x})$  for all  $t \in [T]$ , where  $h_t(\cdot)$  is sampled from a fixed distribution every iteration and  $\sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{X}} \|\nabla c_t(\mathbf{x})\| \leq C_T$ . By definition of  $F_t(\mathbf{x})$ ,

$$F_t(\mathbf{x}) = \mathbb{E}_{f_t \sim \mathfrak{D}_t}[f_t(\mathbf{x})] = \mathbb{E}_{h_t \sim \mathfrak{D}}[h_t(\mathbf{x}) + c_t(\mathbf{x})] = \mathbb{E}_{h_t \sim \mathfrak{D}}[h_t(\mathbf{x})] + c_t(\mathbf{x}). \quad (65)$$

Since  $h_t(\cdot)$  is i.i.d for each  $t$ , their expectations are the same. Then we have

$$\begin{aligned} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 &\leq 2G\|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2 \stackrel{(65)}{=} 2G\|\nabla c_t(\mathbf{x}) - \nabla c_{t-1}(\mathbf{x})\| \\ &\leq 2G(\|\nabla c_t(\mathbf{x})\| + \|\nabla c_{t-1}(\mathbf{x})\|). \end{aligned}$$

Therefore, we have the following upper bound for the cumulative variation:

$$\Sigma_{1:T}^2 = \mathbb{E} \left[ \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right] \leq \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} 2G(\|\nabla c_t(\mathbf{x})\| + \|\nabla c_{t-1}(\mathbf{x})\|) \leq 4GC_T.$$

Besides, we can calculate the variance as

$$\sigma_t^2 = \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{f_t \sim \mathfrak{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2] \stackrel{(65)}{=} \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{h_t \sim \mathfrak{D}_t} [\|\nabla h_t(\mathbf{x}) - \nabla \mathbb{E}_{h_t \sim \mathfrak{D}}[h_t(\mathbf{x})]\|_2^2] = \sigma,$$

where  $\sigma > 0$  is the variance of stochastic gradients. This implies  $\sigma_{1:T}^2 = \mathbb{E} \left[ \sum_{t=1}^T \sigma_t^2 \right] = \sigma T$ .

Combining the above two upper bounds of  $\sigma_{1:T}^2$  and  $\Sigma_{1:T}^2$  with the regret bounds of optimistic OMD in Theorem 3 and Theorem 5 completes the proof.  $\blacksquare$

### C.2 Proof of Corollary 2

**Proof** The difference between ROM and i.i.d. stochastic model is that ROM samples a loss from the loss set without replacement in each round, while i.i.d. stochastic model samples independently and uniformly with replacement in each round. However, following Sachs et al. (2022), we can bound the variance of ROM with respect to  $\mathfrak{D}_t$  for each  $t$  by the variance  $\sigma_1^2$  of the first round, which can also be regarded as the variance of the i.i.d. model for every round. Specifically, for  $\forall \mathbf{x} \in \mathcal{X}$  and every  $t \in [T]$ , we have

$$\mathbb{E}_{f_t \sim \mathfrak{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2] \leq \mathbb{E}_{f_t \sim \mathfrak{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_1(\mathbf{x})\|_2^2]. \quad (66)$$

Since ROM samples losses without replacement, let set  $\Gamma_t$  represent the index set of losses that can be selected in the  $t$ th round, thus  $\Gamma_1 = [T]$ , then we have

$$\mathbb{E}_{f_t \sim \mathfrak{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_1(\mathbf{x})\|_2^2] = \frac{1}{T - (t-1)} \sum_{i \in \Gamma_t} \|\nabla f_i(\mathbf{x}) - \nabla F_1(\mathbf{x})\|_2^2$$

$$\leq \frac{1}{T - (t - 1)} \sum_{i \in \Gamma_1} \|\nabla f_i(\mathbf{x}) - \nabla F_1(\mathbf{x})\|_2^2 \leq \frac{T}{T - (t - 1)} \sigma_1^2.$$

So combining (66) with the above inequality, we get that

$$\mathbb{E}_{f_t \sim \mathfrak{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_t(\mathbf{x})\|_2^2] \leq \frac{T}{T - (t - 1)} \sigma_1^2, \quad \forall \mathbf{x} \in \mathcal{X}, t \in [T]. \quad (67)$$

Besides, from (66), we can also get that

$$\begin{aligned} \mathbb{E} [\sigma_t^2] &\leq \mathbb{E} \left[ \max_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{f_t \sim \mathfrak{D}_t} [\|\nabla f_t(\mathbf{x}) - \nabla F_1(\mathbf{x})\|_2^2] \right] \\ &\leq \mathbb{E} \left[ \mathbb{E}_{f_t \sim \mathfrak{D}_t} \left[ \max_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla F_1(\mathbf{x})\|_2^2 \right] \right] = \Lambda, \end{aligned} \quad (68)$$

where we review that  $\Lambda = \frac{1}{T} \sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \bar{\nabla}_T(\mathbf{x})\|_2^2$ . Then, we use a technique from [Sachs et al. \(2022\)](#) by introduce a variable  $\tau \in [T]$ , which help us upper bound  $\sigma_{1:T}^2$  as

$$\begin{aligned} \sigma_{1:T}^2 &= \mathbb{E} \left[ \sum_{t=1}^T \sigma_t^2 \right] \leq \mathbb{E} \left[ \sum_{t=1}^{\tau} \sigma_t^2 \right] + \mathbb{E} \left[ \sum_{t=\tau+1}^T \sigma_t^2 \right] \stackrel{(67), (68)}{\leq} \sum_{t=1}^{\tau} \frac{T}{T - (t - 1)} \sigma_1^2 + (T - \tau) \Lambda \\ &\leq \sum_{n=T-(\tau-1)}^T \frac{1}{n} T \sigma_1^2 + (T - \tau) \Lambda \leq \left( 1 + \log \frac{T}{T - (\tau - 1)} \right) T \sigma_1^2 + (T - \tau) \Lambda. \end{aligned}$$

If  $T \sigma_1^2 / \Lambda > 2$ , we set  $\tau = T - \lfloor T \sigma_1^2 / \Lambda \rfloor$ , then we have

$$\begin{aligned} \sigma_{1:T}^2 &\leq \left( 1 + \log \frac{T}{\lfloor T \sigma_1^2 / \Lambda \rfloor} \right) T \sigma_1^2 + T \sigma_1^2 \leq \left( 1 + \log \frac{1}{\sigma_1^2 / \Lambda - 1/T} \right) T \sigma_1^2 + T \sigma_1^2 \\ &\leq \left( 1 + \log \frac{2\Lambda}{\sigma_1^2} \right) T \sigma_1^2 + T \sigma_1^2 \leq T \sigma_1^2 \log \left( \frac{2e^2 \Lambda}{\sigma_1^2} \right). \end{aligned}$$

Otherwise, if  $T \sigma_1^2 / \Lambda \leq 2$ , we set  $\tau = T$ , then we can get the regret bound of  $\mathcal{O}(T \sigma_1^2 (1 + \log T))$ . Since we have

$$\mathcal{O}(T \sigma_1^2 (1 + \log T)) \leq \mathcal{O}(T \sigma_1^2 (1 + \log(2\Lambda / \sigma_1^2))) \leq \mathcal{O}(T \sigma_1^2 \log(2e^2 \Lambda / \sigma_1^2)),$$

then the final bound of  $\sigma_{1:T}^2$  is of order  $\mathcal{O} \left( T \sigma_1^2 \log \left( \frac{2e^2 \Lambda}{\sigma_1^2} \right) \right)$ .

Next, we try to bound  $\Sigma_{1:T}^2$ . We suppose that  $k_t = \Gamma_t \setminus \Gamma_{t+1}$  represents the loss selected in round  $t$ , then we have

$$\begin{aligned} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 &= \left\| \frac{1}{T - (t - 1)} \sum_{i \in \Gamma_t} \nabla f_i(\mathbf{x}) - \frac{1}{T - (t - 2)} \sum_{i \in \Gamma_{t-1}} \nabla f_i(\mathbf{x}) \right\|_2^2 \\ &= \left\| \frac{(T - t + 2) - (T - t + 1)}{(T - t + 1)(T - t + 2)} \sum_{i \in \Gamma_t} \nabla f_i(\mathbf{x}) - \frac{1}{T - t + 2} \nabla f_{k_{t-1}}(\mathbf{x}) \right\|_2^2 \end{aligned}$$

$$\begin{aligned}
&\leq \frac{2}{(T-t+2)^2} \left\| \frac{1}{T-t+1} \sum_{i \in \Gamma_t} \nabla f_i(\mathbf{x}) \right\|_2^2 + \frac{2}{(T-t+2)^2} \|\nabla f_{k_{t-1}}(\mathbf{x})\|_2^2 \\
&\leq \frac{4G^2}{(T-t+2)^2},
\end{aligned}$$

where the last inequality is derived from Assumption 1 (boundedness of the gradient norm).

Summing the above inequality over  $t = 1, \dots, T$ , and taking the expectation give

$$\Sigma_{1:T}^2 = \mathbb{E} \left[ \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla F_t(\mathbf{x}) - \nabla F_{t-1}(\mathbf{x})\|_2^2 \right] \leq \sum_{t=1}^T \frac{4G^2}{(T-t+2)^2} \leq 8G^2.$$

Finally, we substitute the bound of  $\sigma_{1:T}^2$  and  $\Sigma_{1:T}$  into Theorem 1, which is for convex and smooth functions, and complete the proof.  $\blacksquare$

### C.3 Proof of Corollary 4

According to the problem setup in Section 5.7, we can bound  $\tilde{\sigma}_{1:T}^2$  as

$$\begin{aligned}
\tilde{\sigma}_{1:T}^2 &= \mathbb{E} \left[ \sum_{t=1}^T \mathbb{E} \left[ \sup_{\mathbf{x} \in \mathcal{X}} \left\| \sum_{k=1}^K ([\hat{\boldsymbol{\mu}}_{y_t}]_k - [\boldsymbol{\mu}_{y_t}]_k) \cdot \nabla F_0^k(\mathbf{x}) \right\|_2^2 \right] \right] \\
&\leq \mathbb{E} \left[ \sum_{t=1}^T \mathbb{E} \left[ G^2 \left\| \sum_{k=1}^K ([\hat{\boldsymbol{\mu}}_{y_t}]_k - [\boldsymbol{\mu}_{y_t}]_k) \right\|_2^2 \right] \right] \leq KG^2 \sum_{t=1}^T \mathbb{E} [\|\hat{\boldsymbol{\mu}}_{y_t} - \boldsymbol{\mu}_{y_t}\|_2^2] \quad (69)
\end{aligned}$$

and bound  $\Sigma_{1:T}^2$  as

$$\begin{aligned}
\Sigma_{1:T}^2 &= \mathbb{E} \left[ \sum_{t=1}^T \sup_{\mathbf{x} \in \mathcal{X}} \left\| \sum_{k=1}^K ([\boldsymbol{\mu}_{y_t}]_k - [\boldsymbol{\mu}_{y_{t-1}}]_k) \cdot \nabla F_0^k(\mathbf{x}) \right\|_2^2 \right] \\
&\leq \mathbb{E} \left[ G^2 \sum_{t=1}^T \left\| \sum_{k=1}^K ([\boldsymbol{\mu}_{y_t}]_k - [\boldsymbol{\mu}_{y_{t-1}}]_k) \right\|_2^2 \right] \leq KG^2 \sum_{t=1}^T \mathbb{E} [\|\boldsymbol{\mu}_{y_t} - \boldsymbol{\mu}_{y_{t-1}}\|_2^2]. \quad (70)
\end{aligned}$$

Next, it is necessary to apply the key technique of Zhang et al. (2020) highlighted in Bai et al. (2022) to deal with the  $P_T$  term in Theorem 9. We know that  $P_T$  is the path length of the comparator sequence  $\mathbf{u}_1, \dots, \mathbf{u}_T$ , where  $\mathbf{u}_t$  can be any point in  $\mathcal{X}$ . While frequently used in exploring dynamic regret, this quantity lacks explicit significance in OLS. To this end, Bai et al. (2022) propose a new quantity  $L_T = \sum_{t=2}^T \|\boldsymbol{\mu}_{y_t} - \boldsymbol{\mu}_{y_{t-1}}\|_1$  to measure the variation of label distributions and consider the dynamic regret against the sequence  $\{\mathbf{x}_t^*\}_{t \in [T]}$ , where  $\mathbf{x}_t^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} F_t(\mathbf{x})$  but  $\mathbf{x}_t^* \notin \arg \min_{\mathbf{x} \in \mathcal{X}} f_t(\mathbf{x})$ . In their footsteps, we first decompose the dynamic regret bound into two parts by introducing a reference sequence that only changes every  $\Delta$  iteration. Specifically, the  $m$ -th time interval is denoted as  $\mathcal{I}_m = [(m-1)\Delta + 1, m\Delta]$

and any comparator  $\mathbf{u}_t$  within  $\mathcal{I}_m$  is considered the optimum decision for the interval, i.e.,  $\mathbf{u}_t = \mathbf{x}_{\mathcal{I}_m}^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} \sum_{t \in \mathcal{I}_m} F_t(\mathbf{x})$  for any  $t \in \mathcal{I}_m$ . Then we have

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \right] = \underbrace{\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \right]}_{\text{term (a)}} + \underbrace{\mathbb{E} \left[ \sum_{m=1}^M \sum_{t \in \mathcal{I}_m} f_t(\mathbf{x}_{\mathcal{I}_m}^*) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \right]}_{\text{term (b)}},$$

where  $M = \lceil T/\Delta \rceil \leq T/\Delta + 1$ .

For term (a), we can directly use (49) in the proof of Theorem 9 to get

$$\text{term (a)} \leq \left( D(\ln N + 4) + 2\sqrt{2(D^2 + 2DP_T)} \right) \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) + 4G^4 + \ln N + 4.$$

Notice we can derive that  $P_T \leq D(M-1) \leq DT/\Delta$  because the comparator sequence  $\{\mathbf{u}_t\}_{t=1}^T$  only changes  $M-1$  times. Hence, we have

$$\text{term (a)} \leq \left( D(\ln N + 4) + 4D\sqrt{\frac{2T}{\Delta}} + 2\sqrt{2}D \right) \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) + 4G^4 + \ln N + 4.$$

For term (b), we follow the analysis of Bai et al. (2022) to show that

$$\begin{aligned} \text{term (b)} &= \sum_{m=1}^M \sum_{t \in \mathcal{I}_m} (F_t(\mathbf{x}_{\mathcal{I}_m}^*) - F_t(\mathbf{x}_t^*)) \leq \sum_{m=1}^M \sum_{t \in \mathcal{I}_m} (F_t(\mathbf{x}_{s_m}^*) - F_t(\mathbf{x}_t^*)) \\ &= \sum_{m=1}^M \sum_{t \in \mathcal{I}_m} (F_t(\mathbf{x}_{s_m}^*) - F_{s_m}(\mathbf{x}_{s_m}^*) + F_{s_m}(\mathbf{x}_{s_m}^*) - F_t(\mathbf{x}_t^*)) \\ &\leq \sum_{m=1}^M \sum_{t \in \mathcal{I}_m} (F_t(\mathbf{x}_{s_m}^*) - F_{s_m}(\mathbf{x}_{s_m}^*) + F_{s_m}(\mathbf{x}_t^*) - F_t(\mathbf{x}_t^*)) \\ &\leq 2\Delta \sum_{m=1}^M \sum_{t \in \mathcal{I}_m} \sup_{\mathbf{x} \in \mathcal{X}} |F_t(\mathbf{x}) - F_{t-1}(\mathbf{x})| = 2\Delta \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} |F_t(\mathbf{x}) - F_{t-1}(\mathbf{x})|, \end{aligned}$$

where  $s_m = (m-1)\Delta + 1$  is the first time step at  $\mathcal{I}_m$ . Since  $\text{term (b)} = 0$  when  $\Delta = 1$ , we can bound it as  $\text{term (b)} \leq \mathbb{1}\{\Delta > 1\} \cdot 2\Delta \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} |F_t(\mathbf{x}) - F_{t-1}(\mathbf{x})|$ . Furthermore, we transform the  $\sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} |F_t(\mathbf{x}) - F_{t-1}(\mathbf{x})|$  term into a term related to  $L_T$  as

$$\begin{aligned} \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} |F_t(\mathbf{x}) - F_{t-1}(\mathbf{x})| &= \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \left| \sum_{k=1}^K ([\boldsymbol{\mu}_{y_t}]_k - [\boldsymbol{\mu}_{y_{t-1}}]_k) F_0^k(\mathbf{x}) \right| \\ &\leq \sum_{t=2}^T B \sum_{k=1}^K |[\boldsymbol{\mu}_{y_t}]_k - [\boldsymbol{\mu}_{y_{t-1}}]_k| = B \sum_{t=2}^T \|\boldsymbol{\mu}_{y_t} - \boldsymbol{\mu}_{y_{t-1}}\|_1 = BL_T, \end{aligned}$$

where  $B \triangleq \sup_{(\mathbf{z}, y) \in \mathcal{Z} \times \mathcal{Y}, \mathbf{x} \in \mathcal{X}} |\ell(h(\mathbf{x}, \mathbf{z}), y)|$  is the upper bound of function values. Thus,  $\text{term (b)} \leq \mathbb{1}\{\Delta > 1\} \cdot 2B\Delta L_T$ . Combining it with the upper bound of term (a) yields

$$\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \right]$$

$$\begin{aligned} &\leq \mathbb{1}\{\Delta > 1\} \cdot 2B\Delta L_T + 4D\sqrt{\frac{2T}{\Delta}} \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) \\ &\quad + \left( \ln N + 4 + 2\sqrt{2} \right) D \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) + 4G^4 + \ln N + 4. \end{aligned}$$

Below, we set different values for  $\Delta$  in two cases and obtain the final regret bound.

**Case 1.** When  $D\sqrt{2T} \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) > BL_T$ , in such a case, we can set  $\Delta = \left[ (D\sqrt{2T}(G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2}))^{\frac{2}{3}} (BL_T)^{-\frac{2}{3}} \right]$ . Then we get that

$$\begin{aligned} &\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \right] \\ &\leq 12(BD^2 L_T T)^{\frac{1}{3}} \left( G + \sqrt{2\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right)^{\frac{2}{3}} + (\ln N + 8) \left( 1 + D \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) \right) + 4G^4 \\ &= \mathcal{O} \left( L_T^{\frac{1}{3}} T^{\frac{1}{3}} \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right)^{\frac{2}{3}} + \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right) \right). \end{aligned}$$

**Case 2.** When  $D\sqrt{2T} \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) \leq BL_T$ , we set  $\Delta = 1$  and get

$$\begin{aligned} &\mathbb{E} \left[ \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_t^*) \right] \\ &\leq 8(BD^2 L_T T)^{\frac{1}{3}} \left( G + \sqrt{2\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right)^{\frac{2}{3}} + (\ln N + 8) \left( 1 + D \left( G + 2\sqrt{2\tilde{\sigma}_{1:T}^2} + 2\sqrt{\Sigma_{1:T}^2} \right) \right) + 4G^4 \\ &= \mathcal{O} \left( L_T^{\frac{1}{3}} T^{\frac{1}{3}} \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right)^{\frac{2}{3}} + \left( \sqrt{\tilde{\sigma}_{1:T}^2} + \sqrt{\Sigma_{1:T}^2} \right) \right). \end{aligned}$$

We end the proof by combining the two cases with the upper bounds of  $\tilde{\sigma}_{1:T}^2$  and  $\Sigma_{1:T}^2$ .

## Appendix D. Technical Lemmas

**Lemma 10** (Bregman proximal inequality ([Chen and Teboulle, 1993](#), Lemma 3.2)). *Let  $\mathcal{X}$  be a convex set in a Banach space. Let  $f : \mathcal{X} \rightarrow \mathbb{R}$  be a closed proper convex function on  $\mathcal{X}$ . Given a convex regularizer  $\psi : \mathcal{X} \rightarrow \mathbb{R}$ , we denote its induced Bregman divergence by  $\mathcal{D}_\psi(\cdot, \cdot)$ . Then, any update of the form  $\mathbf{x}_k = \arg \min_{\mathbf{x} \in \mathcal{X}} \{f(\mathbf{x}) + \mathcal{D}_\psi(\mathbf{x}, \mathbf{x}_{k-1})\}$  satisfies the following inequality for any  $\mathbf{u} \in \mathcal{X}$*

$$f(\mathbf{x}_k) - f(\mathbf{u}) \leq \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{k-1}) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_k) - \mathcal{D}_\psi(\mathbf{x}_k, \mathbf{x}_{k-1}).$$

**Lemma 11.** *Let  $l_1, \dots, l_T$  and  $\delta$  be non-negative real numbers. Then  $\sum_{t=1}^T \frac{l_t}{\sqrt{\delta + \sum_{i=1}^t l_i}} \leq 2\sqrt{\delta + \sum_{t=1}^T l_t}$ , where we define  $0/\sqrt{0} = 0$  for simplicity.*

**Lemma 12.** Let  $l_1, \dots, l_T$  and  $\delta$  be non-negative real numbers. Then  $\sum_{t=1}^T \frac{l_t}{\sqrt{\delta + \sum_{i=1}^{t-1} l_i}} \leq 4\sqrt{\delta + \sum_{t=1}^T l_t} + \max_{t \in [T]} l_t$ , where we define  $0/\sqrt{0} = 0$  for simplicity.

**Lemma 13** (Lemma 12 of Hazan et al. (2007)). Let  $A \succeq B \succ 0$  be positive definite matrices. Then  $\langle A^{-1}, A - B \rangle \leq \ln \frac{|A|}{|B|}$ , where  $|A|$  denotes the determinant of matrix  $A$ .

**Lemma 14.** Let  $\mathbf{u}_t \in \mathbb{R}^d$  ( $t = 1, \dots, T$ ), be a sequence of vectors. Define  $S_t = \sum_{\tau=1}^t \mathbf{u}_\tau \mathbf{u}_\tau^\top + \varepsilon I$ , where  $\varepsilon > 0$ . Then  $\sum_{t=1}^T \mathbf{u}_t^\top S_t^{-1} \mathbf{u}_t \leq d \ln \left( 1 + \frac{\sum_{t=1}^T \|\mathbf{u}_t\|_2^2}{d\varepsilon} \right)$ .

**Proof** Using Lemma 13, we have  $\langle A^{-1}, A - B \rangle \leq \ln \frac{|A|}{|B|}$  for any two positive definite matrices  $A \succeq B \succ 0$ . Following the argument of Luo et al. (2016, Theorem 2), we have

$$\begin{aligned} \sum_{t=1}^T \mathbf{u}_t^\top S_t^{-1} \mathbf{u}_t &= \sum_{t=1}^T \langle S_t^{-1}, \mathbf{u}_t \mathbf{u}_t^\top \rangle = \sum_{t=1}^T \langle S_t^{-1}, S_t - S_{t-1} \rangle \leq \sum_{t=1}^T \ln \frac{|S_t|}{|S_{t-1}|} \\ &= \ln \frac{|S_T|}{|S_0|} = \sum_{i=1}^d \ln \left( 1 + \frac{\lambda_i(\sum_{t=1}^T \mathbf{u}_t \mathbf{u}_t^\top)}{\varepsilon} \right) = d \sum_{i=1}^d \frac{1}{d} \ln \left( 1 + \frac{\lambda_i(\sum_{t=1}^T \mathbf{u}_t \mathbf{u}_t^\top)}{\varepsilon} \right) \\ &\leq d \ln \left( 1 + \frac{\sum_{i=1}^d \lambda_i(\sum_{t=1}^T \mathbf{u}_t \mathbf{u}_t^\top)}{d\varepsilon} \right) = d \ln \left( 1 + \frac{\sum_{t=1}^T \|\mathbf{u}_t\|_2^2}{d\varepsilon} \right), \end{aligned}$$

where the last inequality is due to Jensen's inequality. ■

**Lemma 15** (regret analysis of optimistic FTRL (Orabona, 2019, Theorem 7.35)). Let  $V \in \mathbb{R}^d$  be convex, closed, and non-empty. Denote by  $G_t(\mathbf{x}) = \Psi_t(\mathbf{x}) + \sum_{s=1}^{t-1} \ell_s(\mathbf{x})$ . Assume for  $t = 1, \dots, T$  that  $G_t$  is proper and  $\lambda_t$ -strongly convex with respect to  $\|\cdot\|$ ,  $\ell_t$  and  $\tilde{\ell}_t$  proper and convex ( $\tilde{\ell}_t$  is the predicted next loss), and  $\text{int dom } G_t \cap V \neq \{\}$ . Also, assume that  $\partial \ell_t(\mathbf{x}_t)$  and  $\partial \tilde{\ell}_t(\mathbf{x}_t)$  are non-empty. Then there exists  $\tilde{\mathbf{g}}_t \in \partial \tilde{\ell}_t(\mathbf{x}_t)$  for  $t \in [T]$  such that

$$\begin{aligned} &\sum_{t=1}^T \ell_t(\mathbf{x}_t) - \sum_{t=1}^T \ell_t(\mathbf{x}) \\ &\leq \Psi_{T+1}(\mathbf{x}) - \Psi_1(\mathbf{x}_1) + \sum_{t=1}^T \left( \langle \mathbf{g}_t - \tilde{\mathbf{g}}_t, \mathbf{x}_t - \mathbf{x}_{t+1} \rangle - \frac{\lambda_t}{2} \|\mathbf{x}_t - \mathbf{x}_{t+1}\|^2 + \Psi_t(\mathbf{x}_{t+1}) - \Psi_{t+1}(\mathbf{x}_{t+1}) \right) \end{aligned}$$

for all  $\mathbf{g}_t \in \partial \ell_t(\mathbf{x}_t)$ .

**Lemma 16** (Lemma 13 of Zhao et al. (2024)). Let  $a_1, a_2, \dots, a_T, b$  and  $\bar{c}$  be non-negative real numbers and  $a_t \in [0, B]$  for any  $t \in [T]$ . Let the step size be  $c_t = \min \left\{ \bar{c}, \sqrt{\frac{b}{\sum_{s=1}^t a_s}} \right\}$  and  $c_0 = \bar{c}$ . Then, we have  $\sum_{t=1}^T c_{t-1} a_t \leq 2\bar{c}B + 4\sqrt{b \sum_{t=1}^T a_t}$ .



## References

- Jacob Abernethy, Peter L. Bartlett, Alexander Rakhlin, and Ambuj Tewari. Optimal strategies and minimax lower bounds for online convex games. In *Proceedings of the 21st Annual Conference on Learning Theory (COLT)*, pages 415–423, 2008.
- Kwangjun Ahn, Chulhee Yun, and Suvrit Sra. SGD with shuffling: optimal rates without component convexity and large epoch requirements. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 17526–17535, 2020.
- Idan Amir, Idan Attias, Tomer Koren, Roi Livni, and Yishay Mansour. Prediction with corrupted expert advice. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 14315–14325, 2020.
- Peter Auer, Nicolò Cesa-Bianchi, and Claudio Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 64(1):48–75, 2002.
- Dheeraj Baby and Yu-Xiang Wang. Optimal dynamic regret in proper online learning with strongly convex losses and beyond. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1805–1845, 2022.
- Yong Bai, Yu-Jie Zhang, Peng Zhao, Masashi Sugiyama, and Zhi-Hua Zhou. Adapting to online label shift with provable guarantees. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 29960–29974, 2022.
- Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- Nicolo Campolongo and Francesco Orabona. Temporal variability in implicit online learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 12377–12387, 2020.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Nicolò Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- Gong Chen and Marc Teboulle. Convergence analysis of a proximal-like minimization algorithm using bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, 1993.
- Keyi Chen and Francesco Orabona. Generalized implicit follow-the-regularized-leader. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pages 4826–4838, 2023.
- Sijia Chen, Wei-Wei Tu, Peng Zhao, and Lijun Zhang. Optimistic online mirror descent for bridging stochastic and adversarial online convex optimization. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pages 5002–5035, 2023.

- Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*, pages 6.1–6.20, 2012.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7):2121–2159, 2011.
- Pierre Gaillard and Olivier Wintenberger. Efficient online algorithms for fast-rate regret bounds under sparsity. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 7026–7036, 2018.
- Dan Garber, Gal Korcea, and Kfir Y. Levy. Online convex optimization in the random order model. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 3387–3396, 2020.
- Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization I: A generic algorithmic framework. *Siam Journal on Optimization*, 22(4):1469–1492, 2012.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- Elad Hazan and Satyen Kale. Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization. In *Proceedings of the 24th Annual Conference on Learning Theory (COLT)*, pages 421–436, 2011.
- Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- Bin Hu, Peter Seiler, and Anders Rantzer. A unified analysis of stochastic optimization methods using jump system theory and quadratic constraints. In *Proceedings of the 30th Conference on Learning Theory (COLT)*, pages 1157–1189, 2017.
- Shinji Ito. On optimal robustness to adversarial corruption in online decision problems. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, pages 7409–7420, 2021.
- Andrew Jacobsen and Ashok Cutkosky. Parameter-free mirror descent. In *Proceedings of the 35th Conference on Learning Theory (COLT)*, pages 4160–4211, 2022.
- Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 398–406, 2015.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems 26 (NIPS)*, pages 315–323, 2013.

- Pooria Joulani, András György, and Csaba Szepesvári. A modular analysis of adaptive (non-)convex optimization: Optimism, composite objectives, variance reduction, and variational bounds. *Theoretical Computer Science*, 808:108–138, 2020.
- Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015.
- Tomer Koren and Kfir Levy. Fast rates for exp-concave empirical risk minimization. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pages 1477–1485, 2015.
- Guanghui Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, 2012.
- Zachary C. Lipton, Yu-Xiang Wang, and Alexander J. Smola. Detecting and correcting for label shift with black box predictors. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 3128–3136, 2018.
- Haipeng Luo, Alekh Agarwal, Nicolò Cesa-Bianchi, and John Langford. Efficient second order online learning by sketching. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 902–910, 2016.
- Mehrdad Mahdavi, Lijun Zhang, and Rong Jin. Lower and upper bounds on the generalization of stochastic exponentially concave optimization. In *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*, page 1305–1320, 2015.
- H. Brendan McMahan and Matthew J. Streeter. Adaptive bound optimization for online convex optimization. In *Proceedings of the 23rd Conference on Learning Theory (COLT)*, pages 244–256, 2010.
- Arkadi Nemirovski. Prox-method with rate of convergence  $O(1/t)$  for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2005.
- Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- Gergely Neu and Lorenzo Rosasco. Iterate averaging as regularization for stochastic gradient descent. In *Proceedings of the 31st Annual Conference on Learning Theory (COLT)*, pages 3222–3242, 2018.
- Francesco Orabona. A modern introduction to online learning. *ArXiv preprint*, arXiv:1912.13213, 2019.
- Francesco Orabona, Nicolo Cesa-Bianchi, and Claudio Gentile. Beyond logarithmic bounds in online learning. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 823–831, 2012.

- Erik Ordentlich and Thomas M. Cover. The cost of achieving the best portfolio in hindsight. *Mathematics of Operations Research*, 23(4):960–982, 1998.
- Yu-Yang Qian, Yong Bai, Zhen-Yu Zhang, Peng Zhao, and Zhi-Hua Zhou. Handling new class in online label shift. In *Proceedings of the 23rd IEEE International Conference on Data Mining (ICDM)*, pages 1283–1288, 2023.
- Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Proceedings of the 26th Conference on Learning Theory (COLT)*, pages 993–1019, 2013.
- Sarah Sachs, Hedi Hadiji, Tim van Erven, and Cristóbal A Guzmán. Between stochastic and adversarial online convex optimization: Improved regret bounds via smoothness. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 691–702, 2022.
- Sarah Sachs, Hedi Hadiji, Tim van Erven, and Cristobal Guzman. Accelerated rates between stochastic and adversarial online convex optimization. *ArXiv preprint*, arXiv:2303.03272, 2023.
- Shai Shalev-Shwartz. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, The Hebrew University of Jerusalem, 2007.
- Shai Shalev-Shwartz. SDCA without duality, regularization, and individual convexity. In *Proceedings of the 33th International Conference on Machine Learning (ICML)*, pages 747–754, 2016.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*, page 5, 2009.
- Uri Sherman, Tomer Koren, and Yishay Mansour. Optimal rates for random order online optimization. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, pages 2097–2108, 2021.
- Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low-noise and fast rates. In *Advances in Neural Information Processing Systems 23 (NIPS)*, pages 2199–2207, 2010.
- Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pages 2989–2997, 2015.
- Jing Wang, Miao Yu, Peng Zhao, and Zhi-Hua Zhou. Learning with adaptive resource allocation. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, page to appear, 2024.
- Lijun Zhang and Zhi-Hua Zhou. Stochastic approximation of smooth and strongly convex functions: Beyond the  $O(1/T)$  convergence rate. In *Proceedings of the 32nd Annual Conference on Learning Theory (COLT)*, pages 3160–3179, 2019.

- Lijun Zhang, Mehrdad Mahdavi, and Rong Jin. Linear convergence with condition number independent access of full gradients. In *Advance in Neural Information Processing Systems 26 (NIPS)*, pages 980–988, 2013.
- Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 1323–1333, 2018.
- Lijun Zhang, Guanghui Wang, Jinfeng Yi, and Tianbao Yang. A simple yet universal strategy for online convex optimization. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pages 26605–26623, 2022.
- Yu-Jie Zhang, Peng Zhao, and Zhi-Hua Zhou. A simple online algorithm for competing with dynamic comparators. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 390–399, 2020.
- Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 12510–12520, 2020.
- Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Adaptivity and non-stationarity: Problem-dependent dynamic regret for online convex optimization. *Journal of Machine Learning Research*, 25(98):1 – 52, 2024.
- Zhi-Hua Zhou. Learnability with time-sharing computational resource concerns. *National Science Review*, nwae204, 2024.
- Julian Zimmert and Yevgeny Seldin. Tsallis-INF: An optimal algorithm for stochastic and adversarial bandits. *Journal of Machine Learning Research*, 22(28):1–49, 2021.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 928–936, 2003.