

# Non-stationary Projection-Free Online Learning with Dynamic and Adaptive Regret Guarantees

Yibo Wang<sup>1,2</sup>, Wenhao Yang<sup>1,2</sup>, Wei Jiang<sup>1</sup>, Shiyin Lu<sup>3</sup>,  
Bing Wang<sup>3</sup>, Haihong Tang<sup>3</sup>, Yuanyu Wan<sup>4</sup>, Lijun Zhang<sup>1,2,\*</sup>

<sup>1</sup> National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China

<sup>2</sup> School of Artificial Intelligence, Nanjing University, Nanjing 210023, China

<sup>3</sup> Alibaba Group, Hangzhou 310052, China

<sup>4</sup> School of Software Technology, Zhejiang University, Ningbo 315048, China,

## Abstract

Projection-free online learning has drawn increasing interest due to its efficiency in solving high-dimensional problems with complicated constraints. However, most existing projection-free online methods focus on minimizing the static regret, which unfortunately fails to capture the challenge of changing environments. In this paper, we investigate non-stationary projection-free online learning, and choose *dynamic regret* and *adaptive regret* to measure the performance. Specifically, we first provide a novel dynamic regret analysis for an existing projection-free method named BOGD<sub>IP</sub>, and establish an  $\mathcal{O}(T^{3/4}(1 + P_T))$  dynamic regret bound, where  $P_T$  denotes the path-length of the comparator sequence. Then, we improve the upper bound to  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  by running multiple BOGD<sub>IP</sub> algorithms with different step sizes in parallel, and tracking the best one on the fly. Our results are the first general-case dynamic regret bounds for projection-free online learning, and can recover the existing  $\mathcal{O}(T^{3/4})$  static regret by setting  $P_T = 0$ . Furthermore, we propose a projection-free method to attain an  $\tilde{\mathcal{O}}(\tau^{3/4})$  adaptive regret bound for any interval with length  $\tau$ , which nearly matches the static regret over that interval. The essential idea is to maintain a set of BOGD<sub>IP</sub> algorithms dynamically, and combine them by a meta algorithm. Moreover, we demonstrate that it is also equipped with an  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  dynamic regret bound. Finally, empirical studies verify our theoretical findings.

## Introduction

In many online learning problems, the decision constraint sets are often high-dimensional and complicated, rendering optimization over such sets challenging. In these cases, traditional projection-based methods, such as Online Gradient Descent (OGD) (Zinkevich 2003), often suffer heavy computational costs due to the time-consuming or even intractable projection operations. To address this limitation, projection-free online methods, which replace projections with less expensive computations (e.g., linear optimizations) and thus can be implemented efficiently in many cases of interest, have drawn considerable attention in the online learning community (Hazan and Kale 2012; Garber and Hazan

2016; Huang et al. 2016; Levy and Krause 2019; Chen, Zhang, and Karbasi 2019; Hazan and Minasyan 2020; Wan, Tu, and Zhang 2020; Molinaro 2020; Kalhan et al. 2021; Wan and Zhang 2021; Wan, Xue, and Zhang 2021; Kretzu and Garber 2021; Garber and Kretzu 2022; Mhammedi 2022; Wan et al. 2022; Wang et al. 2023a; Lu et al. 2023; Wan, Zhang, and Song 2023; Garber and Kretzu 2023).

The studies of projection-free online methods follow the framework of Online Convex Optimization (OCO), which can be regarded as a repeated game between a learner against an adversary (Shalev-Shwartz 2012). At round  $t$ , the learner chooses an action  $\mathbf{x}_t$  from a convex domain set  $\mathcal{K}$ , and then suffers an instantaneous loss  $f_t(\mathbf{x}_t)$ , where the convex loss function  $f_t(\cdot) : \mathcal{K} \rightarrow \mathbb{R}$  is chosen by the adversary. The majority of existing projection-free methods, e.g., Online Frank-Wolfe (OFW) (Hazan and Kale 2012), minimize the static regret:

$$\text{Regret}_T = \sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T f_t(\mathbf{x}), \quad (1)$$

which benchmarks the cumulative loss of the online method against that of the best fixed action in hindsight. However, in real-world scenarios such as online recommendation and online traffic scheduling (Hazan 2016), this static metric is unsuitable as the environments are non-stationary and the best action is drifting over time. To tackle this issue, two novel metrics: dynamic regret and adaptive regret, are proposed independently (Zinkevich 2003; Hazan and Seshadhri 2007; Daniely, Gonen, and Shalev-Shwartz 2015).

The dynamic regret stems from Zinkevich (2003), who defines

$$\text{D-Regret}_T(\mathbf{u}_1, \dots, \mathbf{u}_T) = \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t), \quad (2)$$

where  $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{K}$  are any possible comparators. Unfortunately, obtaining a sublinear dynamic regret with arbitrarily varying sequences is impossible. As a result, to establish a meaningful bound, it is common to introduce some regularities of the comparator sequence, such as the path-length

$$P_T = \sum_{t=2}^T \|\mathbf{u}_{t-1} - \mathbf{u}_t\|_2.$$

\*Corresponding author, zhanglj@lamda.nju.edu.cn  
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The adaptive regret is originally introduced by Hazan and Seshadhri (2007), and further strengthened by Daniely, Gonen, and Shalev-Shwartz (2015). Formally, it is defined as

$$\text{SA-Regret}_T(\tau) = \max_{[s, s+\tau-1] \subseteq [T]} \left\{ \sum_{t=s}^{s+\tau-1} f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{K}} \sum_{t=s}^{s+\tau-1} f_t(\mathbf{x}) \right\}, \quad (3)$$

which is the maximum static regret over any interval with the length  $\tau$ . Since in different intervals the best actions can be different, (3) essentially measures the performance of the online method against changing comparators.

In the literature, only a few projection-free online methods (Kalhan et al. 2021; Wan, Xue, and Zhang 2021; Wan, Zhang, and Song 2023) have investigated dynamic regret minimization, but all of them focus on the worst case of (2), where  $\mathbf{u}_t \in \arg \min_{\mathbf{u} \in \mathcal{K}} f_t(\mathbf{u})$  is a minimizer of  $f_t(\cdot)$ . However, the worst-case dynamic regret is too pessimistic, and cannot recover the static regret bound of previous methods (Hazan and Kale 2012; Hazan and Minasyan 2020). Besides, there exist two studies (Garber and Kretzu 2022; Lu et al. 2023) that propose projection-free methods for adaptive regret minimization. However, Garber and Kretzu (2022) only consider a weak form of (3) which does not respect short intervals well, and the method of Lu et al. (2023) could be time-consuming in many popular domains, e.g., bounded trace norm matrices and matroid polytopes (Mhammedi 2022).

In this paper, we choose (2) and (3) as the performance metrics, and propose two novel methods for non-stationary projection-free online learning. Specifically, in the dynamic regret minimization, we first establish a novel dynamic regret bound of  $\mathcal{O}(T^{3/4}(1 + P_T))$  for an existing projection-free variant of Online Gradient Descent, termed as BOGD<sub>IP</sub> (Garber and Kretzu 2022).<sup>1</sup> Then, we improve the upper bound to  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  by proposing a two-layer method named POLD, which maintains multiple BOGD<sub>IP</sub> algorithms with different step sizes, and tracks the best one on the fly by a meta algorithm. In the adaptive regret minimization, we propose a novel projection-free method named POLA, which attains an  $\tilde{\mathcal{O}}(\tau^{3/4})$  adaptive regret bound for any interval with the length  $\tau$ . The key idea is to construct a set of intervals dynamically, run a BOGD<sub>IP</sub> algorithm that aims to minimize the static regret for each interval, and combine them by a meta algorithm. Moreover, we show that our POLA can also minimize the dynamic regret, and ensures an  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  bound. Notably, although POLA can achieve the same dynamic regret bound as POLD, the latter one is still valuable in the sense that it employs a clearer structure and a simpler meta algorithm, rendering it much easier to comprehend and implement.

**Contributions.** We summarize the contributions of this work below.

<sup>1</sup>In Garber and Kretzu (2022), BOGD<sub>IP</sub> is referred to as Blocked Online Gradient Descent with Linear Optimization Oracle (LOO-BOGD).

- For dynamic regret, we first provide a novel analysis for BOGD<sub>IP</sub> (Garber and Kretzu 2022), and establish an  $\mathcal{O}(T^{3/4}(1 + P_T))$  dynamic regret. Then, we improve this bound to  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  by proposing a two-layer method named POLD. Note that the obtained bounds can recover the previous  $\mathcal{O}(T^{3/4})$  static regret (Hazan and Kale 2012) by setting  $P_T = 0$ . To the best of our knowledge, these are the *first* general-case dynamic regret bounds in projection-free online learning.
- For adaptive regret, based on BOGD<sub>IP</sub>, we propose a novel projection-free method named POLA and obtain an  $\tilde{\mathcal{O}}(\tau^{3/4})$  adaptive regret which nearly matches previous static results. Moreover, we show that POLA can also ensure an  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  dynamic regret bound. In other words, it can minimize dynamic regret and adaptive regret simultaneously.
- We conduct experiments on practical problems to verify our theoretical findings in dynamic regret and adaptive regret minimization. Empirical results demonstrate the advantage of proposed methods.

## Related Work

In this section, we briefly review related work in dynamic regret and adaptive regret.

### Dynamic Regret

In the literature, dynamic regret has two different forms. One is the general case (2) introduced by Zinkevich (2003), who defines it as the difference between the cumulative loss of the online method and that of *any* possible comparator sequence. In this seminal work, Zinkevich (2003) establishes the first general-case bound of  $\mathcal{O}(\sqrt{T}(1 + P_T))$  for OGD. Later, Zhang, Lu, and Zhou (2018) improve the upper bound to  $\mathcal{O}(\sqrt{T}(1 + P_T))$ , motivated by the strategy of maintaining multiple step sizes in MetaGrad (van Erven and Koolen 2016; Mhammedi, Koolen, and van Erven 2019; van Erven, Koolen, and van der Hoeven 2021). In recent years, several studies have further investigated the general-case dynamic regret by leveraging the curvature of loss functions, such as exponential concavity (Baby and Wang 2021) and strong convexity (Baby and Wang 2022).

The other is the worst case of (2), which specializes the comparators as the minimizers of loss functions (Besbes, Gur, and Zeevi 2015; Jadbabaie et al. 2015; Mokhtari et al. 2016; Yang et al. 2016; Baby and Wang 2019):

$$\text{D-Regret}_T(\mathbf{u}_1^*, \dots, \mathbf{u}_T^*) = \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t^*), \quad (4)$$

where  $\mathbf{u}_t^* \in \arg \min_{\mathbf{u} \in \mathcal{K}} f_t(\mathbf{u})$  is a minimizer of  $f_t(\cdot)$ . However, as pointed out by Zhang, Lu, and Zhou (2018), the worst-case dynamic regret (4) is too pessimistic and could lead to overfitting in the stationary problems.

In projection-free online learning, several studies (Kalhan et al. 2021; Wan, Xue, and Zhang 2021; Wan, Zhang, and Song 2023) have investigated the dynamic regret recently, but they only consider the worst-case formulation

Method	Loss	Operation	Metric	Bound
Kalhan et al. (2021)	smooth & convex	LO	WD-R	$\mathcal{O}(\sqrt{T}(1 + F_T + \sqrt{D_T}))$
Wan, Zhang, and Song (2023)	smooth & convex	LO	WD-R	$\mathcal{O}(\sqrt{T}(1 + F_T))$
Wan, Xue, and Zhang (2021)	convex	LO	WD-R	$\mathcal{O}(\max\{T^{2/3}F_T^{1/3}, \sqrt{T}\})$
	strongly convex	LO	WD-R	$\mathcal{O}(\max\{\sqrt{TF_T \log T}, \log T\})$
<b>BOGD<sub>IP</sub></b> (this work)	convex	LO	D-R	$\mathcal{O}(T^{3/4}(1 + P_T))$
<b>POLD</b> (this work)	convex	LO	D-R	$\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$
<b>POLA</b> (this work)	convex	LO	D-R	$\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$
Garber and Kretzu (2022)	convex	LO	A-R	$\mathcal{O}(T^{3/4})$
Lu et al. (2023)	convex	MO	SA-R	$\tilde{\mathcal{O}}(\sqrt{\tau})$
<b>POLA</b> (this work)	convex	LO	SA-R	$\tilde{\mathcal{O}}(\tau^{3/4})$

Table 1: Summary of existing methods in non-stationary projection-free online learning. Abbreviations: linear optimization  $\rightarrow$  LO, membership operation  $\rightarrow$  MO, worst-case dynamic regret (4)  $\rightarrow$  WD-R, general-case dynamic regret (2)  $\rightarrow$  D-R, weak adaptive regret (5)  $\rightarrow$  A-R, strongly adaptive regret (3)  $\rightarrow$  SA-R.  $\tau$  denotes the length of an interval  $I$ , i.e.,  $\tau = |I|$ .

(4). Specifically, for smooth and convex losses, Kalhan et al. (2021) establish an  $\mathcal{O}(\sqrt{T}(1 + F_T + \sqrt{D_T}))$  worst-case bound<sup>2</sup>, where  $F_T$  denotes the functional variation (Besbes, Gur, and Zeevi 2015)

$$F_T = \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{K}} |f_t(\mathbf{x}) - f_{t-1}(\mathbf{x})|,$$

and  $D_T$  denotes the gradient variation (Chiang et al. 2012)

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

For convex losses and strongly convex losses, Wan, Xue, and Zhang (2021) develop the  $\mathcal{O}(\max\{T^{2/3}F_T^{1/3}, \sqrt{T}\})$  and  $\mathcal{O}(\max\{\sqrt{TF_T \log T}, \log T\})$  worst-case bounds, respectively. Very recently, Wan, Zhang, and Song (2023) refine the analysis of Kalhan et al. (2021), achieving an improved  $\mathcal{O}(\sqrt{T}(1 + F_T))$  bound. However, due to the weakness of (4), their bounds can be very loose for any other comparators, and cannot recover the static regret of existing methods, e.g.,  $\mathcal{O}(T^{3/4})$  for convex losses (Hazan and Kale 2012).

### Adaptive Regret

Prior studies in adaptive regret minimization mainly focus on the setting of Prediction with Expert Advice (PEA) (Littlestone and Warmuth 1994; Freund et al. 1997; György, Linder, and Lugosi 2012; Luo and Schapire 2015; Adamskiy et al. 2016), and OCO (Hazan and Seshadhri 2007; Daniely, Gonen, and Shalev-Shwartz 2015; Jun et al. 2017a,b; Zhang, Liu, and Zhou 2019). In this section, we specifically introduce the related work of the latter one.

Hazan and Seshadhri (2007) first introduce the notion of adaptive regret, but in a weak form:

$$\text{A-Regret}_T = \max_{[s,e] \subseteq [T]} \left\{ \sum_{t=s}^e f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{K}} \sum_{t=s}^e f_t(\mathbf{x}) \right\}, \tag{5}$$

<sup>2</sup>A recent study (Zhou, Xu, and Tzoumas 2023) obtains the same regret bound while removing the smoothness assumption.

which is the maximum static regret over any contiguous interval. To minimize (5), they propose Follow the Leading History (FLH) with an  $\mathcal{O}(d \log^2 T)$  weak adaptive regret bound for exponentially concave losses where  $d$  denotes the dimensionality. However, (5) could be dominated by long intervals and hence, cannot respect short intervals well. For example, one may obtain an  $\mathcal{O}(\sqrt{T})$  weak adaptive regret for OGD, but this is vacuous for the intervals with length  $o(\sqrt{T})$  (Hazan 2016). For this reason, Daniely, Gonen, and Shalev-Shwartz (2015) put forth the (strongly) adaptive regret (3), and design a two-layer algorithm named Strongly Adaptive Online Learner (SAOL). The basic idea is first to construct a set of Geometric Covering (GC) intervals and for each interval, run an OGD algorithm that can obtain the optimal static regret. Then, SAOL combines the actions of these OGD algorithms by a meta algorithm. We observe that the technique of constructing GC intervals can be traced back to the prior studies (Willems and Krom 1997; György, Linder, and Lugosi 2012).

In projection-free online learning, Garber and Kretzu (2022) study the weak version of adaptive regret (5), and propose a projection-free extension of OGD named BOGD<sub>IP</sub> with an  $\mathcal{O}(T^{3/4})$  bound. Unfortunately, due to the limitation of (5), their bound does not respect short intervals well. Very recently, following the framework of SAOL, Lu et al. (2023) propose a novel two-layer method to minimize (3). Different from previous projection-free algorithms, e.g., OFW (Hazan and Kale 2012), their method circumvents the projections with membership operations (Mhammedi 2022). However, such operations could be inefficient in many practical scenarios, e.g., bounded trace norm matrices and matroid polytopes (Mhammedi 2022). Besides, in each round, their method need to perform  $\mathcal{O}(\log T)$  membership operations for each expert algorithm, which brings heavy computational costs when  $T$  is large.

### Summary

While a few studies have investigated non-stationary projection-free online learning (see Table 1 for details), they

are still unsatisfactory in the following aspects:

- In the dynamic regret minimization, there is no study optimizing the general-case form (2), which is more challenging since it needs to build a universal guarantee over any comparator sequences.
- In the adaptive regret minimization, although Lu et al. (2023) have established bounds for (3), their method is based on the membership operations, instead of the more popular linear optimizations.

## Main Results

In this section, we first introduce the basic assumptions. Then, we present our proposed methods as well as their theoretical guarantees in dynamic regret and adaptive regret minimization. The proofs for theoretical results can be found in the full version (Wang et al. 2023b).

### Assumptions

Similar to previous studies on OCO, we adopt the following standard assumptions (Shalev-Shwartz 2012; Hazan 2016).

**Assumption 1.** The convex decision set  $\mathcal{K}$  contains the origin  $\mathbf{0}$ , and belongs to an Euclidean ball  $R\mathcal{B}$  with the diameter  $D = 2R$ , i.e.,

$$\forall \mathbf{x}, \mathbf{x}' \in \mathcal{K}, \|\mathbf{x} - \mathbf{x}'\|_2 \leq D. \quad (6)$$

**Assumption 2.** At each round  $t$ , the loss function  $f_t(\cdot)$  is  $G$ -Lipschitz over  $\mathcal{K}$ , i.e.,

$$\forall \mathbf{x}, \mathbf{y} \in \mathcal{K}, |f_t(\mathbf{x}) - f_t(\mathbf{y})| \leq G\|\mathbf{x} - \mathbf{y}\|_2. \quad (7)$$

**Assumption 3.** At each round  $t$ , the loss function  $f_t(\cdot)$  is convex over  $\mathcal{K}$ , i.e.,

$$\forall \mathbf{x}, \mathbf{y} \in \mathcal{K}, f_t(\mathbf{y}) \geq f_t(\mathbf{x}) + \nabla f_t(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}). \quad (8)$$

**Assumption 4.** At each round  $t$ , the loss function value  $f_t(\mathbf{x})$  belongs to  $[0, 1]$  for any  $\mathbf{x} \in \mathcal{K}$ , i.e.,

$$\forall \mathbf{x} \in \mathcal{K}, 0 \leq f_t(\mathbf{x}) \leq 1. \quad (9)$$

### Projection-Free Dynamic Regret

We first revisit BOGD<sub>IP</sub> (Garber and Kretzu 2022), of which the key idea is to replace the projection operation with an infeasible projection oracle  $\mathcal{O}_{IP}$ , defined as following.

**Definition 1.** Let  $\mathcal{O}_{IP}$  be an infeasible projection oracle over  $\mathcal{K} \subseteq R\mathcal{B}$ , and  $\epsilon$  be the error tolerance. Then, for any input points  $(\mathbf{x}_0, \mathbf{y}_0) \in \mathcal{K} \times \mathbb{R}^d$ , the infeasible projection oracle returns

$$\mathbf{x}, \tilde{\mathbf{y}} = \mathcal{O}_{IP}(\mathcal{K}, \epsilon, \mathbf{x}_0, \mathbf{y}_0),$$

where  $(\mathbf{x}, \tilde{\mathbf{y}}) \in \mathcal{K} \times R\mathcal{B}$ , and  $\|\mathbf{x} - \tilde{\mathbf{y}}\|_2 \leq \sqrt{3}\epsilon$  and  $\forall \mathbf{z} \in \mathcal{K}, \|\tilde{\mathbf{y}} - \mathbf{z}\|_2 \leq \|\mathbf{y}_0 - \mathbf{z}\|_2$ .

**Remark:**  $\mathcal{O}_{IP}$  can be implemented efficiently by solving linear optimizations. We briefly introduce this implementation in the supplementary material, and refer interested readers to Garber and Kretzu (2022) for a deeper comprehension.

Besides, BOGD<sub>IP</sub> utilizes the blocking technique (Garber and Kretzu 2020; Hazan and Minasyan 2020), which divides

---

### Algorithm 1: Blocked Online Gradient Descent with Infeasible Projections (BOGD<sub>IP</sub>)

---

**Input:** Number of rounds  $T$ , domain set  $\mathcal{K}$ , step size  $\eta$ , infeasible projection oracle  $\mathcal{O}_{IP}$

**Initialization:** Choose arbitrary point  $\mathbf{x}_1 \in \mathcal{K}$  and set  $\tilde{\mathbf{y}}_1 = \mathbf{x}_1$ ,  $m = 1$ , block size  $K = \eta^{-2/3}$  and error tolerance  $\epsilon = \eta^{2/3}$ .

- 1: **for**  $t = 1$  to  $T$  **do**
  - 2:   Submit  $\mathbf{x}_t = \mathbf{x}_m$ , observe  $f_t(\mathbf{x}_t)$  and obtain  $\nabla f_t(\mathbf{x}_t)$
  - 3:   **if**  $t \bmod K = 0$  **then**
  - 4:     Update  $\mathbf{y}_{m+1}$  according to (10)
  - 5:     Set  $\mathbf{x}_{m+1}, \tilde{\mathbf{y}}_{m+1}$  according to (11), and  $m = \lfloor t/K \rfloor + 1$
  - 6:   **end if**
  - 7: **end for**
- 

the time horizon  $T$  into equally-sized blocks and only conducts updating at the end of each block. In other words, for each block  $m$ , BOGD<sub>IP</sub> maintains  $(\mathbf{x}_m, \tilde{\mathbf{y}}_m) \in \mathcal{K} \times R\mathcal{B}$ , and updates them at the last round of block  $m$ . To be precise, BOGD<sub>IP</sub> first performs gradient descent on  $\tilde{\mathbf{y}}_m$  with the step size  $\eta$ :

$$\mathbf{y}_{m+1} = \tilde{\mathbf{y}}_m - \eta \sum_{r=(m-1)K+1}^{mK} \nabla f_r(\mathbf{x}_m), \quad (10)$$

where  $K$  is the block size and  $\sum_{r=(m-1)K+1}^{mK} \nabla f_r(\mathbf{x}_m)$  is the sum of all gradients during the block  $m$ . Then, BOGD<sub>IP</sub> invokes  $\mathcal{O}_{IP}$  to obtain  $\mathbf{x}_{m+1}$  and  $\tilde{\mathbf{y}}_{m+1}$  for the next block:

$$\mathbf{x}_{m+1}, \tilde{\mathbf{y}}_{m+1} = \mathcal{O}_{IP}(\mathcal{K}, \epsilon, \mathbf{x}_m, \mathbf{y}_{m+1}). \quad (11)$$

With appropriate parameters, we can prove that BOGD<sub>IP</sub> requires  $\mathcal{O}(T^{1/2})$  invocations of  $\mathcal{O}_{IP}$ , and each invocation solves  $\mathcal{O}(T^{1/2})$  linear optimizations. As a result, there are at most  $\mathcal{O}(T)$  linear optimizations for the time horizon  $T$ . We summarize the detailed procedure in Algorithm 1.

In the prior study, Garber and Kretzu (2022) have investigated the weak adaptive regret (5). Different from them, we focus on minimizing the general-case dynamic regret (2) and establish an  $\mathcal{O}(T^{3/4}(1 + P_T))$  bound for BOGD<sub>IP</sub> as shown in Theorem 1. The intuition lies in that BOGD<sub>IP</sub> is a projection-free variant of OGD, which is very suitable for dynamic regret minimization (Zinkevich 2003).

**Theorem 1.** Let  $\eta = T^{-3/4}$ ,  $K = \eta^{-2/3} = T^{1/2}$  and  $\epsilon = \eta^{2/3} = T^{-1/2}$ . Under Assumptions 1, 2 and 3, Algorithm 1 guarantees

$$\begin{aligned} \text{D-Regret}_T(\mathbf{u}_1, \dots, \mathbf{u}_T) &\leq \mathcal{O}\left(\eta^{1/3}T + \eta^{-1}(1 + P_T)\right) \\ &= \mathcal{O}\left(T^{3/4}(1 + P_T)\right). \end{aligned}$$

Moreover, the overall number of solving linear optimizations is  $\mathcal{O}(T)$ .

**Remark:** Our result is the first general-case dynamic regret bound in projection-free online learning, and can automatically adapt to the nature of environments. For example,

---

**Algorithm 2: Projection-free Online Learning with Dynamic Regret (POLD)**

---

**Input:** A learning rate  $\alpha$ , a set  $\mathcal{H}$  containing step size  $\eta_i$  for each expert  $E_i$   
**Initialization:** Activate a set of experts  $\{E_i \mid \eta_i \in \mathcal{H}\}$  by invoking BOGD<sub>IP</sub> for each  $\eta_i \in \mathcal{H}$ .

- 1: For each expert  $E_i$ , set  $w_1^i = \frac{C}{i(i+1)}$  where  $C = 1 + \frac{1}{N}$
- 2: **for**  $t = 1$  to  $T$  **do**
- 3:   Receive  $\mathbf{x}_t^i$  from each expert  $E_i$
- 4:   Compute  $\mathbf{x}_t$  according to (14)
- 5:   Submit  $\mathbf{x}_t$ , and update the weight  $w_{t+1}^i$  for each expert  $E_i$  according to (15)
- 6:   Send  $f_t(\cdot)$  to each expert  $E_i$
- 7: **end for**

---

when the comparators are fixed (i.e.,  $P_T = 0$ ), our dynamic regret degenerates to  $\mathcal{O}(T^{3/4})$ , which matches the static regret bound of Hazan and Kale (2012). To be specific, we have the following corollary, which can also be derived from Theorem 3 of Garber and Kretzu (2022)

**Corollary 1.** *Under Assumptions 1, 2 3, Algorithm 1 with the same parameter setting in Theorem 1 guarantees a static regret bound of*

$$\text{Regret}_T \leq \mathcal{O}(T^{3/4}). \quad (12)$$

### Improved Projection-Free Dynamic Regret

Note that the linear dependency on  $P_T$  in Theorem 1 is too loose and the obtained bound can be vacuous with  $P_T = \Omega(T^{1/4})$ . To address this issue, we propose a two-layer method, termed as Projection-free Online Learning with Dynamic Regret (POLD), with a tighter bound of  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$ . To help understanding, we first briefly introduce the motivation behind POLD.

Let us consider a given sequence  $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_T \in \mathcal{K}$  with the path-length  $\tilde{P}_T = \sum_{t=2}^T \|\tilde{\mathbf{u}}_{t-1} - \tilde{\mathbf{u}}_t\|_2$ . According to Theorem 1, we can choose the step size  $\tilde{\eta} = \mathcal{O}(T^{-3/4}(1 + \tilde{P}_T)^{3/4})$  and achieve a tighter  $\mathcal{O}(T^{3/4}(1 + \tilde{P}_T)^{1/4})$  bound. This indicates that if the path-length is known, we can actually tune the step size to obtain an improved bound. To deal with the uncertainty of the path-length, we adopt the strategy of maintaining multiple step sizes (van Erven and Koolen 2016; Zhang, Lu, and Zhou 2018), and leverage the two-layer structure: running multiple BOGD<sub>IP</sub> algorithms with different step sizes and combining them by a meta algorithm. In the following, we describe the detailed procedure.

First, we create a set of step sizes

$$\mathcal{H} = \left\{ \eta_i = 2^{i-1} \left( \frac{7D^2}{2G^2T} \right)^{3/4} \mid i = 1, \dots, N \right\}, \quad (13)$$

where  $N = \lceil \frac{3}{4} \log_2(1 + 4T/7) \rceil + 1$ . Then, we activate a set of experts  $\{E_i \mid \eta_i \in \mathcal{H}\}$ , each of which is an instance of BOGD<sub>IP</sub> with the step size  $\eta_i$  chosen from  $\mathcal{H}$ . For each expert  $E_i$ , we initiate its weight  $w_1^i = \frac{C}{i(i+1)}$  where

---

**Algorithm 3: Projection-free Online Learning with Adaptive Regret (POLA)**

---

- 1: **for**  $t = 1$  to  $T$  **do**
- 2:   **for**  $I \in \mathcal{C}_t$  **do**
- 3:     Create an expert  $E_I$  which runs BOGD<sub>IP</sub> from an arbitrary initial point with  $\eta = |I|^{-3/4}$
- 4:     For the expert  $E_I$ , set  $R_{t-1,I} = C_{t-1,I} = 0$
- 5:     Add expert  $E_I$  to the set of active experts  $\mathcal{A}_t$
- 6:   **end for**
- 7:   From  $\mathcal{A}_t$ , remove all experts who end at the round  $t$
- 8:   Receive the action  $\mathbf{x}_{t,I}$  of each expert  $E_I \in \mathcal{A}_t$  and calculate its weight  $w_{t,I}$  according to (17)
- 9:   Submit  $\mathbf{x}_t$  defined in (18) and then receive  $f_t(\cdot)$
- 10:   For each  $E_I \in \mathcal{A}_t$ , update

$$R_{t,I} = R_{t-1,I} + f_t(\mathbf{x}_t) - f_t(\mathbf{x}_{t,I})$$

$$C_{t,I} = C_{t-1,I} + |f_t(\mathbf{x}_t) - f_t(\mathbf{x}_{t,I})|$$

- 11:   Send  $f_t(\cdot)$  to each expert  $E_I \in \mathcal{A}_t$
- 12: **end for**

---

$C = 1 + \frac{1}{N}$ . Next, inspired by the Hedge algorithm (Freund and Schapire 1997), we combine the actions of experts in a weighted-average fashion. Concretely, in each round  $t$ , POLD receives the action  $\mathbf{x}_t^i$  from expert  $E_i$ , and computes the weighted average action:

$$\mathbf{x}_t = \sum_{i \in \mathcal{H}} w_t^i \mathbf{x}_t^i, \quad (14)$$

where  $w_t^i$  is the weight assigned to  $E_i$ . After that, POLD updates the weight of  $E_i$  by

$$w_{t+1}^i = \frac{w_t^i e^{-\alpha f_t(\mathbf{x}_t^i)}}{\sum_{\mu \in \mathcal{H}} w_t^\mu e^{-\alpha f_t(\mathbf{x}_t^\mu)}}, \quad (15)$$

where  $\alpha$  denotes the learning rate of the meta algorithm. Finally, POLD reveals the function  $f_t(\cdot)$  to all experts so that they can update their actions for the next round. We summarize all the procedure in Algorithm 2, and present the following theorem.

**Theorem 2.** *Let  $\alpha = \sqrt{8/T}$  and  $\mathcal{H}$  be defined as (13). Under Assumptions 1, 2, 3 and 4, Algorithm 2 guarantees*

$$\text{D-Regret}_T(\mathbf{u}_1, \dots, \mathbf{u}_T) \leq \mathcal{O}\left(T^{3/4}(1 + P_T)^{1/4}\right).$$

**Remark:** Compared with the upper bound in Theorem 1, the dependence on the path-length is reduced from  $P_T$  to  $P_T^{1/4}$ .

### Projection-Free Adaptive Regret

As mentioned before, besides the dynamic regret (2), there do exist another metric called (strongly) adaptive regret (3) in the non-stationary environments. In this section, we proceed to investigate minimizing (3) and present Projection-free Online Learning with Adaptive Regret (POLA). Following existing studies on adaptive regret (Hazan and Seshadhri 2007; Daniely, Gonen, and Shalev-Shwartz 2015),




Figure 1: Geometric Covering (GC) intervals. In the figure, each interval is denoted by  $[ \ ]$ .

POLA contains three parts: an expert algorithm, a set of intervals, and a meta algorithm. In the following, we specify them separately.

First, we take BOGD<sub>IP</sub> as the expert algorithm, since it is projection-free and ensures an  $\mathcal{O}(|I|^{3/4})$  static regret for a given interval  $I$  as shown in Corollary 1. Then, we build the GC intervals (Daniely, Gonen, and Shalev-Shwartz 2015) shown in Figure 1:

$$\mathcal{I} = \bigcup_{k \in \mathbb{N} \cup \{0\}} \mathcal{I}_k, \quad \mathcal{I}_k = \{[i \cdot 2^k, (i + 1) \cdot 2^k - 1] : i \in \mathbb{N}\}. \tag{16}$$

For each interval  $I$ , we maintain an instance of BOGD<sub>IP</sub>, denoted as the expert  $E_I$ , to minimize the static regret over that interval. According to Corollary 1, we set the step size  $\eta = |I|^{-3/4}$  to obtain the  $\mathcal{O}(|I|^{3/4})$  static regret bound over the interval  $I$ .

Next, to track the best expert on the fly, we choose AdaNormalHedge (Luo and Schapire 2015) as the meta algorithm since it naturally supports the setting that the number of experts varies over time (Zhang, Liu, and Zhou 2019). The key ingredient of AdaNormalHedge is the potential function:  $\Phi(R, C) = \exp([R]_+^2 / 3C)$ , where  $[x]_+ = \max(0, x)$ ,  $\Phi(0, 0) = 1$  and  $R, C$  are two variables maintained by each expert. Based on  $\Phi(R, C)$ , we can compute the weight for each expert according to the following weight function:

$$w(R, C) = \frac{1}{2} (\Phi(R + 1, C + 1) - \Phi(R + 1, C - 1)).$$

Putting all pieces together, we obtain POLA for adaptive regret minimization. Below, we describe the detailed procedure, which is also summarized in Algorithm 3.

For brevity, we denote the set of all active experts as  $\mathcal{A}_t$  for the round  $t$ , and the set of intervals that start from the round  $t$  as  $\mathcal{C}_t = \{I \mid I \in \mathcal{I}, t \in I, (t - 1) \notin I\}$ . In Step 3, we create an instance of BOGD<sub>IP</sub> as the expert  $E_I$  for each  $I \in \mathcal{C}_t$ , and initiate it from an arbitrary initial point with the step size  $\eta = |I|^{-3/4}$ . In Step 4, we set the variables  $R_{t-1, I} = C_{t-1, I} = 0$  for  $E_I$ , where  $R_{t-1, I} = \sum_{u=\min I}^{t-1} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_{t, I})$  denotes the regret of  $E_I$  up to round  $t - 1$ , and  $C_{t-1, I} = \sum_{u=\min I}^{t-1} |f_t(\mathbf{x}_t) - f_t(\mathbf{x}_{t, I})|$  denotes the sum of the absolute value of instantaneous regrets, and  $\min I$  denotes the beginning round of  $I$ . In Step 5, the new expert  $E_I$  is added to  $\mathcal{A}_t$ . Then, we remove all experts from  $\mathcal{A}_t$ , who end at the round  $t$  (Step 7). After receiving the action  $\mathbf{x}_{t, I}$  from  $E_I$ , we update its corresponding weight as following:

$$w_{t, I} = \frac{w(R_{t-1, I}, C_{t-1, I})}{\sum_{E_I \in \mathcal{A}_t} w(R_{t-1, I}, C_{t-1, I})}. \tag{17}$$

In Step 9, we submit the weighted action

$$\mathbf{x}_t = \sum_{E_I \in \mathcal{A}_t} w_{t, I} \mathbf{x}_{t, I}, \tag{18}$$

and receive the loss function  $f_t(\cdot)$ . In Step 10, for each  $E_I \in \mathcal{A}_t$ , we compute its corresponding variables  $R_{t, I}$  and  $C_{t, I}$ . At the end, we reveal  $f_t(\cdot)$  to all active experts, so that they can update their actions for the next round (Step 11). We present the adaptive regret bound of POLA below.

**Theorem 3.** *Under Assumptions 1, 2, 3 and 4, Algorithm 3 guarantees*

$$\text{SA-Regret}_T(\tau) \leq \mathcal{O}(\sqrt{\tau \log T} + \tau^{3/4}) = \tilde{\mathcal{O}}(\tau^{3/4}).$$

**Remark:** Compared to existing methods (Garber and Kretzu 2022; Lu et al. 2023) for adaptive regret minimization, POLA has following advantages.

- POLA enjoys an  $\tilde{\mathcal{O}}(\tau^{3/4})$  strongly adaptive regret, and thus can still perform well on short intervals. In contrast, Garber and Kretzu (2022) minimize the weak adaptive regret (5), which only promises a performance guarantee on long intervals.
- For each expert, POLA performs only  $\mathcal{O}(1)$  linear optimizations per round on average, whereas Lu et al. (2023) require a significantly higher number of  $\mathcal{O}(\log T)$  membership operations. Moreover, their operations could be inefficient compared to linear optimizations in many popular domains. For example, the trace norm constraints  $\mathcal{K} = \{X \mid \|X\|_* \leq \delta, X \subset \mathbb{R}^{m \times n}\}$  incurs a membership operation cost of  $\mathcal{O}(mn^2)$  while the linear optimization cost is  $\mathcal{O}(nnz(X))$ , where  $nnz(X)$  denotes the number of non-zero entries (Mhammedi 2022).

Moreover, we note that previous studies on projection-based online learning (Zhang, Lu, and Yang 2020; Zhang et al. 2022; Cutkosky 2020) have shown that it is possible to design a single algorithm to minimize dynamic regret and adaptive regret simultaneously. In particular, our POLA shares a similar two-layer structure with the method of Zhang, Lu, and Yang (2020), inspiring us to investigate the performance of POLA for dynamic regret minimization. The following theorem shows that POLA also enjoys an  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  dynamic regret bound.

**Theorem 4.** *Under Assumptions 1, 2, 3 and 4, Algorithm 3 guarantees*

$$\text{D-Regret}_T(\mathbf{u}_1, \dots, \mathbf{u}_T) \leq \mathcal{O}\left(T^{3/4}(1 + P_T)^{1/4}\right).$$

**Remark:** Although POLA achieves the same dynamic regret bound as POLD, this does not imply that the latter one is insignificant. Compared with POLA, POLD employs a simpler meta algorithm and does not need to construct GC intervals, making it much easier to comprehend and implement.




Figure 2: Experimental results for dynamic regret minimization.

## Experiments

In this section, we present experimental results that verify our theoretical findings in dynamic regret. Empirical studies on adaptive regret can be found in the full version (Wang et al. 2023b).

**Setup.** To evaluate our methods (i.e. BOGD<sub>IP</sub>, POLD and POLA) in dynamic regret minimization, we study the problem of online matrix completion, of which the goal is to produce a matrix  $X$  from the trace norm ball in an online fashion to approximate the target matrix  $M \in \mathbb{R}^{m \times n}$ . Specifically, in each round  $t$ , the learner receive a sampled data  $(i, j)$  with the value  $M_{ij}$  from the entry set  $OB$  of  $M$ . Then, the learner chooses  $X$  from the trace norm ball  $\mathcal{K} = \{X \mid \|X\|_* \leq \delta, X \in \mathbb{R}^{m \times n}\}$  where  $\delta$  is the parameter, and suffers the online loss  $f_t(X) = \sum_{(i,j) \in OB} |X_{ij} - M_{ij}|$ . We conduct the experiments with  $\delta = 10^4$  on the public dataset: MovieLens 100K<sup>3</sup>, which contains 100000 ratings from 943 users on 1682 movies. Following Wan, Xue, and Zhang (2021), we slightly modify the dataset to simulate the non-stationary environments. Concretely, we generate an extended datasets  $\{(i_k, j_k, M_{i_k j_k})\}_{k=1}^{300000}$  by merging three copies of MovieLens 100K. For entries corresponding to  $k = 100001, \dots, 200000$ , we negate the original values  $M_{i_k j_k}$  to obtain  $-M_{i_k j_k}$ . For simplicity, we divide the extended datasets into  $T = 3000$  partitions. In this way, the target matrix  $M$  drifts every 1000 rounds.

**Contenders.** We choose the projection-free algorithm: Multi-OCG (Wan, Xue, and Zhang 2021), and the projection-based algorithm: Ader (Zhang et al. 2018) as the contenders. All parameters of each method are set according to the theoretical suggestions. For instance, the learning rate of the  $i$ -th expert is set as  $\eta_i = c(2^{i-1})^{-1/2}$  in Multi-OCG, and  $\eta_i = c2^{i-1}T^{-1/2}$  in Ader, and  $\eta_i = c2^{i-1}T^{-3/4}$  in POLD, where  $c$  is the hyper-parameters selected from  $\{10^{-1}, 10^0, \dots, 10^6\}$ .

**Results.** We report the average instantaneous loss, the cumulative loss and the runtime (in seconds) against the number of rounds for each method in Figure 2. As evident from

the results, projection-free methods are significantly more efficient compared to the projection-based approach (i.e. Ader), albeit with a slight compromise on cumulative loss. This observation is reasonable in the sense that (i) the cost of linear optimization over the trace norm ball is  $\mathcal{O}(nmz(X))$  whereas projection operation suffers a much higher  $\mathcal{O}(mn^2)$  cost; (ii) our methods ensure an  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  bound against the  $\mathcal{O}(\sqrt{T(1 + P_T)})$  bound of Ader. Moreover, owing to the inherent advantage in minimizing the general-case dynamic regret, our methods yield a lower cumulative loss compared to the projection-free contender Multi-OCG.

## Conclusion and Future Work

In this paper, we investigate non-stationary projection-free online learning with dynamic regret and adaptive regret guarantees. Specifically, in the dynamic regret minimization, we provide a novel dynamic regret analysis for BOGD<sub>IP</sub> (Garber and Kretzu 2022), and establish the first  $\mathcal{O}(T^{3/4}(1 + P_T))$  general-case dynamic regret. Then, we improve this bound to  $\mathcal{O}(T^{3/4}(1 + P_T)^{1/4})$  by proposing POLD, which runs a set of BOGD<sub>IP</sub> algorithms with different step sizes in parallel and tracks the best one on the fly. In the adaptive regret minimization, we present our method POLA with an  $\tilde{\mathcal{O}}(\tau^{3/4})$  strongly adaptive regret bound. The essential idea is to construct the GC intervals, maintain an instance of BOGD<sub>IP</sub> to minimize the static regret for each interval, and then combine actions of instances by a meta algorithm. Furthermore, we show that POLA can also minimize the dynamic regret and achieve the same bound as that of POLD. Empirical studies on dynamic regret and adaptive regret minimization have verified our theoretical findings.

Currently, both POLD and POLA need to maintain  $\mathcal{O}(\log T)$  experts, which leads to  $\mathcal{O}(\log T)$  linear optimizations per round. Therefore, a natural question arises: is it possible to further reduce the number of linear optimizations in each round, i.e., from  $\mathcal{O}(\log T)$  to  $\mathcal{O}(1)$ ? We note that in non-stationary projection-based online learning,  $\mathcal{O}(\log T)$  projection operations can indeed be reduced to  $\mathcal{O}(1)$  (Zhao et al. 2022). But in the projection-free setting, it seems highly non-trivial and we leave it as a future work.

<sup>3</sup><https://grouplens.org/datasets/movielens/100k/>

## Acknowledgments

This work was partially supported by NSFC (62361146852, 62306275), and the Collaborative Innovation Center of Novel Software Technology and Industrialization. We would like to thank the anonymous reviewers for their constructive suggestions.

## References

- Adamskiy, D.; Koolen, W. M.; Chernov, A.; and Vovk, V. 2016. A Closer Look at Adaptive Regret. *Journal of Machine Learning Research*, 17(23): 1–21.
- Baby, D.; and Wang, Y.-X. 2019. Online Forecasting of Total-variation-bounded Sequences. In *Advances in Neural Information Processing Systems* 32, 11071–11081.
- Baby, D.; and Wang, Y.-X. 2021. Optimal Dynamic Regret in Exp-Concave Online Learning. In *Proceedings of the 34th Conference on Learning Theory*, 359–409.
- Baby, D.; and Wang, Y.-X. 2022. 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*, 1805–1845.
- Besbes, O.; Gur, Y.; and Zeevi, A. J. 2015. Non-stationary Stochastic Optimization. *Operations Research*, 63(5): 1227–1244.
- Chen, L.; Zhang, M.; and Karbasi, A. 2019. Projection-free Bandit Convex Optimization. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, 2047–2056.
- Chiang, C.-K.; Yang, T.; Lee, C.-J.; Mahdavi, M.; Lu, C.-J.; Jin, R.; and Zhu, S. 2012. Online Optimization with Gradual Variations. In *Proceedings of the 25th Conference on Learning Theory*, 6.1–6.20.
- Cutkosky, A. 2020. Parameter-free, Dynamic, and Strongly-Adaptive Online Learning. In *Proceedings of the 37th International Conference on Machine Learning*, 2250–2259.
- Daniely, A.; Gonen, A.; and Shalev-Shwartz, S. 2015. Strongly Adaptive Online Learning. In *Proceedings of the 32nd International Conference on Machine Learning*, 1405–1411.
- Freund, Y.; and Schapire, R. E. 1997. A Decision-Theoretic Generalization of Online Learning and An Application to Boosting. *Journal of Computer and System Sciences*, 55(1): 119–139.
- Freund, Y.; Schapire, R. E.; Singer, Y.; and Warmuth, M. K. 1997. Using and Combining Predictors That Specialize. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, 334–343.
- Garber, D.; and Hazan, E. 2016. A Linearly Convergent Conditional Gradient Algorithm with Applications to Online and Stochastic Optimization. *SIAM Journal on Optimization*, 26(3): 1493–1528.
- Garber, D.; and Kretzu, B. 2020. Improved Regret Bounds for Projection-free Bandit Convex Optimization. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2196–2206.
- Garber, D.; and Kretzu, B. 2022. New Projection-free Algorithms for Online Convex Optimization with Adaptive Regret Guarantees. In *Proceedings of the 35th Conference on Learning Theory*, 2326–2359.
- Garber, D.; and Kretzu, B. 2023. Projection-free Online Exp-concave Optimization. In *Proceedings of the 36th Conference on Learning Theory*, 1259–1284.
- György, A.; Linder, T.; and Lugosi, G. 2012. Efficient Tracking of Large Classes of Experts. *IEEE Transactions on Information Theory*, 58(11): 6709–6725.
- Hazan, E. 2016. Introduction to Online Convex Optimization. *Foundations and Trends in Optimization*, 2(3–4): 157–325.
- Hazan, E.; and Kale, S. 2012. Projection-free Online Learning. In *Proceedings of the 29th International Conference on Machine Learning*, 1843–1850.
- Hazan, E.; and Minasyan, E. 2020. Faster Projection-free Online Learning. In *Proceedings of the 33rd Conference on Learning Theory*, 1877–1893.
- Hazan, E.; and Seshadhri, C. 2007. Adaptive Algorithms for Online Decision Problems. *Electronic Colloquium on Computational Complexity*, 14(8).
- Huang, R.; Lattimore, T.; György, A.; and Szepesvari, C. 2016. Following the Leader and Fast Rates in Linear Prediction: Curved Constraint Sets and Other Regularities. In *Advances in Neural Information Processing Systems* 29, 4970–4978.
- Jadbabaie, A.; Rakhlin, A.; Shahrampour, S.; and Sridharan, K. 2015. Online Optimization: Competing with Dynamic Comparators. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 398–406.
- Jun, K.-S.; Orabona, F.; Willett, R.; and Wright, S. 2017a. Improved Strongly Adaptive Online Learning Using Coin Betting. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 943–951.
- Jun, K.-S.; Orabona, F.; Wright, S.; and Willett, R. 2017b. Online Learning for Changing Environments Using Coin Betting. *Electronic Journal of Statistics*, 11(2): 5282–5310.
- Kalhan, D. S.; Bedi, A. S.; Koppel, A.; Rajawat, K.; Hassani, H.; Gupta, A. K.; and Banerjee, A. 2021. Dynamic Online Learning via Frank-Wolfe Algorithm. *IEEE Transactions on Signal Processing*, 69: 932–947.
- Kretzu, B.; and Garber, D. 2021. Revisiting Projection-free Online Learning: the Strongly Convex Case. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, 3592–3600.
- Levy, K.; and Krause, A. 2019. Projection Free Online Learning over Smooth Sets. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, 1458–1466.
- Littlestone, N.; and Warmuth, M. K. 1994. The Weighted Majority Algorithm. *Information and Computation*, 108(2): 212–261.
- Lu, Z.; Brukhim, N.; Gradu, P.; and Hazan, E. 2023. Projection-free Adaptive Regret with Membership Oracles.



- In *Proceedings of the 34th International Conference on Algorithmic Learning Theory*, 1055–1073.
- Luo, H.; and Schapire, R. E. 2015. Achieving All with No Parameters: Adanormalhedge. In *Proceedings of the 28th Conference on Learning Theory*, 1286–1304.
- Mhammedi, Z. 2022. Efficient Projection-Free Online Convex Optimization with Membership Oracle. In *Proceedings of the 35th Conference on Learning Theory*, 5314–5390.
- Mhammedi, Z.; Koolen, W. M.; and van Erven, T. 2019. Lipschitz Adaptivity with Multiple Learning Rates in Online Learning. In *Proceedings of the 32nd Conference on Learning Theory*, 2490–2511.
- Mokhtari, A.; Shahrampour, S.; Jadbabaie, A.; and Ribeiro, A. 2016. Online Optimization in Dynamic Environments: Improved Regret Rates for Strongly Convex Problems. In *55th IEEE Conference on Decision and Control*, 7195–7201.
- Molinaro, M. 2020. Curvature of Feasible Sets in Offline and Online Optimization. *ArXiv e-prints*, arXiv:2002.03213.
- Shalev-Shwartz, S. 2012. Online Learning and Online Convex Optimization. *Foundations and Trends in Machine Learning*, 4(2): 107–194.
- van Erven, T.; and Koolen, W. M. 2016. MetaGrad: Multiple Learning Rates in Online Learning. In *Advances in Neural Information Processing Systems 29*, 3666–3674.
- van Erven, T.; Koolen, W. M.; and van der Hoeven, D. 2021. MetaGrad: Adaptation Using Multiple Learning Rates in Online Learning. *Journal of Machine Learning Research*, 22(161): 1–61.
- Wan, Y.; Tu, W.-W.; and Zhang, L. 2020. Projection-free Distributed Online Convex Optimization with  $\mathcal{O}(\sqrt{T})$  Communication Complexity. In *Proceedings of the 37th International Conference on Machine Learning*, 9818–9828.
- Wan, Y.; Wang, G.; Tu, W.-W.; and Zhang, L. 2022. Projection-free Distributed Online Learning with Sublinear Communication Complexity. *Journal of Machine Learning Research*, 23(172): 1–53.
- Wan, Y.; Xue, B.; and Zhang, L. 2021. Projection-free Online Learning in Dynamic Environments. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, 10067–10075.
- Wan, Y.; and Zhang, L. 2021. Projection-free Online Learning over Strongly Convex Set. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, 10076–10084.
- Wan, Y.; Zhang, L.; and Song, M. 2023. Improved Dynamic Regret for Online Frank-Wolfe. In *Proceedings of the 36th Conference on Learning Theory*, 3304–3327.
- Wang, Y.; Wan, Y.; Zhang, S.; and Zhang, L. 2023a. Distributed Projection-Free Online Learning for Smooth and Convex Losses. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, 10226–10234.
- Wang, Y.; Yang, W.; Jiang, W.; Lu, S.; Wang, B.; Tang, H.; Wan, Y.; and Zhang, L. 2023b. Non-stationary Projection-free Online Learning with Dynamic and Adaptive Regret Guarantees. *ArXiv e-prints*, arXiv:2305.11726.
- Willems, F.; and Krom, M. 1997. Live-and-Die Coding for Binary Piecewise I.I.D. Sources. In *Proceedings of IEEE International Symposium on Information Theory*, 68.
- Yang, T.; Zhang, L.; Jin, R.; and Yi, J. 2016. Tracking Slowly Moving Clairvoyant: Optimal Dynamic Regret of Online Learning with True and Noisy Gradient. In *Proceedings of the 33rd International Conference on Machine Learning*, 449–457.
- Zhang, L.; Jiang, W.; Yi, J.; and Yang, T. 2022. Smoothed Online Convex Optimization Based on Discounted-Normal-Predictor. In *Advances in Neural Information Processing Systems 35*, 4928–4942.
- Zhang, L.; Liu, T.-Y.; and Zhou, Z.-H. 2019. Adaptive Regret of Convex and Smooth Functions. In *Proceedings of the 36th International Conference on Machine Learning*, 7414–7423.
- Zhang, L.; Lu, S.; and Yang, T. 2020. Minimizing Dynamic Regret and Adaptive Regret Simultaneously. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 309–319.
- Zhang, L.; Lu, S.; and Zhou, Z.-H. 2018. Adaptive Online Learning in Dynamic Environments. In *Advances in Neural Information Processing Systems 31*, 1323–1333.
- Zhang, L.; Yang, T.; Jin, R.; and Zhou, Z.-H. 2018. Dynamic Regret of Strongly Adaptive Methods. In *Proceedings of the 35th International Conference on Machine Learning*, 5877–5886.
- Zhao, P.; Xie, Y.-F.; Zhang, L.; and Zhou, Z.-H. 2022. Efficient Methods for Non-stationary Online Learning. In *Advances in Neural Information Processing Systems 35*, 11573–11585.
- Zhou, H.; Xu, Z.; and Tzoumas, V. 2023. Efficient Online Learning with Memory via Frank-Wolfe Optimization: Algorithms with Bounded Dynamic Regret and Applications to Control. *ArXiv e-prints*, arXiv:2301.00497.
- Zinkevich, M. 2003. Online Convex Programming and Generalized Infinitesimal Gradient Ascent. In *Proceedings of the 20th International Conference on Machine Learning*, 928–936.