

Revisiting Projection-Free Online Learning with Time-Varying Constraints

Yibo Wang^{1,2}, Yuanyu Wan^{3,1,*}, Lijun Zhang^{1,2,*}

¹ National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China

² School of Artificial Intelligence, Nanjing University, Nanjing 210023, China

³ School of Software Technology, Zhejiang University, Ningbo 315048, China
{wangyb, zhanglj}@lamda.nju.edu.cn, wanyy@zju.edu.cn

Abstract

We investigate constrained online convex optimization, in which decisions must belong to a fixed and typically complicated domain, and are required to approximately satisfy additional time-varying constraints over the long term. In this setting, the commonly used projection operations are often computationally expensive or even intractable. To avoid the time-consuming operation, several projection-free methods have been proposed with an $\mathcal{O}(T^{3/4}\sqrt{\log T})$ regret bound and an $\mathcal{O}(T^{7/8})$ cumulative constraint violation (CCV) bound for general convex losses. In this paper, we improve this result and further establish *novel* regret and CCV bounds when loss functions are strongly convex. The primary idea is to first construct a composite surrogate loss, involving the original loss and constraint functions, by utilizing the Lyapunov-based technique. Then, we propose a parameter-free variant of the classical projection-free method, namely online Frank-Wolfe (OFW), and run this new extension over the online-generated surrogate loss. Theoretically, for general convex losses, we achieve an $\mathcal{O}(T^{3/4})$ regret bound and an $\mathcal{O}(T^{3/4}\log T)$ CCV bound, both of which are order-wise tighter than existing results. For strongly convex losses, we establish new guarantees of an $\mathcal{O}(T^{2/3})$ regret bound and an $\mathcal{O}(T^{5/6})$ CCV bound. Moreover, we also extend our methods to a more challenging setting with bandit feedback, obtaining similar theoretical findings. Empirically, experiments on real-world datasets have demonstrated the effectiveness of our methods.

Introduction

Online convex optimization (OCO) has become a popular paradigm for modeling online decision-making problems (Shalev-Shwartz 2012; Hazan 2016; Orabona 2019), such as online portfolio optimization (Agarwal et al. 2006) and online advertisement system (McMahan et al. 2013). Formally, OCO can be viewed as a structured iterative game between a learner and an adversary. Specifically, at each round t , the learner first chooses a decision \mathbf{x}_t from a convex and fixed domain $\mathcal{K} \subseteq \mathbb{R}^d$. Then, the adversary reveals a convex loss function $f_t(\cdot) : \mathcal{K} \mapsto \mathbb{R}$, and the learner suffers the cost $f_t(\mathbf{x}_t)$. The goal of the learner is to minimize the regret:

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

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

defined as the difference between the cumulative loss of the learner and that of the best fixed decision.

In the literature, there have been abundant theoretical appeals for OCO, such as the $\mathcal{O}(\sqrt{T})$ regret bound for general convex losses (Zinkevich 2003) and the $\mathcal{O}(\log T)$ regret bound for strongly convex losses (Hazan, Agarwal, and Kale 2007). In practice, besides the *hard and fixed* domain \mathcal{K} , the decisions made by the learner are typically governed by a series of *soft and time-varying* constraints, which may be violated in several rounds but should be satisfied on average over the long term. For example, in wireless communication systems, operators manage varying transmission power consumption to ensure the reception of messages (Mannor, Tsitsiklis, and Yu 2009); in online advertisement systems, advertisers employ dynamic budgets to maximize the click-through-rates for their advertisements (Liakopoulos et al. 2019). These practical applications thus motivate the development of constrained online convex optimization (COCO) (Mahdavi, Jin, and Yang 2012; Neely and Yu 2017).

In the framework of COCO, the time-varying constraints are typically captured by the inequality $g_t(\mathbf{x}) \leq 0$ where $g_t(\cdot) : \mathcal{K} \mapsto \mathbb{R}$ is a convex function revealed by the adversary at the end of each round t . Consequently, in addition to minimizing (1), the learner also aims to ensure the cumulative constraint violation (CCV):

$$Q_T = \sum_{t=1}^T g_t^+(\mathbf{x}_t) \quad (2)$$

to be sublinear with respect to the time horizon T , where $g_t^+(\mathbf{x}) \triangleq \max\{0, g_t(\mathbf{x})\}$. To optimize (1) and (2) concurrently, various efforts have been made recently (Cao and Liu 2019; Yu and Neely 2020; Yi et al. 2021; Guo et al. 2022; Yi et al. 2023; Sinha and Vaze 2024), and established plentiful guarantees, including the regret and CCV bounds of $\mathcal{O}(\sqrt{T})$ for general convex losses (Yu, Neely, and Wei 2017).

The key operation in these COCO methods is the projection that pulls an infeasible decision back into the hard constraint \mathcal{K} . In many practical scenarios, the domain \mathcal{K} is typically high-dimensional and complex, rendering projections onto \mathcal{K} computationally expensive or even intractable, which significantly limits the applicability of these methods. To address this issue, several studies (Lee, Ho-Nguyen, and Lee 2023; Garber and Kretzu 2024) propose projection-free methods for COCO, which replace the time-consuming

projection with the more efficient linear optimization operation. One prominent example is online semidefinite optimization, where the hard constraint is a positive semidefinite cone with the bounded trace. In this case, the linear optimization has been proven at least an order of magnitude faster than projection (Hazan and Kale 2012). Unfortunately, existing projection-free methods can only guarantee the $\mathcal{O}(T^{3/4}\sqrt{\log T})$ regret bound and the $\mathcal{O}(T^{7/8})$ CCV bound for general convex losses (Garber and Kretzu 2024).

In this paper, we improve the above bounds and introduce *new* theoretical guarantees for strongly convex losses. The key idea is to first construct a composite surrogate loss that consists of the original loss $f_t(\cdot)$ and the time-varying constraint $g_t(\cdot)$, based on a carefully designed Lyapunov function. Rigorous analysis reveals that both (1) and (2) are simultaneously controlled by the regret in terms of the surrogate losses, so that we can directly apply classical projection-free methods, e.g., online Frank-Wolfe (OFW) (Hazan and Kale 2012), over the surrogate losses to minimize the two metrics. Notably, since the surrogate loss is generated in an online manner, essential prior knowledge for OFW (e.g., the gradient norm bound) is unavailable beforehand. Therefore, we need to employ the methods that are agnostic to the prior parameters about the surrogate loss. To this end, we propose the *first* parameter-free variant of OFW for general convex losses based on the doubling trick technique (Cesa-Bianchi et al. 1997). By running the parameter-free variant over the composite surrogate losses, we establish an $\mathcal{O}(T^{3/4})$ regret bound and an $\mathcal{O}(T^{3/4} \log T)$ CCV bound for the general convex loss. Both of our results are better than the state-of-the-art bounds achieved by Garber and Kretzu (2024). Additionally, we further investigate the strongly convex loss and achieve an $\mathcal{O}(T^{2/3})$ regret bound and an $\mathcal{O}(T^{5/6})$ CCV bound, by constructing the surrogate loss based on a different Lyapunov function and running the strongly convex variant of OFW (Wan and Zhang 2021).

Furthermore, to handle the more challenging bandit setting, we combine our proposed methods with the classical one-point estimator (Flaxman, Kalai, and McMahan 2005), which can approximate the gradient with only the loss value. Theoretically, for general convex losses, we establish the $\mathcal{O}(T^{3/4})$ regret bound and the $\mathcal{O}(T^{3/4} \log T)$ CCV bound. For strongly convex losses, we achieve the $\mathcal{O}(T^{2/3} \log T)$ regret bound and the $\mathcal{O}(T^{5/6} \log T)$ CCV bound.

Contributions. We summarize our contributions below.

- For general convex losses, we deliver an $\mathcal{O}(T^{3/4})$ regret bound and an $\mathcal{O}(T^{3/4} \log T)$ CCV bound, both of which improve the previous results of Garber and Kretzu (2024). During the analysis, we propose the *first* parameter-free variant of OFW, which may be an independent of interest;
- For strongly convex losses, we establish the *novel* results of an $\mathcal{O}(T^{2/3})$ regret bound and an $\mathcal{O}(T^{5/6})$ CCV bound for projection-free COCO;
- We extend our methods to the bandit setting and achieve similar bounds as those in the full-information setting;

- We verify our theoretical findings by conducting experiments on real-world datasets. The empirical results have demonstrated the effectiveness of our methods.

Related Work

In this section, we briefly overview the recent progress on projection-free and constrained online convex optimization.

Projection-Free Online Convex Optimization

The pioneering work of Hazan and Kale (2012) introduces the first projection-free online method, namely online Frank-Wolfe (OFW), which is an online extension of the classical Frank-Wolfe algorithm (Frank and Wolfe 1956). The basic idea is to replace the time-consuming projection with the following linear optimization steps:

$$\mathbf{v}_t = \underset{\mathbf{x} \in \mathcal{K}}{\operatorname{argmin}} \langle \nabla F_t(\mathbf{x}_t), \mathbf{x} \rangle, \quad \mathbf{x}_{t+1} = \mathbf{x}_t + \sigma_t(\mathbf{v}_t - \mathbf{x}_t)$$

where $\sigma_t > 0$ denotes the step size, and $F_t(\mathbf{x})$ is defined as

$$F_t(\mathbf{x}) = \eta \sum_{\tau=1}^{t-1} \nabla f_\tau(\mathbf{x}_\tau)^\top \mathbf{x} + \|\mathbf{x} - \mathbf{x}_1\|_2^2$$

parameterized by $\eta > 0$. With the prior knowledge about $f_t(\cdot)$ (e.g., the gradient norm bound) and appropriate configurations on η and σ_t , OFW ensures an $\mathcal{O}(T^{3/4})$ regret bound for general convex losses.

Based on OFW, plenty of investigations deliver tighter regret bounds by utilizing additional properties on $f_t(\cdot)$, such as the smoothness (Hazan and Minasyan 2020), the strong convexity (Wan and Zhang 2021; Kretzu and Garber 2021), and the exponential concavity (Garber and Kretzu 2023; Mhammedi 2024). Moreover, several efforts improve the regret bounds by leveraging special structures of \mathcal{K} (Garber and Hazan 2016; Levy and Krause 2019; Molinaro 2020; Wan and Zhang 2021; Mhammedi 2022; Gatmiry and Mhammedi 2023). Additionally, there exist other studies exploring more practical scenarios, e.g., the bandit feedback (Chen, Zhang, and Karbasi 2019; Garber and Kretzu 2020; Zhang et al. 2024), the delayed feedback (Wan et al. 2022b), the distributed setting (Zhang et al. 2017; Wan, Tu, and Zhang 2020; Wan et al. 2022a; Wang et al. 2023; Wan et al. 2024) and non-stationary environments (Kalhan et al. 2021; Wan, Xue, and Zhang 2021; Garber and Kretzu 2022; Lu et al. 2023; Wan, Zhang, and Song 2023; Wang et al. 2024).

Constrained Online Convex Optimization

In the literature, there are two lines of research for COCO. One is the time-invariant setting, where the soft constraints are assumed to be fixed, i.e., $g_t(\cdot) = g(\cdot)$, and known to the learner at the beginning round. In this setting, for general convex losses, Mahdavi, Jin, and Yang (2012) originally develop an $\mathcal{O}(\sqrt{T})$ regret bound and an $\mathcal{O}(T^{3/4})$ CCV bound. Then, subsequent studies generalize the results, and obtain tighter bounds for both regret and CCV under additional conditions (Jenatton, Huang, and Archambeau 2016; Yuan and Lamperski 2018; Yu and Neely 2020; Yi et al. 2021).

The other is the time-variant setting, where the soft constraints change over time and are only revealed after the

Methods	Losses	Constraints	Feedback	Regret	CCV
Lee, Ho-Nguyen, and Lee (2023)	cvx	sto	full-info	$\mathcal{O}(T^{5/6})$	$\mathcal{O}(T^{5/6})$
	cvx	adv	full-info	$\mathcal{O}(T^{5/6+\alpha})$	$\mathcal{O}(T^{11/12-\alpha/2})$
Garber and Kretzu (2024)	cvx	adv	full-info	$\mathcal{O}(T^{3/4}\sqrt{\log T})$	$\mathcal{O}(T^{7/8})$
Theorem 1 (this work)	cvx	adv	full-info	$\mathcal{O}(T^{3/4})$	$\mathcal{O}(T^{3/4} \log T)$
Theorem 2 (this work)	str-cvx	adv	full-info	$\mathcal{O}(T^{2/3})$	$\mathcal{O}(T^{5/6})$
Garber and Kretzu (2024)	cvx	adv	bandits	$\mathcal{O}(T^{3/4}\sqrt{\log T})$	$\mathcal{O}(T^{7/8} \log T)$
	cvx	adv	bandits	$\mathcal{O}(T^{3/4})$	$\mathcal{O}(T^{3/4} \log T)$
Theorem 3 (this work)	cvx	adv	bandits	$\mathcal{O}(T^{2/3} \log T)$	$\mathcal{O}(T^{5/6} \log T)$
Theorem 4 (this work)	str-cvx	adv	bandits	$\mathcal{O}(T^{2/3} \log T)$	$\mathcal{O}(T^{5/6} \log T)$

Table 1: Comparisons of our results with existing projection-free methods for COCO. Abbreviations: convex \rightarrow cvx, strongly convex \rightarrow str-cvx, stochastic \rightarrow sto, adversarial \rightarrow adv, full-information \rightarrow full-info.

learner submits the decision. Under the stochastic time-varying constraints and the Slater’s condition, Yu, Neely, and Wei (2017) deliver an $\mathcal{O}(\sqrt{T})$ regret bound and an $\mathcal{O}(\sqrt{T})$ CCV bound. Subsequently, extensive studies focus on the more general adversarial time-varying constraints and attempt to remove the Slater’s condition (Neely and Yu 2017; Sun, Dey, and Kapoor 2017; Liakopoulos et al. 2019; Cao and Liu 2019; Guo et al. 2022; Yi et al. 2023; Sinha and Vaze 2024). One of the key techniques in these work is to analyze a refined bound based on the Lyapunov drift of a virtual queue, which partially inspires our methods. To the best of our knowledge, the state-of-the-art results in this setting are delivered by Sinha and Vaze (2024), who establish the $\mathcal{O}(\sqrt{T})$ regret bound and the $\mathcal{O}(\sqrt{T} \log T)$ CCV bound for general convex losses, and the $\mathcal{O}(\log T)$ regret bound and the $\mathcal{O}(\sqrt{T} \log T)$ CCV bound for strongly convex losses.

As mentioned before, the above methods still rely on the inefficient projection for decision updates, which thereby motivates the development of projection-free COCO. Lee, Ho-Nguyen, and Lee (2023) first obtain an $\mathcal{O}(T^{5/6+\alpha})$ regret bound and an $\mathcal{O}(T^{11/12-\alpha/2})$ CCV bound with the parameter $\alpha \in (0, 1)$ for general convex losses and the full-information feedback. Later, Garber and Kretzu (2024) propose to apply a recent projection-free method, named LOO-BOGD (Garber and Kretzu 2022), under the drift-plus-penalty framework (Neely 2010) that is extensively used in previous COCO methods (Yu, Neely, and Wei 2017; Guo et al. 2022), and thus deliver an $\mathcal{O}(T^{3/4}\sqrt{\log T})$ regret bound and an $\mathcal{O}(T^{7/8})$ CCV bound. When only the bandit feedback (i.e., the function value) is accessible, they obtain the same regret bound and a slightly worse $\mathcal{O}(T^{7/8} \log T)$ CCV bound. More details can be found in Table 1.

Preliminaries

In this section, we recall the basic assumptions and definitions that are commonly used in prior studies (Mahdavi, Jin, and Yang 2012; Hazan and Kale 2012; Agrawal and Devanur 2014).

Assumption 1. *The convex decision set \mathcal{K} contains the ball of radius r centered at the origin $\mathbf{0}$, and is contained in an ball with the diameter $D = 2R$, i.e., $r\mathcal{B} \subseteq \mathcal{K} \subseteq R\mathcal{B}$ where*

$$\mathcal{B} = \{\mathbf{x} \in \mathbb{R}^d \mid \|\mathbf{x}\|_2 \leq 1\}.$$

Assumption 2. *At each round t , the loss function $f_t(\cdot)$ and the constraint function $g_t(\cdot)$ are G -Lipschitz over \mathcal{K} , i.e., $\forall \mathbf{x}, \mathbf{y} \in \mathcal{K}$, $|f_t(\mathbf{x}) - f_t(\mathbf{y})| \leq G\|\mathbf{x} - \mathbf{y}\|_2$ and $|g_t(\mathbf{x}) - g_t(\mathbf{y})| \leq G\|\mathbf{x} - \mathbf{y}\|_2$.*

Assumption 3. *At each round t , the loss function value $f_t(\mathbf{x})$ is bounded over \mathcal{K} , i.e., $\forall \mathbf{x} \in \mathcal{K}$, $|f_t(\mathbf{x})| \leq M$.*

Definition 1. *Let $\Phi(x) : \mathbb{R}^+ \mapsto \mathbb{R}$ be a convex function. It is called Lyapunov if $\Phi(x)$ satisfies (i) $\Phi(0) = 0$; (ii) $\Phi(x) \geq 0, \forall x \in \mathbb{R}^+$; (iii) $\Phi(x)$ is non-decreasing.*

Definition 2. *Let $f(\mathbf{x}) : \mathcal{K} \mapsto \mathbb{R}$ be a function over \mathcal{K} . It is called α_f -strongly convex if for all $\mathbf{x}, \mathbf{y} \in \mathcal{K}$*

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\alpha_f}{2} \|\mathbf{y} - \mathbf{x}\|_2^2.$$

In analysis, we will make use of the following property of strongly convex functions (Garber and Hazan 2015).

Lemma 1. *Let $f(\mathbf{x})$ be an α_f -strongly convex function over \mathcal{K} and $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}} f(\mathbf{x})$. Then, for any $\mathbf{x} \in \mathcal{K}$*

$$\frac{\alpha_f}{2} \|\mathbf{x} - \mathbf{x}^*\|_2^2 \leq f(\mathbf{x}) - f(\mathbf{x}^*). \quad (3)$$

Main Results

In this section, we initially present our methods as well as their theoretical guarantees for the full-information setting. Then, we extend our investigations to the bandit setting. Due to the limitation of space, all proofs are deferred in the supplementary material.

Algorithms for Full-Information Setting

Overall, we first construct a composite surrogate loss function based on the loss $f_t(\mathbf{x})$, the constraint $g_t(\mathbf{x})$ and a specially designed Lyapunov function that depends on the type of $f_t(\mathbf{x})$. Then, we employ parameter-free variants of OFW to optimize the surrogate loss.

Specifically, let Q_t be the cumulative constraint violation at the round t , and $\Phi(\cdot)$ be a convex Lyapunov function. According to (2), Q_t can be formalized recursively as

$$Q_t = Q_{t-1} + g_t^+(\mathbf{x}_t), \forall t \geq 1 \quad (4)$$

Algorithm 1: Online Frank-Wolfe with Time-Varying Constraints (OFW-TVC)

Input: Hyper-parameters β, γ , and the function $\Phi(\cdot)$

```

1: Choose any  $\mathbf{x}_1 \in \mathcal{K}$ , and set  $\tilde{G}_1 = s_1 = k = 1$ 
2: for  $t = 1$  to  $T$  do
3:   Play  $\mathbf{x}_t$ , and suffer  $f_t(\mathbf{x}_t)$  and  $g_t(\mathbf{x}_t)$ 
4:   Construct  $Q_t$  and  $\tilde{f}_t(\mathbf{x})$  according to (4) and (6)
5:   while  $\tilde{G}_k < \beta G(\gamma + \Phi'(\beta Q_t))$  do
6:     Set  $\tilde{G}_{k+1} = 2\tilde{G}_k, s_{k+1} = t, k = k + 1$ 
7:   end while
8:   Set  $\eta_k$  and  $F_{s_k:t}(\mathbf{x})$  according to (9) and (10)
9:   Compute  $\mathbf{v}_t$  and  $\sigma_{s_k,t}$  according to (11) and (12)
10:  Update  $\mathbf{x}_{t+1}$  according to (13)
11: end for

```

and $Q_0 = 0$. By utilizing the convexity of $\Phi(\cdot)$, the Lyapunov drift of Q_t at the round t , i.e., $\Phi(\beta Q_t) - \Phi(\beta Q_{t-1})$, is upper bounded by

$$\begin{aligned} \Phi(\beta Q_t) - \Phi(\beta Q_{t-1}) &\leq \Phi'(\beta Q_t) \beta [Q_t - Q_{t-1}] \\ &\stackrel{(4)}{=} \Phi'(\beta Q_t) \beta g_t^+(\mathbf{x}_t) \end{aligned} \quad (5)$$

where $\beta > 0$ denotes a hyper-parameter. To simultaneously minimize $f_t(\mathbf{x})$ and $g_t(\mathbf{x})$, we follow the drift-plus-penalty framework (Neely 2010), and construct the surrogate loss function $\tilde{f}_t(\mathbf{x})$ by combining the loss $f_t(\mathbf{x})$ and the upper bound of the Lyapunov drift in (5):

$$\tilde{f}_t(\mathbf{x}) = \gamma \beta f_t(\mathbf{x}) + \Phi'(\beta Q_t) \beta g_t^+(\mathbf{x}) \quad (6)$$

where $\gamma > 0$ denotes a hyper-parameter. In fact, it can be verified that the regret in terms of $\tilde{f}_t(\mathbf{x})$, denoted by Regret'_T , concurrently captures (1) and (2) (Sinha and Vaze 2024):

$$\text{Regret}'_T \stackrel{(5),(6)}{\geq} \gamma \beta \text{Regret}_T + \Phi(\beta Q_T). \quad (7)$$

With an appropriate configuration on $\Phi(\cdot)$, (1) and (2) can be decoupled from (7), delivering corresponding theoretical guarantees. It should be noticed that the specific choice of $\Phi(\cdot)$ is quite involved, since (i) it is employed to construct the surrogate loss in (6), necessitating a simple form that does not incur expensive computational costs; (ii) it appears in (7) and is required to adeptly balance the regret and CCV.

In the following, we investigate the general convex losses and the strongly convex losses.

General Convex Losses. Given the favorable property of minimizing $\tilde{f}_t(\mathbf{x})$ shown in (7), one may attempt to apply the classical OFW method over $\tilde{f}_t(\mathbf{x})$ for the simultaneous minimization on (1) and (2). However, such a straightforward application is not suitable, since $\tilde{f}_t(\mathbf{x})$ is generated in an online manner, and thus the prior knowledge required by OFW is unavailable beforehand. For example, the ℓ_2 -norm of the subgradient $\nabla \tilde{f}_t(\mathbf{x})$ is bounded by:

$$\begin{aligned} \|\nabla \tilde{f}_t(\mathbf{x})\|_2 &\leq \gamma \beta \|\nabla f_t(\mathbf{x})\|_2 + \Phi'(\beta Q_t) \beta \|\nabla g_t(\mathbf{x})\|_2 \\ &\leq \beta G(\gamma + \Phi'(\beta Q_T)) \triangleq \tilde{G}, \end{aligned} \quad (8)$$

in which the last step follows the fact that $\Phi(\cdot)$ is convex and hence its derivative $\Phi'(\cdot)$ is non-decreasing. From (8), it can be observed that \tilde{G} is unknown due to the uncertainty of Q_T at the round t . For this reason, we propose the first parameter-free variant of OFW, which is agnostic to \tilde{G} , and thereby can be employed to minimize $\tilde{f}_t(\mathbf{x})$. The basic idea is to utilize an estimation of \tilde{G} for decision updating. If the estimation is too low, we repeatedly double the current guess and employ the first valid value for updates. We summarize our method in Algorithm 1.

Specifically, at the Step 1, we choose any point $\mathbf{x}_1 \in \mathcal{K}$ as the decision for the first round and make the estimation $\tilde{G}_1 = 1$. Then, at each round t , we submit the decision \mathbf{x}_t , suffer the cost $f_t(\mathbf{x}_t)$ and the constraint $g_t(\mathbf{x}_t)$ (Step 3). At the Step 4, we construct Q_t and the surrogate loss function $\tilde{f}_t(\mathbf{x})$ according to (4) and (6), respectively. Next, we verify the feasibility of the estimation \tilde{G}_k . If it is lower than $\beta G(\gamma + \Phi'(\beta Q_t))$, we continuously double the current estimation until an appropriate value is found (Steps 5-7). After that, we set the learning rate

$$\eta_k = D(2\tilde{G}_k T^{3/4})^{-1}, \quad (9)$$

and construct the function

$$F_{s_k:t}(\mathbf{x}) = \eta_k \sum_{\tau=s_k}^t \langle \nabla \tilde{f}_\tau(\mathbf{x}_\tau), \mathbf{x} \rangle + \|\mathbf{x} - \mathbf{x}_{s_k}\|_2^2 \quad (10)$$

based on the historical gradients $\nabla \tilde{f}_t(\mathbf{x}_t)$ since the round s_k (Step 8), where s_k denotes the first round that utilizes the estimation \tilde{G}_k . At the Step 9, we compute \mathbf{v}_t according to

$$\mathbf{v}_t \in \underset{\mathbf{x} \in \mathcal{K}}{\text{argmin}} \langle \nabla F_{s_k:t}(\mathbf{x}_t), \mathbf{x} \rangle, \quad (11)$$

and set the step size as

$$\sigma_{s_k,t} = 2(t - s_k + 1)^{-1/2}. \quad (12)$$

Finally, we update the decision \mathbf{x}_{t+1} for the next round as shown below (Step 10):

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \sigma_{s_k,t} (\mathbf{v}_t - \mathbf{x}_t). \quad (13)$$

By choosing the exponential Lyapunov function $\Phi(x) = \exp(2^{-1} T^{-3/4} x) - 1$, we establish the following theorem.

Theorem 1. *Let $\beta = (2^6 G D)^{-1}$ and $\gamma = 1$. Under Assumptions 1 and 2, if the loss functions and the constraint functions are general convex, Algorithm 1 ensures the bounds of*

$$\text{Regret}_T = \mathcal{O}(T^{3/4}), \quad Q_T = \mathcal{O}(T^{3/4} \log T).$$

Remark. Compared to the $\mathcal{O}(T^{3/4} \sqrt{\log T})$ regret bound and the $\mathcal{O}(T^{7/8})$ CCV bound in Garber and Kretzu (2024), our results for both metrics are tighter. The underlying reasons can be attributed to: (i) the choice of projection-free methods. Under the drift-plus-penalty framework, Garber and Kretzu (2024) choose to run the projection-free LOO-BOGD method, which, due to its complex design, necessitates additional effort to balance the costs of linear optimization and the performance. In contrast, our proposed method

Algorithm 2: Strongly Convex Variant of OFW with Time-Varying Constraints (SCOFW-TVC)

Input: Hyper-parameters β, γ , and the function $\Phi(\cdot)$, and the modulus of strong convexity α_f

- 1: Choose any $\mathbf{x}_1 \in \mathcal{K}$
 - 2: **for** $t = 1$ to T **do**
 - 3: Play \mathbf{x}_t , and suffer $f_t(\mathbf{x}_t)$ and $g_t(\mathbf{x}_t)$
 - 4: Construct Q_t and $\tilde{f}_t(\mathbf{x})$ according to (4) and (6)
 - 5: Set $F_t^{sc}(\mathbf{x})$ according to (14)
 - 6: Compute \mathbf{v}_t and σ_t^{sc} according to (15) and (16)
 - 7: Update \mathbf{x}_{t+1} according to (17)
 - 8: **end for**
-

is inherently simpler, naturally requiring only one linear optimization per round; (ii) the specification of $\Phi(x)$. Garber and Kretzu (2024) implicitly choose $\Phi(x) = x$, which potentially fails to balance regret and CCV for general convex losses, leading to looser results. Furthermore, it should be emphasized that even if $\Phi(x)$ in Garber and Kretzu (2024) is replaced with the exponential function, the complex management of linear optimization costs in LOO-BOGD still prevents it from yielding the same results as ours.

Strongly Convex Losses. In this case, note that for the α_f -strong convex $f_t(\mathbf{x})$, the surrogate loss $\tilde{f}_t(\mathbf{x})$ defined in (6) is $\gamma\beta\alpha_f$ -strongly convex. Therefore, we can employ the strongly convex variant of OFW to minimize $\tilde{f}_t(\mathbf{x})$. In this paper, we choose the SCOFW method proposed by Wan and Zhang (2021), because of its simplicity and agnosticism to \tilde{G} . The detailed procedures are given in Algorithm 2.

Specifically, we first choose any decision $\mathbf{x}_1 \in \mathcal{K}$ for initialization (Step 1). Then, at each round t , we make the decision \mathbf{x}_t , suffer the loss $f_t(\mathbf{x}_t)$ and the constraint $g_t(\mathbf{x}_t)$, and construct Q_t and $\tilde{f}_t(\mathbf{x})$ according to (4) and (6) (Steps 3-4). At Steps 5-6, we construct $F_t^{sc}(\mathbf{x})$ in the following way:

$$F_t^{sc}(\mathbf{x}) = \sum_{\tau=1}^t \left[\left\langle \nabla \tilde{f}_\tau(\mathbf{x}_\tau), \mathbf{x} \right\rangle + C_1 \|\mathbf{x} - \mathbf{x}_\tau\|_2^2 \right] \quad (14)$$

where we denote $C_1 = \gamma\beta\alpha_f/2$ for brevity, and compute \mathbf{v}_t according to:

$$\mathbf{v}_t \in \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}} \langle \nabla F_t^{sc}(\mathbf{x}_t), \mathbf{x} \rangle \quad (15)$$

and σ_t^{sc} according to

$$\sigma_t^{sc} = \operatorname{argmin}_{\sigma \in [0,1]} F_t^{sc}(\mathbf{x}_t + \sigma(\mathbf{v}_t - \mathbf{x}_t)). \quad (16)$$

At the Step 7, we update the decision \mathbf{x}_{t+1} for the next round according to

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \sigma_{s_k,t}(\mathbf{v}_t - \mathbf{x}_t). \quad (17)$$

By choosing the quadratic Lyapunov function $\Phi(x) = x^2 + x$, we establish the following theoretical results.

Theorem 2. *Let $\beta = G^{-1}D^{-1}T^{-2/3}$ and $\gamma = G/(G + \alpha_f D)$. Under Assumptions 1 and 2, if the loss functions are α_f -strongly convex, and the constraint functions are general convex, Algorithm 2 ensures the bounds of*

$$\text{Regret}_T = \mathcal{O}(T^{2/3}), \quad Q_T = \mathcal{O}(T^{5/6}).$$

Remark. Theorem 2 provides the *first* theoretical guarantees for the strongly convex losses in projection-free COCO, which are tighter than those in Garber and Kretzu (2024) for general convex losses.

Algorithms for Bandit Setting

In this section, we investigate the bandit setting, where only the function value is available. To handle the more challenging setting, we introduce the one-point gradient estimator (Flaxman, Kalai, and McMahan 2005), which can approximate the gradient with a single function value.

One-Point Gradient Estimator. For a function $f(\mathbf{x})$, we define its δ -smooth version as

$$\hat{f}_\delta(\mathbf{x}) = \mathbb{E}_{\mathbf{u} \sim \mathcal{B}^d} [f(\mathbf{x} + \delta \mathbf{u})] \quad (18)$$

which satisfies the following lemma (Flaxman, Kalai, and McMahan 2005, Lemma 1).

Lemma 2. *Let $\delta > 0$, $\hat{f}_\delta(\mathbf{x})$ defined in (18) ensures*

$$\nabla \hat{f}_\delta(\mathbf{x}) = \mathbb{E}_{\mathbf{u} \sim \mathcal{S}^d} [(d/\delta) f(\mathbf{x} + \delta \mathbf{u}) \mathbf{u}] \quad (19)$$

where \mathcal{S}^d denotes the unit sphere in \mathbb{R}^d .

To exploit the one-point gradient estimator, we define the shrunk set of \mathcal{K} as stated below

$$\mathcal{K}_\delta = (1 - \delta/r)\mathcal{K} = \{(1 - \delta/r)\mathbf{x} \mid \mathbf{x} \in \mathcal{K}\},$$

where $0 < \delta < r$ denotes the shrunk parameter.

Compared to our methods for the full-information setting, we make the following modifications:

- At each round t , the decision \mathbf{x}_t consists of two parts:

$$\mathbf{x}_t = \mathbf{y}_t + \delta \mathbf{u}_t \quad (20)$$

where $\mathbf{y}_t \in \mathcal{K}_\delta$ denotes an auxiliary decision learned from historical information, and $\mathbf{u}_t \sim \mathcal{S}^d$ is uniformly sampled from \mathcal{S}^d ;

- The gradient of $\tilde{f}_t(\mathbf{x}_t)$ is approximated by the one-point gradient estimator:

$$\tilde{\nabla}_t = (d/\delta) [\tilde{f}_t(\mathbf{x}_t)] \mathbf{u}_t, \quad (21)$$

so that we can adhere to the update rules in our previous methods;

- To manage the approximate error introduced by (21), we employ the blocking technique (Garber and Kretzu 2020; Hazan and Minasyan 2020) for decision updates, i.e., dividing the time horizon T into equally-sized blocks and only updating decisions at the end of each block.

In the bandit setting, we also investigate the general convex losses and the strongly convex losses.

General Convex Losses. In this case, we incorporate the modifications (20) and (21), and the blocking technique into Algorithm 1. The detailed procedures are summarized in Algorithm 3. Specifically, for initialization, we set $\tilde{G}_1 = m = 1$, and choose any $\hat{\mathbf{y}}_1 \in \mathcal{K}_\delta$ (Step 1). At each round t , we update $\mathbf{y}_t = \hat{\mathbf{y}}_m$ where $\hat{\mathbf{y}}_m \in \mathcal{K}_\delta$ is the auxiliary decision used in the block m , make the decision \mathbf{x}_t according to (20),

Algorithm 3: Bandit Frank-Wolfe with Time-Varying Constraints (BFW-TVC)

Input: Hyper-parameters $\beta, \gamma, c, K, \epsilon$, and the function $\Phi(\cdot)$

```

1: Choose any  $\hat{\mathbf{y}}_1 \in \mathcal{K}_{\delta_1}$ , and set  $\tilde{G}_1 = m = 1$ 
2: for  $t = 1$  to  $T$  do
3:   Set  $\mathbf{y}_t = \hat{\mathbf{y}}_m$ , and play  $\mathbf{x}_t$  according to (20)
4:   Suffer  $f_t(\mathbf{x}_t)$  and  $g_t(\mathbf{x}_t)$ 
5:   Construct  $Q_t$  and  $\tilde{f}_t(\mathbf{x})$  according to (4) and (6)
6:   Compute  $\tilde{\nabla}_t$  according to (21)
7:   if  $t \bmod K = 0$  then
8:     while  $\tilde{G}_k < \beta G(\gamma + \Phi'(\beta Q_\tau)), \forall \tau$  in block do
9:       Set  $\tilde{G}_{k+1} = 2\tilde{G}_k, k = k + 1$ 
10:    end while
11:    Compute  $\hat{\nabla}_m = \sum_{\tau=t-K+1}^t \tilde{\nabla}_\tau$ 
12:    Set  $\eta_k$  and  $F_{b_k:m}(\mathbf{y})$  according to (22) and (23)
13:    Set  $\tilde{\mathbf{y}}_1 = \hat{\mathbf{y}}_m$  and  $\tau = 0$ 
14:    repeat
15:      Set  $\tau = \tau + 1$ 
16:      Update  $\mathbf{v}_\tau$  and  $\sigma_\tau$  according to (24) and (25)
17:      Compute  $\tilde{\mathbf{y}}_{\tau+1}$  according to (26)
18:    until  $\langle \nabla F_{b_k:m}(\tilde{\mathbf{y}}_\tau), \tilde{\mathbf{y}}_\tau - \mathbf{v}_\tau \rangle \leq \epsilon$ 
19:    Set  $\hat{\mathbf{y}}_{m+1} = \tilde{\mathbf{y}}_{\tau+1}$  and  $m = m + 1$ 
20:  end if
21: end for

```

and suffer $f_t(\mathbf{x}_t)$ and $g_t(\mathbf{x}_t)$ (Steps 3-4). Then, we construct the CCV Q_t , the surrogate loss function $\tilde{f}_t(\mathbf{x})$ and the gradient estimation $\tilde{\nabla}_t$ according to (4), (6) and (21), respectively (Steps 5-6). At the end of the block m , we update our decision. To be precise, we first evaluate the current guess \tilde{G}_k for the gradient norm bound: if it is unsuitable, we double the value until an appropriate \tilde{G}_k is found (Steps 8-10). At the Step 11, we compute the cumulative gradient estimation $\hat{\nabla}_m = \sum_{\tau=t-K+1}^t \tilde{\nabla}_\tau$ where K denotes the block size. At the Step 12, we set η_k according to

$$\eta_k = cD(dM\tilde{G}_k T^{3/4})^{-1}, \quad (22)$$

and construct $F_{b_k:m}$ according to

$$F_{b_k:m}(\mathbf{y}) = \eta_k \sum_{\tau=b_k}^m \langle \hat{\nabla}_\tau, \mathbf{y} \rangle + \|\mathbf{y} - \mathbf{y}_{s_k}\|_2^2, \quad (23)$$

where b_k denotes the first block that utilizes the estimation \tilde{G}_k , and s_k denotes the first round of b_k . Next, we update the auxiliary decision for the next block, and set $\tilde{\mathbf{y}}_1 = \mathbf{y}_m$ and $\tau = 0$ (Step 12). At Steps 14-18, we repeat the following procedures: updating $\tau = \tau + 1$, computing \mathbf{v}_τ according to

$$\mathbf{v}_\tau \in \operatorname{argmin}_{\mathbf{y} \in \mathcal{K}_\delta} \langle \nabla F_{b_k:m}(\tilde{\mathbf{y}}_\tau), \mathbf{y} \rangle, \quad (24)$$

and σ_τ according to

$$\sigma_\tau = \operatorname{argmin}_{\sigma \in [0,1]} F_{b_k:m}(\tilde{\mathbf{y}}_\tau + \sigma(\mathbf{v}_\tau - \tilde{\mathbf{y}}_\tau)), \quad (25)$$

and updating $\tilde{\mathbf{y}}_{\tau+1}$ according to

$$\tilde{\mathbf{y}}_{\tau+1} = \tilde{\mathbf{y}}_\tau + \sigma_\tau(\mathbf{v}_\tau - \tilde{\mathbf{y}}_\tau), \quad (26)$$

Algorithm 4: Strongly Convex Variant of BFW with Time-Varying Constraints (SCBFW-TVC)

Input: Hyper-parameters $\beta, \gamma, \delta, K, L$, and the function $\Phi(\cdot)$, and the modulus of strong convexity α_f

```

1: Choose any  $\hat{\mathbf{y}}_1 \in \mathcal{K}_\delta$ , and set  $m = 1$ 
2: for  $t = 1$  to  $T$  do
3:   Set  $\mathbf{y}_t = \hat{\mathbf{y}}_m$ , and play  $\mathbf{x}_t$  according to (20)
4:   Suffer  $f_t(\mathbf{x}_t)$  and  $g_t(\mathbf{x}_t)$ 
5:   Construct  $Q_t$  and  $\tilde{f}_t(\mathbf{x})$  according to (4) and (6)
6:   Compute  $\tilde{\nabla}_t$  according to (21)
7:   if  $t \bmod K = 0$  then
8:     Compute  $\hat{\nabla}_m = \sum_{\tau=t-K+1}^t \tilde{\nabla}_\tau$ 
9:     Set  $F_m^{sc}(\mathbf{y})$  according to (27), and  $\tilde{\mathbf{y}}_1 = \mathbf{y}_m$ 
10:    for  $\tau = 1$  to  $L$  do
11:      Compute  $\mathbf{v}_\tau^{sc}$  according to (28)
12:      Calculate  $\sigma_\tau^{sc}$  according to (29)
13:      Update  $\tilde{\mathbf{y}}_{\tau+1}$  according to (30)
14:    end for
15:    Set  $\hat{\mathbf{y}}_{m+1} = \tilde{\mathbf{y}}_{L+1}$  and  $m = m + 1$ 
16:  end if
17: end for

```

until the stop condition $\langle \nabla F_{b_k:m}(\tilde{\mathbf{y}}_\tau), \tilde{\mathbf{y}}_\tau - \mathbf{v}_\tau \rangle \leq \epsilon$ is satisfied. After that, we set the auxiliary decision $\hat{\mathbf{y}}_{m+1} = \tilde{\mathbf{y}}_{\tau+1}$ for the next block.

With the configurations of the exponential Lyapunov function $\Phi(x) = \exp(2^{-1}T^{-3/4}x) - 1$ and suitable parameters, we obtain the following theorem.

Theorem 3. Let $\gamma = 1, K = T^{1/2}, \epsilon = 4D^2T^{-1/2}$ and $c > 0$ be a constant satisfying $\delta = cT^{-1/4} \leq r$, and $\beta = C_2^{-1}$ where $C_2 = 2^4G(cD/r + 3c + 1 + 2cD/(dM) + dMD/c)$. Under Assumptions 1 and 2, if the loss functions and the constraint functions are general convex, Algorithm 3 ensures the bounds of

$$\mathbb{E}[\operatorname{Regret}_T] = \mathcal{O}(T^{3/4}), \quad Q_T = \mathcal{O}(T^{3/4} \log T).$$

Remark. Theorem 3 presents tighter regret and CCV bounds, compared to the $\mathcal{O}(T^{3/4} \sqrt{\log T})$ regret bound and the $\mathcal{O}(T^{7/8} \log T)$ CCV bound in Garber and Kretzu (2024).

Strongly Convex Losses. In this case, we also employ the one-point gradient estimator and the blocking technique, and summarize the procedures in Algorithm 4. Overall, our method for strongly convex losses is similar to Algorithm 3, with the primary difference in the update of decisions. Specifically, at the end of each block m , we first compute the cumulative gradient estimation $\hat{\nabla}_m = \sum_{\tau=t-K+1}^t \tilde{\nabla}_\tau$ (Step 8), and then construct $F_m^{sc}(\mathbf{y})$ as shown below:

$$F_m^{sc}(\mathbf{y}) = \sum_{\tau=1}^m \langle \hat{\nabla}_\tau, \mathbf{y} \rangle + C_3 \|\mathbf{y}\|_2^2 \quad (27)$$

where we denote $C_3 = \gamma\beta\alpha_f t/2$ for brevity, and set $\tilde{\mathbf{y}}_1 = \mathbf{y}_m$ (Step 9). Next, we repeat the following procedures for L times to refine the auxiliary decision (Steps 10-14): at the iteration $\tau \in [L]$, computing \mathbf{v}_τ^{sc} according to

$$\mathbf{v}_\tau^{sc} \in \operatorname{argmin}_{\mathbf{y} \in \mathcal{K}_\delta} \langle \nabla F_m^{sc}(\tilde{\mathbf{y}}_\tau), \mathbf{y} \rangle, \quad (28)$$

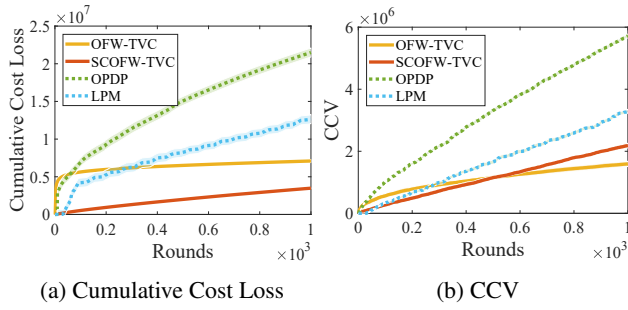


Figure 1: Experimental results on the MovieLens dataset.

calculating σ_t^{sc} according to

$$\sigma_\tau^{sc} = \underset{\sigma \in [0,1]}{\operatorname{argmin}} F_m^{sc}(\tilde{\mathbf{y}}_\tau + \sigma(\mathbf{v}_\tau^{sc} - \tilde{\mathbf{y}}_\tau)), \quad (29)$$

and updating $\tilde{\mathbf{y}}_{\tau+1}$ according to

$$\tilde{\mathbf{y}}_{\tau+1} = \tilde{\mathbf{y}}_\tau + \sigma_\tau^{sc}(\mathbf{v}_\tau^{sc} - \tilde{\mathbf{y}}_\tau). \quad (30)$$

Finally, we set the auxiliary decision for the next block as $\mathbf{y}_{m+1} = \tilde{\mathbf{y}}_{L+1}$ (Step 15).

By setting the quadratic Lyapunov function $\Phi(x) = x^2$ and proper parameters, we obtain the following theorem.

Theorem 4. Let $\beta = G^{-1}D^{-1}T^{-2/3}$, $\gamma = G/(G + \alpha_f D)$, $K = L = T^{2/3}$, and $\delta = cT^{-1/3}$ with $c > 0$ satisfying $cT^{-1/3} < r$, and $\gamma = \mathcal{O}(T^{2/3})$. Under Assumptions 1 and 2, if the loss functions are α_f -strongly convex, and the constraint functions are general convex, Algorithm 4 ensures the bounds of

$$\mathbb{E}[\operatorname{Regret}_T] = \mathcal{O}(T^{2/3} \log T), \quad Q_T = \mathcal{O}(T^{5/6} \log T).$$

Remark. Theorem 4 provides the *first* regret and CCV bounds for the strongly convex case with bandit feedback in projection-free COCO. By utilizing the strong convexity of $f_t(\cdot)$, both of our results are tighter than those established for the general convex losses in Garber and Kretzu (2024).

Experiments

In this section, we conduct empirical studies on real-world datasets to evaluate our theoretical findings.

General Setup. We investigate the online matrix completion problem (Hazan and Kale 2012; Lee, Ho-Nguyen, and Lee 2023), the goal of which is to generate a matrix X in an online manner to approximate the target matrix $M \in \mathbb{R}^{m \times n}$. Specifically, at each round t , the learner receives a sampled data (i, j) with the value M_{ij} from the observed subset O of M . Then, the learner chooses a matrix X from the trace norm ball $\mathcal{K} = \{X \mid \|X\|_* \leq \delta, X \in \mathbb{R}^{m \times n}\}$ where $\delta > 0$ is the parameter, and suffers the strongly convex cost loss $f_t(X_t) = \sum_{(i,j) \in O} (X_{ij} - M_{ij})^2/2$ and the constraint loss $g_t(X_t) = \operatorname{Tr}(P_t X_t)$ where P_t is uniformly sampled from $[-1, 1]^{n \times m}$. The experiments are conducted with $\delta = 10^4$ on two real-world datasets: MovieLens¹ for the full-information setting, and Film Trust (Guo, Zhang, and Yorke-Smith 2013) for the bandit setting.

¹<https://grouplens.org/datasets/movielens/100k/>

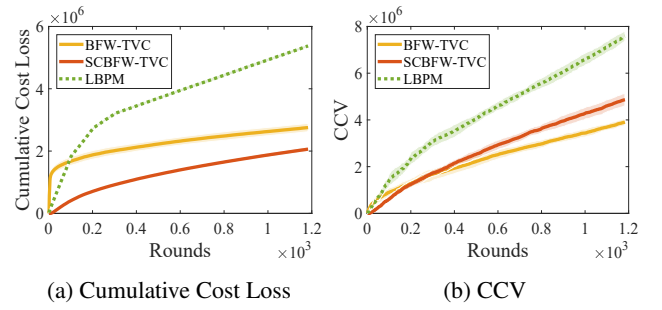


Figure 2: Experimental results on the Film Trust dataset.

Baselines. We choose three projection-free COCO methods as the contenders: (i) OPDP (Lee, Ho-Nguyen, and Lee 2023, Algorithm 1) and LPM (Garber and Kretzu 2024, Algorithm 4) for the full-information setting; (ii) LBPM (Garber and Kretzu 2024, Algorithm 5) for the bandit setting. All parameters of each method are set according to their theoretical suggestions, and we choose the best hyper-parameters from the range of $[10^{-5}, 10^{-4}, \dots, 10^4, 10^5]$.

Results. All experiments are repeated 10 times and we report experimental results (mean and standard deviation) in Figures 1 and 2. As evident from the results, in the full-information setting, OFW-TVC outperforms its competitors significantly in terms of both two metrics. Moreover, by utilizing the strong convexity of $f_t(\cdot)$, our SCOFW-TVC yields the lowest cumulative cost loss, albeit with a slight compromise on CCV. Similarly, in the bandit setting, it can be observed that our methods consistently outperform others, aligning with the theoretical guarantees.

Conclusion and Future Work

In this paper, we investigate projection-free COCO and propose a series of methods for the full-information and bandit settings. The key idea is to utilize the Lyapunov-based technique to construct a composite surrogate loss, consisting of the original cost and the constraint loss, and employ parameter-free variants of OFW running over the surrogate loss to simultaneously optimize the regret and CCV. In this way, we improve previous results for general convex cost losses and establish *novel* regret and CCV bounds for strongly convex cost losses. During the analysis, we propose the *first* parameter-free variant of OFW for general convex losses, which may hold independent interest. Finally, empirical studies have verified our theoretical findings.

Currently, for strongly convex losses, we improve the regret bound from $\mathcal{O}(T^{3/4})$ to $\mathcal{O}(T^{2/3})$, but sacrifice another metric CCV with a marginally looser bound of $\mathcal{O}(T^{5/6})$, compared to our results for general convex losses. This phenomenon may be due to the potential impropriety of the quadratic Lyapunov function. Hence, one possible solution is to choose other more powerful functions, which seems highly non-trivial, and we leave it as future work.

Acknowledgments

This work was partially supported by National Science and Technology Major Project (2022ZD0114801), and NSFC (U23A20382, 62306275). We would like to thank the anonymous reviewers for their constructive suggestions.

References

- Agarwal, A.; Hazan, E.; Kale, S.; and Schapire, R. E. 2006. Algorithms for Portfolio Management Based on the Newton Method. In *Proceedings of the 23rd International Conference on Machine Learning*, 9–16.
- Agrawal, S.; and Devanur, N. R. 2014. Fast Algorithms for Online Stochastic Convex Programming. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, 1405–1424.
- Cao, X.; and Liu, K. J. R. 2019. Online Convex Optimization with Time-Varying Constraints and Bandit Feedback. *IEEE Transactions on Automatic Control*, 64(7): 2665–2680.
- Cesa-Bianchi, N.; Freund, Y.; Haussler, D.; Helmbold, D. P.; Schapire, R. E.; and Warmuth, M. K. 1997. How to Use Expert Advice. *Journal of the ACM*, 44(3): 427–485.
- 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.
- Flaxman, A. D.; Kalai, A. T.; and McMahan, H. B. 2005. Online Convex Optimization in the Bandit Setting: Gradient Descent without a Gradient. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, 385–394.
- Frank, M.; and Wolfe, P. 1956. An Algorithm for Quadratic Programming. *Naval Research Logistics Quarterly*, 3(1–2): 95–110.
- Garber, D.; and Hazan, E. 2015. Faster Rates for the Frank-Wolfe Method over Strongly-Convex Sets. In *Proceedings of the 32nd International Conference on Machine Learning*, 541–549.
- 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.
- Garber, D.; and Kretzu, B. 2024. Projection-Free Online Convex Optimization with Time-Varying Constraints. In *Proceedings of the 41st International Conference on Machine Learning*, 14988–15005.
- Gatmiry, K.; and Mhammedi, Z. 2023. Projection-Free Online Convex Optimization via Efficient Newton Iterations. In *Advances in Neural Information Processing Systems 36*, 986–1008.
- Guo, G.; Zhang, J.; and Yorke-Smith, N. 2013. A Novel Bayesian Similarity Measure for Recommender Systems. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, 2619–2625.
- Guo, H.; Liu, X.; Wei, H.; and Ying, L. 2022. Online Convex Optimization with Hard Constraints: Towards the Best of Two Worlds and Beyond. In *Advances in Neural Information Processing Systems 35*, 36426–36439.
- Hazan, E. 2016. Introduction to Online Convex Optimization. *Foundations and Trends in Optimization*, 2(3–4): 157–325.
- Hazan, E.; Agarwal, A.; and Kale, S. 2007. Logarithmic Regret Algorithms for Online Convex Optimization. *Foundations and Trends in Machine Learning*, 69(2): 169–192.
- 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.
- Jenatton, R.; Huang, J.; and Archambeau, C. 2016. Adaptive Algorithms for Online Convex Optimization with Long-term Constraints. In *Proceedings of the 33rd International Conference on Machine Learning*, 402–411.
- 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.
- Lee, D.; Ho-Nguyen, N.; and Lee, D. 2023. Projection-Free Online Convex Optimization with Stochastic Constraints. *ArXiv e-prints*, arXiv:2305.01333.
- 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.
- Liakopoulos, N.; Destounis, A.; Paschos, G.; Spyropoulos, T.; and Mertikopoulos, P. 2019. Cautious Regret Minimization: Online Optimization with Long-Term Budget Constraints. In *Proceedings of the 36th International Conference on Machine Learning*, 3944–3952.
- 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.
- Mahdavi, M.; Jin, R.; and Yang, T. 2012. Trading Regret for Efficiency: Online Convex Optimization with Long Term

- Constraints. *Journal of Machine Learning Research*, 13(81): 2503–2528.
- Mannor, S.; Tsitsiklis, J. N.; and Yu, J. Y. 2009. Online Learning with Sample Path Constraints. *Journal of Machine Learning Research*, 10: 569–590.
- McMahan, H. B.; Holt, G.; Sculley, D.; Young, M.; Ebner, D.; Grady, J.; Nie, L.; Phillips, T.; Davydov, E.; Golovin, D.; Chikkerur, S.; Liu, D.; Wattenberg, M.; Hrafnkelsson, A. M.; Boulos, T.; and Kubica, J. 2013. Ad Click Prediction: A View from the Trenches. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1222–1230.
- 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. 2024. Online Convex Optimization with a Separation Oracle. *ArXiv e-prints*, arXiv:2410.02476.
- Molinaro, M. 2020. Curvature of Feasible Sets in Offline and Online Optimization. *ArXiv e-prints*, arXiv:2002.03213.
- Neely, M. J. 2010. Stochastic Network Optimization with Application to Communication and Queueing Systems. *Synthesis Lectures on Communication Networks*, 3(1): 1–211.
- Neely, M. J.; and Yu, H. 2017. Online Convex Optimization with Time-Varying Constraints. *ArXiv e-prints*, arXiv:1702.04783.
- Orabona, F. 2019. A Modern Introduction to Online Learning. *ArXiv e-prints*, arXiv:1912.13213v6.
- Shalev-Shwartz, S. 2012. Online Learning and Online Convex Optimization. *Foundations and Trends in Machine Learning*, 4(2): 107–194.
- Sinha, A.; and Vaze, R. 2024. Optimal Algorithms for Online Convex Optimization with Adversarial Constraints. In *Advances in Neural Information Processing Systems 37*.
- Sun, W.; Dey, D.; and Kapoor, A. 2017. Safety-Aware Algorithms for Adversarial Contextual Bandit. In *Proceedings of the 34th International Conference on Machine Learning*, 3280–3288.
- Wan, Y.; Tu, W.-W.; and Zhang, L. 2020. Projection-free Distributed Online Convex Optimization with $O(\sqrt{T})$ Communication Complexity. In *Proceedings of the 37th International Conference on Machine Learning*, 9818–9828.
- Wan, Y.; Wang, G.; Tu, W.-W.; and Zhang, L. 2022a. Projection-free Distributed Online Learning with Sublinear Communication Complexity. *Journal of Machine Learning Research*, 23(172): 1–53.
- Wan, Y.; Wang, Y.; Yao, C.; Tu, W.-W.; and Zhang, L. 2022b. Projection-free Online Learning with Arbitrary Delays. *ArXiv e-prints*, arXiv:2204.04964.
- Wan, Y.; Wei, T.; Xue, B.; Song, M.; and Zhang, L. 2024. Optimal and Efficient Algorithms for Decentralized Online Convex Optimization. *ArXiv e-prints*, arXiv:2402.09173.
- 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 Advances*, 10067–10075.
- Wan, Y.; and Zhang, L. 2021. Projection-free Online Learning over Strongly Convex Sets. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence Advances*, 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. 2023. 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. 2024. Non-stationary Projection-Free Online Learning with Dynamic and Adaptive Regret Guarantees. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 15671–15679.
- Yi, X.; Li, X.; Yang, T.; Xie, L.; Chai, T.; and Johansson, K. 2021. Regret and Cumulative Constraint Violation Analysis for Online Convex Optimization with Long Term Constraints. In *Proceedings of the 38th International Conference on Machine Learning*, 11998–12008.
- Yi, X.; Li, X.; Yang, T.; Xie, L.; Chai, T.; and Johansson, K. H. 2023. Regret and Cumulative Constraint Violation Analysis for Distributed Online Constrained Convex Optimization. *IEEE Transactions on Automatic Control*, 68(5): 2875–2890.
- Yu, H.; and Neely, M. J. 2020. A Low Complexity Algorithm with $O(\sqrt{T})$ Regret and $O(1)$ Constraint Violations for Online Convex Optimization with Long Term Constraints. *Journal of Machine Learning Research*, 21(1): 1–24.
- Yu, H.; Neely, M. J.; and Wei, X. 2017. Online Convex Optimization with Stochastic Constraints. In *Advances in Neural Information Processing Systems 30*, 1428–1437.
- Yuan, J.; and Lamperski, A. 2018. Online Convex Optimization for Cumulative Constraints. In *Advances in Neural Information Processing Systems 31*, 6140–6149.
- Zhang, C.; Wang, Y.; Tian, P.; Chen, X.; Wan, Y.; and Song, M. 2024. Projection-Free Bandit Convex Optimization over Strongly Convex Sets. In *Proceedings of the 28th Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 118–129.
- Zhang, W.; Zhao, P.; Zhu, W.; Hoi, S. C. H.; and Zhang, T. 2017. Projection-free Distributed Online Learning in Networks. In *Proceedings of the 34th International Conference on Machine Learning*, 4054–4062.
- Zinkevich, M. 2003. Online Convex Programming and Generalized Infinitesimal Gradient Ascent. In *Proceedings of the 20th International Conference on Machine Learning*, 928–936.