Online Nonsubmodular Optimization with Delayed Feedback in the Bandit Setting

Sifan Yang^{1,2}, Yuanyu Wan^{3,1,*}, Lijun Zhang^{1,3,*}

¹National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China ²School of Artificial Intelligence, Nanjing University, Nanjing, China ²State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, China {yangsf,zhanglj}@lamda.nju.edu.cn, wanyy@zju.edu.cn

Abstract

We investigate the online nonsubmodular optimization with delayed feedback in the bandit setting, where the loss function is α -weakly DR-submodular and β -weakly DRsupermodular. Previous work has established an (α, β) -regret bound of $\mathcal{O}(nd^{1/3}T^{2/3})$, where n is the dimensionality and d is the maximum delay. However, its regret bound relies on the maximum delay and is thus sensitive to irregular delays. Additionally, it couples the effects of delays and bandit feedback as its bound is the product of the delay term and the $\mathcal{O}(nT^{2/3})$ regret bound in the bandit setting without delayed feedback. In this paper, we develop two algorithms to address these limitations, respectively. Firstly, we propose a novel method, namely DBGD-NF, which employs the one-point gradient estimator and utilizes all the available estimated gradients in each round to update the decision. It achieves a better $\mathcal{O}(n\bar{d}^{1/3}T^{2/3})$ regret bound, which is relevant to the average delay $\bar{d} = \frac{1}{T} \sum_{t=1}^{T} d_t \leq d$. Secondly, we extend DBGD-NF by employing a blocking update mechanism to decouple the joint effect of the delays and bandit feedback, which enjoys an $\mathcal{O}(n(T^{2/3} + \sqrt{dT}))$ regret bound. When $d = \mathcal{O}(T^{1/3})$, our regret bound matches the $\mathcal{O}(nT^{2/3})$ bound in the bandit setting without delayed feedback. Compared to our first $\mathcal{O}(n\bar{d}^{1/3}T^{2/3})$ bound, it is more advantageous when the maximum delay $d = o(d^{2/3}T^{1/3})$. Finally, we conduct experiments on structured sparse learning to demonstrate the superiority of our methods.

Introduction

Online learning is a powerful framework that has been used to model various sequential prediction problems (Shalev-Shwartz 2012). It can address scenarios in which decisions are made from a small set (Hazan and Kale 2012), a continuous space (Hazan et al. 2016), or a combinatorial feasible domain (Cesa-Bianchi and Lugosi 2012). In this paper, we study the online nonsubmodular optimization (Lin et al. 2022), an emerging branch of the online learning, which appears in many machine learning tasks like structured sparse learning (El Halabi and Cevher 2015), Bayesian optimization (González et al. 2016), and column subset selection (Sviridenko, Vondrák, and Ward 2017), etc. Similar to the classical online convex optimization (OCO) (Zinkevich 2003), it is typically formulated as a game between an online player and an adversary. In each round $t \in [T]$, the player begins by selecting a subset decision $S_t \subseteq [n]$. When the player submits its decision S_t , the adversary chooses a nonsubmodular loss function $f_t(\cdot) : 2^{[n]} \mapsto \mathbb{R}$ and then the player suffers a loss $f_t(S_t)$, where $f_t(\cdot) = \overline{f_t}(\cdot) - \underline{f_t}(\cdot)$, $\overline{f_t}(\cdot)$ is α -weakly diminishing return submodular (DR-submodular), $\underline{f_t}(\cdot)$ is β -weakly diminishing return supermodular (DR-supermodular) and $2^{[n]}$ represents all the subset of [n].¹

The player aims to minimize the cumulative loss over T rounds, equivalently minimizing the regret:

$$\operatorname{Reg}(T) \triangleq \sum_{t=1}^{T} f_t(S_t) - \min_{S \subseteq [n]} \sum_{t=1}^{T} f_t(S), \qquad (1)$$

which compares the excess loss suffered by the player with that of the best decision chosen in hindsight. As pointed out by El Halabi and Jegelka (2020), the optimization problem $\min_{S\subseteq[n]}\sum_{t=1}^{T} f_t(S)$ is NP-hard, thus it is impossible to find the optimal decision in polynomial time. For this reason, we follow the previous work (Lin et al. 2022) and apply the (α, β) -regret to measure the performance of the online player, which is defined as

$$\operatorname{Reg}_{\alpha,\beta}(T) \triangleq \sum_{t=1}^{T} f_t(S_t) - \left(\frac{1}{\alpha} \bar{f}_t(S^{\star}) - \beta \underline{f}_t(S^{\star})\right), (2)$$

where $S^* = \arg\min_{S \subseteq [n]} \sum_{t=1}^{T} f_t(S)$ and (α, β) are the approximation factors achieved by a certain offline algorithm. Lin et al. (2022) are the first to investigate the online nonsubmodular optimization and develop a method that achieves an (α, β) -regret bound of $\mathcal{O}(\sqrt{nT})$, building on the Lovász extension (Lovász 1983) and the convex relaxation model (El Halabi and Jegelka 2020), where n is the dimensionality.

In lots of real-world scenarios, there may be a potential delay between the query of the player and the corresponding response (Quanrud and Khashabi 2015; Wan et al. 2022, 2023). To address the delayed scenario, Lin et al. (2022) also

^{*}Corresponding author

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

¹The definition of the α -weakly DR-submodular and β -weakly DR-supermodular can be found in Definition 2.

Setting	Method	(α,β) -regret bound
bandit	DBAGD (Lin et al. 2022)	$\mathcal{O}\left(nd^{1/3}T^{2/3} ight)$
bandit	BDGD-NF (Theorem 1)	$\mathcal{O}\left(n \bar{d}^{1/3} T^{2/3} ight)$
bandit	BDBGD-NF (Theorem 3)	$\mathcal{O}\left(n(T^{2/3}+\sqrt{dT})\right)$
full	DOAGD (Lin et al. 2022)	$\mathcal{O}\left(\sqrt{ndT} ight)$
full	DOGD-NF (Theorem 2)	$\mathcal{O}\left(\sqrt{n \bar{d} T}\right)$

Table 1: Summary of results for online nonsubmodular optimization under different settings, where n is the dimensionality, d is the maximum delay and $\bar{d} = \frac{1}{T} \sum_{t=1}^{T} d_t$ is the average delay over T rounds. For simplicity, we use the abbreviations: full \rightarrow full-information setting with delayed feedback, bandit \rightarrow bandit setting with delayed feedback.

explore the problem of online nonsubmodular optimization with delayed feedback, where the online player incurs an arbitrary delay $d_t \ge 1$ in receiving the response. To handle the delayed feedback, they utilize the pooling strategy (Héliou, Mertikopoulos, and Zhou 2020) to propose delay online approximate gradient descent (DOAGD), which enjoys an $\mathcal{O}(\sqrt{ndT})$ regret bound, where d is the maximum delay. Moreover, they consider a more challenging setting, bandit feedback, in which the online player does not receive any additional information about the loss function $f_t(\cdot)$ (e.g., its gradient) beyond the value $f_t(S_t)$, and develop delay online approximate gradient descent (DBAGD), achieving an $\mathcal{O}(nd^{1/3}T^{2/3})$ regret bound. However, their result for bandit setting with delayed feedback, summarized in Table 1. has two limitations. Firstly, it relies on the maximum delay d, which renders it sensitive to irregular delays. Secondly, it is the product of the delay term and the $\mathcal{O}(nT^{2/3})$ regret bound in the non-delayed bandit setting. This arises from DBAGD coupling the effects of the delays and bandit feedback, resulting in a discontented regret bound.

To overcome these limitations, we revisit the online nonsubmodular optimization with delayed feedback in the bandit setting. Specifically, we first develop a delayed algorithm to establish a regret bound that is relevant to the average delay. Our proposed method, named delayed bandit gradient descent for nonsubmodular function (DBGD-NF), achieves an enhanced $\mathcal{O}(n\bar{d}^{1/3}T^{2/3})$ regret bound, where $\bar{d} = \frac{1}{T} \sum_{t=1}^{T} d_t \leq d$ represents the average delay. The primary idea is to employ the one-point gradient estimator (Hazan and Kale 2012; Lin et al. 2022) and use all available estimated gradients to update the decision in each round, instead of utilizing the oldest one like DBAGD. Furthermore, as a by-product, if the online player has access to the gradient of the loss function, we can substitute the estimated gradient in DBGD-NF with the true gradient. Our algorithm, referred to as delayed online gradient descent for nonsubmodular function (DOGD-NF), enjoys a better $\mathcal{O}(\sqrt{ndT})$

regret bound for the full-information setting with delayed feedback.

In our pursuit of decoupling the joint effect of delayed feedback and the gradient estimator, we develop the blocking delayed bandit gradient descent for nonsubmodular function (BDBGD-NF). Drawing inspiration from Wan et al. (2024), we adopt the blocking update mechanism (Zhang et al. 2019; Garber and Kretzu 2020; Wang et al. 2023, 2024) with DBGD-NF. Particularly, we divide the total Trounds into several blocks of size K and update the decision at the end of each block using the estimated gradients from the blocks where all gradients are available. By setting a appropriate block size K, we can reduce the variance of the one-point gradient estimator. Leveraging this technique, BDBGD-NF achieves a superior $\mathcal{O}(n(T^{2/3} + \sqrt{dT}))$ regret bound. When the algorithm faces small delays, i.e., maximum delay $d = \mathcal{O}(T^{1/3})$, this regret bound matches the existing $\mathcal{O}(nT^{2/3})$ regret bound in the non-delayed setting (Lin et al. 2022), benefiting from the blocking update mechanism

On the other hand, when the impact of the delayed feedback d is substantial, i.e., $d = \Omega(T^{1/3})$, our regret bound is on the same order as the $\mathcal{O}(\sqrt{ndT})$ bound we establish for the full-information setting in terms of d and T under the worst case, where $\bar{d} = \Theta(d)$. Moreover, it is better than our first $\mathcal{O}(n\bar{d}^{1/3}T^{2/3})$ bound when the maximum delay $d = o(\bar{d}^{2/3}T^{1/3})$. Notably, compare to Wan et al. (2024), BDBGD-NF is specifically designed for online nonsubmodular optimization, making it more challenging to analyze. Finally, we compare our algorithms with the state-of-the-art (SOTA) methods through numerical experiments to demonstrate the robustness and effectiveness in handling delayed and bandit feedback effects.

To summarize, this paper makes the following contributions to online nonsubmodular optimization with delayed feedback:

- We propose two algorithms for online nonsubmodular optimization with delayed feedback to derive the regret bounds that are relevant to the average delay. Our methods reduce the regret bounds to $\mathcal{O}(\sqrt{ndT})$ and $\mathcal{O}(nd^{1/3}T^{2/3})$ for full-information and bandit feedback settings, respectively.
- To decouple the joint effect of the delayed and bandit feedback, we develop a novel algorithm by utilizing a blocking update technique, which enjoys an $\mathcal{O}(n(T^{2/3} + \sqrt{dT}))$ regret bound.

Related Work

In this section, we briefly introduce the related work of submodular optimization and nonsubmodular optimization.

Submodular Optimization

Submodular optimization has garnered increasing interest in various practical applications, such as sparse reconstruction (Bach 2010; Das, Dasgupta, and Kumar 2012; Liao et al. 2023), graph inference (Gomez-Rodriguez and Schölkopf

2012; Defazio and Caetano 2012), video analysis (Zheng et al. 2014), object detection (Song et al. 2014), etc. The primary property of a submodular function is diminishing returns, meaning that adding an element to a larger set provides less incremental gain than adding it to a smaller set. Early research (Badanidiyuru and Vondrák 2014; Krause and Golovin 2014) on submodular optimization mainly focuses on the offline setting, which may be unsuitable for sequential decision-making problems. To overcome this issue, Hazan and Kale (2012) investigate the online submodular optimization, where the player chooses a subset from a ground set of n elements in round t, and then observe a submodular loss function. Based on the Lovász extension (Lovász 1983), they extend online gradient descent (OGD) (Zinkevich 2003) and bandit gradient descent (BGD) (Flaxman, Kalai, and McMahan 2005) into submodular optimization in both full-information and bandit feedback settings, establishing $\mathcal{O}(\sqrt{nT})$ and $\mathcal{O}(nT^{2/3})$ regret bounds, respectively, where n is the dimensionality.

Nonsubmodular Optimization

Although submodularity is a natural assumption, the objective function is not always exactly submodular in various applications, such as structured sparse learning (El Halabi and Cevher 2015), Bayesian optimization (González et al. 2016), and column subset selection (Sviridenko, Vondrák, and Ward 2017), etc. Instead, it satisfies a weaker version of the DR property, like α -weakly DR-submodular and β weakly DR-supermodular (Lehmann, Lehmann, and Nisan 2006). El Halabi and Jegelka (2020) provide the first approximation guarantee for nonsubmodular minimization by developing an approximate projected subgradient method. Nevertheless, they only focus on the offline setting.

Lin et al. (2022) pioneer the study of online nonsubmodular optimization and introduce the (α, β) -regret to measure the performance of the online player, which is defined in (2). Based on the Lovász extension (Lovász 1983) and the convex relaxation model (El Halabi and Jegelka 2020), they propose online approximate gradient descent (OAGD), which obtains an $\mathcal{O}(\sqrt{nT})$ regret bound. Particularly, it uses the subgradient of the convex relaxation function to perform a gradient descent step over the Lovász extension domain. Then it samples the decision S_t from a certain distribution over all possible sets at round t. Moreover, they also consider the full-information and bandit settings with delay feedback, which are more complex to analyze. To handle the delay feedback, they further extend OAGD by adopting a pooling strategy (Héliou, Mertikopoulos, and Zhou 2020). Their method, namely DOAGD, which keeps a pool to store all the available delayed information and utilizes the oldest received but not utilized gradient to update the decision in each round, achieves an $\mathcal{O}(\sqrt{ndT})$ regret bound, where d is the maximum delay. Nevertheless, it depends on the maximum delay, and thus is sensitive to irregular delays. In the bandit setting, since the player can only observe the loss value of its decision S_t , they employ the importance weighting technique to estimate the gradient. In particular, S_t is chosen from a distribution that is related to the decision with probability $1 - \mu$ ($0 < \mu < 1$) and a random distribution with probability μ for exploration, ensuring the variance of the gradient is upper bounded by $\mathcal{O}(n^2/\mu)$. While they establish an $\mathcal{O}(nd^{1/3}T^{2/3})$ regret bound, their method couples the effects of the delays and gradient estimator.

Preliminary

In this section, we will provide essential definitions and basic setup for optimization of the nonsubmodular functions.

Definitions and Assumptions

Definition 1 For any function $f(\cdot) : 2^{[n]} \to \mathbb{R}$, we define $f(i | S) = f(\{i\} \bigcup S) - f(S)$ to denote the marginal gain of adding an element *i* to *S*. Moreover, $f(\cdot)$ is normalized if and only if $f(\emptyset) = 0$ and nondecreasing if and only if $f(A) \le f(B)$ for any $A \subseteq B$. $\Pi_{[0,1]^n}$ is the projection onto the domain $[0,1]^n$, which can be efficiently implemented as a simple clipping operation.

Definition 2 A function $f(\cdot) : 2^{[n]} \mapsto \mathbb{R}$ is α -weakly DR-submodular with $\alpha > 0$ if

$$f(i \mid A) \ge \alpha f(i \mid B), \quad \text{for all } A \subseteq B, i \in [n] \setminus B.$$

If this inequality holds when $\alpha = 1$, $f(\cdot)$ is submodular.

Similarly, $f(\cdot): 2^{[n]} \mapsto \mathbb{R}$ is β -weakly DR-supermodular with $\beta > 0$ if

$$f(i \mid B) \ge \beta f(i \mid A), \quad \text{for all } A \subseteq B, i \in [n] \setminus B.$$

 $f(\cdot)$ is supermodular when $\beta = 1$.

We define that $f(\cdot)$ is (α, β) -weakly DR-modular if both of the above inequalities hold simultaneously.

Building on the above definitions, we formulate the problem of minimizing structured nonsubmodular functions (El Halabi and Jegelka 2020; Lin et al. 2022):

$$f(S) := \overline{f}(S) - f(S), \tag{3}$$

where $S \subseteq [n]$. Afterwards, we introduce two common assumptions in the online nonsubmodular optimization.

Assumption 1 All the nonsubmodular functions $f(\cdot)$ defined in (3) satisfy $\overline{f}(S) + f(S) \leq L$ for all $S \subseteq [n]$.

Assumption 2 $f(\cdot)$ and $\underline{f}(\cdot)$ defined in (3) are normalized and non-decreasing. $\overline{f}(\cdot)$ is α -weakly DR-submodular and $f(\cdot)$ is β -weakly DR-supermodular.

Then we give an application of the online nonsubmodular optimization to provide practical insights of these assumptions, which is also used in our later experiments.

Structured sparse learning. This problem aims to learn a sparse parameter vector whose support satisfies a specific structure, such as group-sparsity, clustering, tree-structure, or diversity (Kyrillidis et al. 2015). It is typically formulated as $\min_{\mathbf{x} \in \mathbb{R}^n} \ell(\mathbf{x}) + \gamma F(\operatorname{supp}(\mathbf{x}))$, where $\ell(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}$ is the loss function, $F(\cdot) : 2^{[n]} \mapsto \mathbb{R}$ is a set function that imposes restrictions on the support set and γ is a tradeoff parameter. Previous approaches (Bach 2010; El Halabi and Cevher 2015) often replace the discrete regularizer $F(\operatorname{supp}(\mathbf{x}))$ with its closest convex relaxation, which is computationally tractable only when $F(\cdot)$ is submodular. El Halabi and Jegelka (2020) introduce an alternative formulation by using a nonsubmodular regularizer, which is better in practice, defined as

$$\min_{S \subseteq [n]} H(S) = \gamma F(S) - G(S), \tag{4}$$

where $G(S) = \ell(0) - \min_{supp(\mathbf{x}) \subseteq S} \ell(\mathbf{x})$ is a normalized non-decreasing set function. El Halabi and Jegelka (2020) have pointed out that when $\ell(\cdot)$ is smooth, strongly convex and is generated from random data, $G(\cdot)$ is weakly DRmodular. Moreover, if $F(\cdot)$ is α -weakly DR-submodular, the equation (4) can be transformed into (3) so that we can handle it directly. For example, $F(\cdot)$ is often chosen as the range cost function (Bach 2010) ($\alpha = \frac{1}{n-1}$ and n is the dimensionality), applied in the time-series and cancer diagnosis (Rapaport, Barillot, and Vert 2008), and the cost function ($\alpha = \frac{1+a}{1+b-a}$ and a, b are cost parameters), applied in the healthcare (Sakaue 2019).

Since nonsubmodular functions are defined over the discrete domain, determining their minimum values is a challenging task. Therefore, we introduce the Lovász extension (Lovász 1983) to transform a function $f(\cdot)$ defined over a discrete domain [n] to a new function $f_L(\cdot)$ over the unit hypercube $[0, 1]^n$. The extended function $f_L(\cdot)$ is convex if and only if $f(\cdot)$ is submodular. For nonsubmodular functions, we can exploit the convex relation (El Halabi and Jegelka 2020), enabling the use of convex optimization algorithms on the transformed function.

Lovász Extension and Convex Relaxation

Lovász extension ensures that the minima of the function over the domain $[0, 1]^n$ also recover the minima of the original function $f(\cdot)$. In this way, we can reduce the complex optimization task over domain [n] to a simpler convex optimization problem. To clarify this reduction process, we start with some necessary definitions.

Definition 3 A max chain of subsets of [n] is a collection of sets $\{A_0, ..., A_n\}$, and $\emptyset = A_0 \subseteq A_1 \subseteq ... \subseteq A_n = [n]$.

For any $\mathbf{x} \in [0, 1]^n = \mathcal{X}$, we introduce a unique associated permutation $\pi : [n] \mapsto [n]$ such that $\pi(i) = j$, meaning that \mathbf{x}_j is the *i*-th largest number in \mathbf{x} . Notably, we have $1 \ge \mathbf{x}_{\pi(1)} \ge ... \ge \mathbf{x}_{\pi(n)} \ge 0$ and let $\mathbf{x}_{\pi(0)} = 1, \mathbf{x}_{\pi(n+1)} = 0$ for simplicity. If we set $A_i = \{\pi(1), ..., \pi(i)\}$ for all $i \in [n]$ and $A_0 = \emptyset$, the vector \mathbf{x} can be expressed as a convex combination, i.e., $\mathbf{x} = \sum_{i=0}^n \lambda_i \chi(A_i)$, where $\lambda_i = \mathbf{x}_{\pi(i)} - \mathbf{x}_{\pi(i+1)}$ and $\sum_{i=0}^n \lambda_i = 1, \lambda_i \in [0, 1]$ (Hazan and Kale 2012). For any set $S \subseteq [n], \chi(\cdot) : 2^{[n]} \mapsto \{0, 1\}^n$ is an indicator function $\chi(S)_i = 1$ for all $i \in S$ and $\chi(S)_i = 0$ for all $i \notin S$. Next, we give the definition of the Lovász extension.

Definition 4 For any submodular function $f(\cdot)$, its Lovász extension $f_L(\cdot) : \mathcal{X} = [0,1]^n \mapsto \mathbb{R}$ is defined as $f_L(\mathbf{x}) = \sum_{i=0}^n (\mathbf{x}_{\pi(i)} - \mathbf{x}_{\pi(i+1)}) f(A_i) = \sum_{i=0}^n \mathbf{x}_{\pi(i)} f(\pi(i) \mid A_{i-1}).$

It is not hard to verify that $f_L(\chi(S)) = f(S)$ for any $S \subseteq [n]$. Therefore, minimizing the Lovász extension is equivalent to minimizing the original submodular function over all

possible sets. Moreover, Edmonds (1970) has pointed out that the subgradient g of $f_L(\mathbf{x})$ can be computed by

$$\mathbf{g}_{\pi(i)} = f\left(A_{i}\right) - f\left(A_{i-1}\right) \text{ for all } i \in [n].$$
(5)

However, when $f(\cdot)$ is not submodular, many properties break down. For example, $f_L(\cdot)$ is non-convex, which is harder to analyze. To tackle the nonsubmodular functions, we adopt the convex closure $f_C(\cdot)$, which is defined as:

Definition 5 The convex closure $f_C(\cdot) : [0,1]^n \mapsto \mathbb{R}$ for a nonsubmodular function $f(\cdot)$ is the point-wise largest convex function which always lower bounds $f(\cdot)$. Additionally, $f_C(\cdot)$ is the tightest convex extension of $f(\cdot)$ and $\min_{S \subseteq [n]} f(S) = \min_{x \in [0,1]^n} f_C(\mathbf{x})$.

Definition 5 gives us a simpler way to analyze the nonsubmodular function. Unfortunately, it is NP-hard to evaluate and optimize $f_C(\cdot)$ (Vondrák 2007). Nevertheless, we can utilize the proposition 3.1 in Lin et al. (2022) to derive the approximation:

Lemma 1 Assuming $f(\cdot)$ satisfies Assumption 2 and g is calculated according to (5) for all $A \subseteq [n]$ and $\mathbf{x} \in \mathcal{X}$, we have the following guarantees

$$f_{L}(\mathbf{x}) = \langle \mathbf{g}, \mathbf{x} \rangle \ge f_{C}(\mathbf{x}),$$
$$\sum_{i \in A} \mathbf{g}_{i} \le \frac{1}{\alpha} \bar{f}(A) - \beta \underline{f}(A),$$
$$f_{C}(\mathbf{x}) \le f_{L}(\mathbf{x}) = \langle \mathbf{g}, \mathbf{x} \rangle \le \frac{1}{\alpha} \bar{f}_{C}(\mathbf{x}) - \beta \underline{f}_{C}(\mathbf{x})$$

Remark 1 Lemma 1 demonstrates how Lovász extension $f_L(\mathbf{x})$ approximates the convex closure $f_C(\mathbf{x})$ so that the subgradient of $f_L(\mathbf{x})$ can serve as the approximate subgradient for $f_C(\mathbf{x})$ (El Halabi and Jegelka 2020), which plays an important role in our analysis.

Problem Setup

We consider the online nonsubmodular optimization with delayed feedback in the bandit setting, where the loss function $f_t(\cdot)$ is defined in (3), and satisfies Assumption 1 and Assumption 2. In each round t, the player makes a decision $S_t \subseteq [n]$ and then triggers a delay d_t when receiving the loss value. The response will arrive at round $t + d_t - 1$ and the player receives $\{f_k(S_k)|k \in \mathcal{F}_t\}$, where $\mathcal{F}_t = \{k \mid k + d_k - 1 = t\}$ represents the index set of received loss values in round t. To measure the performance of the online player, we follow the previous work (Lin et al. 2022) and apply the (α, β) -regret defined in (2). It compares the loss of the player's decisions to the result returned by an offline algorithm that approximately solves the optimization problem $\min_{S \subseteq [n]} \sum_{t=1}^{T} f_t(S)$ in polynomial time, which is different from the vanilla regret (Zinkevich 2003).

Main Results

In this section, we first develop a delayed algorithm for bandit setting to establish a regret bound that is relevant to the average delay. As a by-product, we demonstrate that it can be slightly adjusted into the full-information setting to obtain a better regret bound. Finally, we present our blocking method to decouple the joint effect of the delays and bandit feedback, further enhancing the regret bound.

Algorithm 1: DBGD-NF

Require: Learning rate η

1: Initialize $\mathbf{x}_1 \in [0,1]^n$

- 2: **for** t = 1 to T **do**
- 3: Let $1 \ge \mathbf{x}_{t,\pi_{(1)}} \ge \mathbf{x}_{t,\pi_{(2)}} \ge \dots \mathbf{x}_{t,\pi_{(n)}} \ge 0$ be the sorted entries in decreasing order with $A_{t,i} = \{\pi_{(1)}, \dots, \pi_{(i)}\}$ for all $i \in [n]$ and $A_{t,0} = \emptyset$. Define $\mathbf{x}_{t,\pi_{(0)}} = 1, \mathbf{x}_{t,\pi_{(n+1)}} = 0$
- 4: For $0 \le i \le n$, calculate $\lambda_{t,i} = \mathbf{x}_{t,\pi_{(i)}} \mathbf{x}_{t,\pi_{(i+1)}}$
- 5: Sample S_t from the distribution defined in (6)
- 6: Observe the loss $f_t(S_t)$
- 7: Calculate $\hat{f}_{t,i}$ according to (7)
- Compute the estimated gradient ĝ by (8) and incur a delay dt ≥ 1
- 9: Receive the gradient set $\{\hat{\mathbf{g}}_k | k \in \mathcal{F}_t\}$
- 10: Update \mathbf{x}_t according to (9)
- 11: end for

Results Related to the Average Delay

Existing literature (Lin et al. 2022) on online nonsubmodular optimization with delayed feedback in the bandit setting adopts a pooling strategy (Héliou, Mertikopoulos, and Zhou 2020) to handle the arbitrary delays, which only uses the oldest available information in a gradient pool. Since the gradient may be delayed by d rounds, its regret bound relies on the maximum delay d. To mitigate the effect of delayed feedback, motivated by previous work (Quanrud and Khashabi 2015), we utilize all the gradients received in round t to update the decision, rather than the oldest one. In the bandit setting, the online player only has access to the loss value. To deal with this issue, we employ the one-point estimator (Hazan and Kale 2012; Lin et al. 2022) to compute the unbiased gradient. Our method, DBGD-NF, is detailed in Algorithm 1. In each round t, S_t is sampled from the distribution

$$P(S_t = A_{t,i}) = (1 - \mu)\lambda_{t,i} + \frac{\mu}{n+1},$$
(6)

where $\lambda_{t,i} = \mathbf{x}_{t,\pi_{(i)}} - \mathbf{x}_{t,\pi_{(i+1)}}$ and $\mu \in (0,1)$ is the exploration probability. Then we utilize the one-point estimator to derive the gradient; that is, we compute

$$\hat{f}_{t,i} = \frac{\mathbf{1} \left(S_t = A_{t,i} \right)}{\left(1 - \mu \right) \lambda_{t,i} + \frac{\mu}{n+1}} f_t \left(S_t \right), \tag{7}$$

where $\mathbf{1}(\cdot)$ is an indicator function, and calculate the unbiased gradient

$$\hat{\mathbf{g}}_{t,\pi(i)} = \hat{f}_{t,i} - \hat{f}_{t,i-1}.$$
 (8)

Notably, we do not assume that the information is immediately available. Instead, owing to incurring a delay d_t , we only receive the information at the end of the round $t+d_t-1$, and we only present the calculation of the gradient for simplicity. Then we use all the available estimated gradients to update the decision

$$\mathbf{x}_{t+1} = \Pi_{[0,1]^n} \left[\mathbf{x}_t - \sum_{k \in \mathcal{F}_t} \hat{\mathbf{g}}_k \right]$$
(9)

Next, we establish the regret bound of Algorithm 1.

Algorithm 2: DOGD-NF

Require: Learning rate η

- 1: Initialize $\mathbf{x}_1 \in [0,1]^n$
- 2: **for** t = 1 to T **do**
- 3: Let $1 \ge \mathbf{x}_{t,\pi_{(1)}} \ge \mathbf{x}_{t,\pi_{(2)}} \ge \dots \mathbf{x}_{t,\pi_{(n)}} \ge 0$ be the sorted entries in decreasing order with $A_{t,i} = \{\pi_{(1)}, \dots, \pi_{(i)}\}$ for all $i \in [n]$ and $A_{t,0} = \emptyset$. Define $\mathbf{x}_{t,\pi_{(0)}} = 1, \mathbf{x}_{t,\pi_{(n+1)}} = 0$ 4: For $0 \le i \le n$, calculate $\lambda_{t,i} = \mathbf{x}_{t,\pi_{(i)}} - \mathbf{x}_{t,\pi_{(i+1)}}$
- 4: For $0 \le i \le n$, calculate $\lambda_{t,i} = \mathbf{x}_{t,\pi_{(i)}} \mathbf{x}_{t,\pi_{(i+1)}}$ 5: Sample S_t from the distribution defined in (10)
- 6: Observe the loss $f_t(S_t)$
- 7: Compute the estimated gradient \mathbf{g}_t by (11) and incur a delay $d_t \ge 1$
- 8: Receive the gradient set $\{\mathbf{g}_k | k \in \mathcal{F}_t\}$
- 9: Update \mathbf{x}_t according to (12)

10: end for

Theorem 1 Under Assumption 1 and Assumption 2, by setting $\mu = \frac{n\bar{d}^{1/3}}{T^{1/3}}, \eta = \frac{1}{L\bar{d}^{1/3}T^{2/3}}, DBGD-NF$ ensures

$$\mathbb{E}\left[\operatorname{\textit{Reg}}_{\alpha,\beta}(T)\right] \leq \mathcal{O}\left(n\bar{d}^{1/3}T^{2/3}\right)$$

where *n* is the dimensionality and $\bar{d} = \frac{1}{T} \sum_{t=1}^{T} d_t$ is the average delay.

Remark 2 Compared to the existing $O(nd^{1/3}T^{2/3})$ regret bound, DBGD-NF reduces the effect of delay and achieves a better $O(n\bar{d}^{1/3}T^{2/3})$ regret bound. It is worth noting that DBGD-NF requires the prior knowledge of the average delay to set the learning rate. Quanrud and Khashabi (2015) also encounter this issue and introduce a simple solution by utilizing the doubling trick (Cesa-Bianchi et al. 1997) to adaptively adjust the learning rate, thereby overcoming this limitation, which we can also employ to attain an equivalent $O(n\bar{d}^{1/3}T^{2/3})$ bound.

Additionally, we also observe that DBGD-NF can be slightly adjusted to the full-information setting to derive a regret bound that relies on the average delay. Our modified method, DOGD-NF, is summarized in Algorithm 2. In each round t, we sample S_t from the distribution

$$P(S_t = A_{t,i}) = \lambda_{t,i},\tag{10}$$

where $A_{t,i} = {\pi_{(1)}, \ldots, \pi_{(i)}}, A_{t,0} = \emptyset$ and $\lambda_{t,i} = \mathbf{x}_{t,\pi_{(i)}} - \mathbf{x}_{t,\pi_{(i+1)}}$. Similar to DOAGD (Lin et al. 2022), we also employ the convex relaxation based on the Lovász extension to compute the approximate subgradient

$$\mathbf{g}_{t,\pi(i)} = f_t(A_{t,i}) - f_t(A_{t,i-1}).$$
(11)

Finally, we use all the available gradients to perform a gradient descent step

$$\mathbf{x}_{t+1} = \Pi_{[0,1]^n} \left[\mathbf{x}_t - \sum_{k \in \mathcal{F}_t} \mathbf{g}_k \right].$$
(12)

Then we present the theoretical guarantee of Algorithm 2.

Theorem 2 Under Assumption 1 and Assumption 2, by setting $\eta = \frac{\sqrt{n}}{1 \sqrt{\frac{3}{2}T}}$, DOGD-NF ensures

$$\mathbb{E}\left[\operatorname{Reg}_{\alpha,\beta}(T)\right] \leq \mathcal{O}(\sqrt{ndT}).$$

To better showcase the improvement of our results, we provide an example that clearly demonstrates the enhancement through the better exponents of T.

Example 1 Consider a situation: $d_{1:T}$ satisfy $d_{1:\sqrt{T}} = \frac{T}{2}$ and $d_{\sqrt{T}+1:T} = 1$. Our methods achieve $\mathcal{O}(\sqrt{nT^{3/4}})$ and $\mathcal{O}(nT^{5/6})$ regret bounds, while the bounds of DOAGD and DBAGD are $\mathcal{O}(\sqrt{nT})$ and $\mathcal{O}(nT)$, respectively.

Decoupled Result

While BOGD-NF achieves a SOTA regret bound with respect to the average delay \overline{d} , its dependence on the maximum delay d remains suboptimal. In the following, we introduce a modification of BOGD-NF aimed at improving the regret dependence on the maximum delay. Inspired by Wan et al. (2024), we combine DBGD-NF with a blocking update technique (Zhang et al. 2019; Garber and Kretzu 2020; Wang et al. 2023, 2024) to reduce the effect of delays. The essential idea is dividing total T rounds into blocks with size K and choosing the decision S_t from the same mixture distribution per block, so that we can reduce the bound of the estimated gradients suffered within each block. It is not hard to verify that the bound on the sum of gradients within a block is

$$\mathbb{E}\left[\left\|\sum_{t=mK+1}^{(m+1)K} \hat{\mathbf{g}}_t\right\|\right] \leq \mathcal{O}\left(n\left(\sqrt{K/\mu}+K\right)\right).$$

By choosing a proper block size $K = \Theta(1/\mu)$, this bound can be enhanced to $\mathcal{O}(nK)$, which is smaller than the $\mathcal{O}(nK/\sqrt{\mu})$ bound of BDAGD. Based on this blocking update mechanism, we reduce the effect of delays on the regret from $\mathcal{O}(n\eta dT \sqrt{\mu})$ to $\mathcal{O}(n\eta dT)$, so that we decouple the effect of the delays and gradient estimator.

Our designed algorithm, BDBGD-NF, is presented in Algorithm 3. We first simply initialize x_1, y_1 to be any point in $[0,1]^n$ and available gradient set for the each block $\mathcal{P}_i = \emptyset$. In each block m, we sample the decision set S_t from the same mixture distribution (6) related to y_m and employ the one-point gradient estimator to derive the gradient.

At the end of each block m, we will identify the block where all the queries are obtained and use the all gradients from available blocks to update y_m

$$\mathbf{y}_{m+1} = \Pi_{[0,1]^n} \left[\mathbf{y}_m - \sum_{i \in \mathcal{A}_m} \sum_{\hat{\mathbf{g}}_k \in \mathcal{P}_i} \hat{\mathbf{g}}_k \right].$$
(13)

In the following, we present the following theorem to establish the theoretical guarantee for BDBGD-NF.

Theorem 3 Under Assumption 1 and Assumption 2, by setting $K = T^{\frac{1}{3}}, \mu = nT^{-1/3}, \eta = \min\{\frac{1}{LT^{2/3}}, \frac{1}{L\sqrt{dT}}\},$ **BDBGD-NF** ensures

$$\mathbb{E}\left[\operatorname{Reg}_{\alpha,\beta}(T)\right] \leq \mathcal{O}\left(n(T^{\frac{2}{3}} + \sqrt{dT})\right).$$

Algorithm 3: BDBGD-NF

Require: $\eta, \mu \in (0, 1), K$

- 1: Initialize point $\mathbf{x}_1 \in [0,1]^n$, set $\mathbf{y}_1 = \mathbf{x}_1$ and set the gradient pool for each block $\mathcal{P}_i = \emptyset, i = 1, ..., [T/K]$
- 2: for m = 1 to $\lfloor T/K \rfloor$ do
- 3: Set block pool $\mathcal{A}_m = \emptyset$
- 4: for time step t = (m-1)K + 1 to min $\{mK, T\}$ do
- 5:
- Choose $\mathbf{x}_t = \mathbf{y}_m$ Let $1 \ge \mathbf{x}_{t,\pi_{(1)}} \ge \mathbf{x}_{t,\pi_{(2)}} \ge \dots \mathbf{x}_{t,\pi_{(n)}} \ge 0$ be the sorted entries in decreasing order with $A_{t,i} =$ 6: $\{\pi_{(1)},\ldots,\pi_{(i)}\}$ for all $i \in [n]$ and $A_{t,0} = \emptyset$. Define $\mathbf{x}_{t,\pi_{(0)}} = 1$, $\mathbf{x}_{t,\pi_{(n+1)}} = 0$ For $0 \le i \le n$, calculate $\lambda_{t,i} = \mathbf{x}_{t,\pi_{(i)}} - \mathbf{x}_{t,\pi_{(i+1)}}$
- 7:
- Sample S_t from the distribution defined in (6) 8:
- 9: Observe the loss $f_t(S_t)$
- 10: Calculate $f_{t,i}$ according to (7)
- Compute the estimated gradient $\hat{\mathbf{g}}_t$ by (8) and incur 11: a delay $d_t > 1$
- Receive the set $\{\hat{\mathbf{g}}_k | k \in \mathcal{F}_t\}$ and update each gra-12: dient to its gradient pool $\mathcal{P}_j = \mathcal{P}_j \bigcup \{ \hat{\mathbf{g}}_k \}$, where $j = \lfloor k/K \rfloor$ is the block that $\hat{\mathbf{g}}_k$ belongs to
- end for 13:
- 14: If $|\mathcal{P}_i| = K$, $\mathcal{A}_m = \mathcal{A}_m \bigcup \{i\}$, which denotes the index of the block where all the queries arrive.
- 15: Perform gradient descent to y_{m+1} according to (13)
- For $i \in \mathcal{A}_m$, set $\mathcal{P}_i = \emptyset$ 16:

17: end for

Remark 3 When $d = O(T^{1/3})$, the regret bound of our method matches the previous $\mathcal{O}(nT^{2/3})$ regret bound (Lin et al. 2022) in the bandit setting without delayed feedback. Otherwise, this regret bound is also on the same order in terms of d and T with the $\mathcal{O}(\sqrt{ndT})$ regret in the fullinformation setting with delayed feedback in the worst case, where $\bar{d} = \Theta(d)$. Moreover, it is better than the former $\mathcal{O}(n\bar{d}^{1/3}T^{2/3})$ bound when $d = o(\bar{d}^{2/3}T^{1/3})$.

Experiments

In this section, we evaluate the effectiveness of our proposed methods through numerical experiments. We compare our methods with the SOTA methods, DOAGD and DBAGD (Lin et al. 2022) on structured sparse learning with delayed feedback. When it comes to hyper-parameter tuning, we set $\eta = \frac{\sqrt{n}}{L\sqrt{dT}}$ for DOGD-NF, $\eta = \frac{1}{Ld^{1/3}T^{2/3}}$ and $\mu = \frac{qnd^{1/3}}{T^{1/3}}$ for DBGD-NF and $\eta = \min\{\frac{1}{LT^{2/3}}, \frac{1}{L\sqrt{dT}}\},$ $\mu = \frac{qn}{T^{1/3}}$ and $K = T^{1/3}$ for BDBGD-NF. As for other methods, we choose $\eta = \frac{\sqrt{n}}{L\sqrt{dT}}$ for DOAGD, and $\eta = \frac{1}{Ld^{1/3}T^{2/3}}$ and $\mu = \frac{qnd^{1/3}}{T^{1/3}}$ for DBAGD, which are according to Theorem 5.2 and Theorem 5.4 in Lin et al. (2022). Moreover, we perform a grid search to select the parameter q from the set $\{0.01, 0.1, 1\}$. All the experiments are conducted in Python 3.7 with two 3.1 GHz Intel Xeon Gold 6346 CPUs and 32GB memory.

Setup. We conduct experiments on structured sparse



Figure 1: Results for full-information.



Figure 2: Results for bandit feedback.

learning defined in (4). Following the setup in Lin et al. (2022), we choose $F_t(S) = \max(S) - \min(S) + 1$ for all $S \neq \emptyset$ and $F(\emptyset) = 0$. We consider a simple linear regression problem, where the sparse optimal solution $\mathbf{x}^* \in \mathbb{R}^n$ only consists k consecutive 1, with the remaining positions being 0. We compute $\mathbf{y}_t = A_t \mathbf{x}^* + \epsilon_t$, where each row of $A_t \in \mathbb{R}^{s \times n}$ is a vector i.i.d. sampled from a Gaussian distribution and $\epsilon_t \in \mathbb{R}^s$ is a Gaussian noise vector with standard deviation 0.1. We choose the square loss $\ell_t(\mathbf{x}) = ||A_t\mathbf{x} - \mathbf{y}_t||^2/2$ and random delays in our experiments, i.e., $1 \leq d_t \leq d$, where d is the max delay. For all experiments, we set rounds T = 8000 such that the block K is 20, dimension n = 10, number of samples s = 128, trade-off parameter $\gamma = 0.1$ and sparse parameter k = 2.

Results. We report the regret against the number of rounds in full-information and bandit settings under different maximum delays d within the set {10, 20, 500} in Figure 1 and Figure 2, respectively. For the delays d_t , we sample them uniformly at random from the range [1, d]. We observe that DOGD-NF suffers less loss compared to DOAGD under full-information with delayed feedback because it utilizes all available gradients rather than just the oldest one. When the maximum delay increases, the superiority becomes more pronounced due to the large gap between the average delay \overline{d} and maximum delay d. As evident from Figure 2, our DBGD-NF also experiences less loss, which is consistent with our theories. Additionally, BDBGD-NF obtains the best performance and consistently yields lowest regret as the de-

lay changes in the bandit setting, which aligns with its theoretical guarantee.

Conclusion and Future Work

In this paper, we study the online nonsubmodular optimization with delayed feedback in the bandit setting and develop several algorithms to improve the existing regret bound. Firstly, our BDGD-NF and DOGD-NF achieve better $\mathcal{O}(n\bar{d}^{1/3}T^{2/3})$ and $\mathcal{O}(\sqrt{n\bar{d}T})$ regret bounds for bandit and full-information settings, respectively, which are relevant to the average delay. Furthermore, to decouple the joint effect of delays and bandit feedback, we combine BDGD-NF with a blocking update technique. Our BDBGD-NF obtains a superior $\mathcal{O}(n(T^{2/3} + \sqrt{dT}))$ regret bound. Finally, the experimental results also demonstrate the effectiveness of our methods.

One might notice that the regret bound of BDBGD-NF depends on the maximum delay rather than the average delay, unlike the former results. In the future, we will investigate how to develop a decoupled algorithm to obtain the regret bound that relies on the average delay. Moreover, we will try to deal with the online nonsubmodular optimization in the non-stationary environments.

Acknowledgments

This work was partially supported by NSFC (U23A20382, 62361146852), the Collaborative Innovation Center of

Novel Software Technology and Industrialization, and the Open Research Fund of the State Key Laboratory of Blockchain and Data Security, Zhejiang University. The authors would like to thank the anonymous reviewers for their constructive suggestions.

References

Bach, F. 2010. Structured sparsity-inducing norms through submodular functions. In *Advances in Neural Information Processing Systems 23*, 118–126.

Badanidiyuru, A.; and Vondrák, J. 2014. Fast algorithms for maximizing submodular functions. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1497–1514.

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.

Cesa-Bianchi, N.; and Lugosi, G. 2012. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5): 1404–1422.

Das, A.; Dasgupta, A.; and Kumar, R. 2012. Selecting diverse features via spectral regularization. In *Advances in Neural Information Processing Systems* 25, 3026–3034.

Defazio, A.; and Caetano, T. S. 2012. A convex formulation for learning scale-free networks via submodular relaxation. In *Advances in Neural Information Processing Systems* 25, 1250–1258.

Edmonds, J. 1970. Submodular functions, matroids, and certain polyhedra. In *Combinatorial structures and their applications*, 69–87.

El Halabi, M.; and Cevher, V. 2015. A totally unimodular view of structured sparsity. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 223–231.

El Halabi, M.; and Jegelka, S. 2020. Optimal approximation for unconstrained non-submodular minimization. In *Proceedings of the 37th International Conference on Machine Learning*, 3961–3972.

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.

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.

Gomez-Rodriguez, M.; and Schölkopf, B. 2012. Submodular inference of diffusion networks from multiple trees. In *Proceedings of the 29th International Conference on Machine Learning*, 489–496.

González, J.; Dai, Z.; Hennig, P.; and Lawrence, N. 2016. Batch bayesian optimization via local penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, 648–657. Hazan, E.; and Kale, S. 2012. Online submodular minimization. *Journal of Machine Learning Research*, 13(93): 2903– 2922.

Hazan, E.; et al. 2016. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4): 157–325.

Héliou, A.; Mertikopoulos, P.; and Zhou, Z. 2020. Gradientfree online learning in continuous games with delayed rewards. In *Proceedings of the 37th International Conference on Machine Learning*, 4172–4181.

Krause, A.; and Golovin, D. 2014. Submodular function maximization. *Tractability*, 3(71-104): 3.

Kyrillidis, A.; Baldassarre, L.; Halabi, M. E.; Tran-Dinh, Q.; and Cevher, V. 2015. Structured sparsity: Discrete and convex approaches. In *Compressed Sensing and its Applications: MATHEON Workshop 2013*, 341–387.

Lehmann, B.; Lehmann, D.; and Nisan, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2): 270–296.

Liao, Y.; Wan, Y.; Yao, C.; and Song, M. 2023. Improved Projection-free Online Continuous Submodular Maximization. *arXiv preprint arXiv:2305.18442*.

Lin, T.; Pacchiano, A.; Yu, Y.; and Jordan, M. 2022. Online nonsubmodular minimization with delayed costs: from full information to bandit feedback. In *Proceedings of the 39th International Conference on Machine Learning*, 13441–13467.

Lovász, L. 1983. Submodular functions and convexity. *Mathematical Programming The State of the Art: Bonn* 1982, 235–257.

Quanrud, K.; and Khashabi, D. 2015. Online learning with adversarial delays. In *Advances in Neural Information Processing Systems* 28, 1270–1278.

Rapaport, F.; Barillot, E.; and Vert, J.-P. 2008. Classification of arrayCGH data using fused SVM. *Bioinformatics*, 24(13): i375–i382.

Sakaue, S. 2019. Greedy and iht algorithms for non-convex optimization with monotone costs of non-zeros. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 206–215.

Shalev-Shwartz. 2012. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2): 107–194.

Song, H. O.; Girshick, R.; Jegelka, S.; Mairal, J.; Harchaoui, Z.; and Darrell, T. 2014. On learning to localize objects with minimal supervision. In *Proceedings of the 31th International Conference on Machine Learning*, 1611–1619.

Sviridenko, M.; Vondrák, J.; and Ward, J. 2017. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research*, 42(4): 1197–1218.

Vondrák, J. 2007. Submodularity in combinatorial optimization. Springer.

Wan, Y.; Wang, Y.; Yao, C.; Tu, W.-W.; and Zhang, L. 2022. Projection-free Online Learning with Arbitrary Delays. *arXiv preprint arXiv:2204.04964*.

Wan, Y.; Yao, C.; Song, M.; and Zhang, L. 2023. Nonstationary online convex optimization with arbitrary delays. *arXiv preprint arXiv:2305.12131*.

Wan, Y.; Yao, C.; Song, M.; and Zhang, L. 2024. Improved regret for bandit convex optimization with delayed feedback. *arXiv preprint arXiv:2402.09152*.

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.

Zhang, M.; Chen, L.; Hassani, H.; and Karbasi, A. 2019. Online continuous submodular maximization: From fullinformation to bandit feedback. In *Advances in Neural Information Processing Systems 32*, 9210–9221.

Zheng, J.; Jiang, Z.; Chellappa, R.; and Phillips, J. P. 2014. Submodular attribute selection for action recognition in video. In *Advances in Neural Information Processing Systems* 27, 481–489.

Zinkevich, M. 2003. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the* 20th International Conference on Machine Learning, 928– 936.