
Simultaneously Learning Stochastic and Adversarial Bandits with General Graph Feedback

Fang Kong¹ Yichi Zhou² Shuai Li¹

Abstract

The problem of online learning with graph feedback has been extensively studied in the literature due to its generality and potential to model various learning tasks. Existing works mainly study the adversarial and stochastic feedback separately. If the prior knowledge of the feedback mechanism is unavailable or wrong, such specially designed algorithms could suffer great loss. To avoid this problem, Erez & Koren (2021) try to optimize for both environments. However, they assume the feedback graphs are undirected and each vertex has a self-loop, which compromises the generality of the framework and may not be satisfied in applications. With a general feedback graph, the observation of an arm may not be available when this arm is pulled, which makes the exploration more expensive and the algorithms more challenging to perform optimally in both environments. In this work, we overcome this difficulty by a new trade-off mechanism with a carefully-designed proportion for exploration and exploitation. We prove the proposed algorithm simultaneously achieves $\text{poly } \log T$ regret in the stochastic setting and minimax-optimal regret of $\tilde{O}(T^{2/3})$ in the adversarial setting where T is the horizon and \tilde{O} hides parameters independent of T as well as logarithmic terms. To our knowledge, this is the first best-of-both-worlds result for general feedback graphs.

1. Introduction

The online learning problem can be formulated by a repeated game between the learner and an unknown environment

¹John Hopcroft Center for Computer Science, Shanghai Jiao Tong University, Shanghai, China ²Microsoft Research Asia, Beijing, China. Correspondence to: Shuai Li <shuaili8@sjtu.edu.cn>.

(Cesa-Bianchi & Lugosi, 2006; Lattimore & Szepesvári, 2020). The environment contains K arms. In each round, the learner selects one among these arms and observes some feedback. At the end of the round, it obtains the reward of the selected arm. The objective of the learner is to minimize the cumulative regret over a specified horizon, defined as the difference between the cumulative reward of the optimal arm and that of the arms selected.

To achieve this goal, the learner should adjust its strategy according to the feedback it observes. One of the most widely-studied feedback models is *bandit feedback*, where the learner only observes the reward of the selected arm in each round (Auer et al., 2002a; Lattimore & Szepesvári, 2020). The other extreme is called *full feedback*, where the learner is able to observe the reward of all arms no matter which arm it selects (Cesa-Bianchi & Lugosi, 2006; Eban et al., 2012). Both these two extremes are special cases of a general feedback model, where the feedback is characterized by a directed graph $G = ([K], E)$ (Mannor & Shamir, 2011). An edge $(i, j) \in E$ implies that the learner can observe arm j when selecting arm i . Under this framework, the graph of bandit feedback is composed of K disjoint self-loops and that of full feedback is a clique with self-loops. Benefiting from its generality and potential to model various learning tasks, graph feedback has been widely studied in the literature (Wu et al., 2015; Alon et al., 2015; Li et al., 2020; Chen et al., 2021).

Previous works on graph feedback mainly study two standard environments: stochastic (Buccapatnam et al., 2014; Wu et al., 2015; Tossou et al., 2017; Li et al., 2020) and adversarial (Mannor & Shamir, 2011; Alon et al., 2015; Chen et al., 2021). In the stochastic setting, the rewards of an arm are drawn independently from a fixed distribution. While in the adversarial setting, the rewards can be chosen arbitrarily. All of these works consider observable feedback graphs, i.e., each arm can be observed by selecting some arms. It is not hard to see that observability is necessary to guarantee sub-linear regret (Alon et al., 2015). With a general observable graph, the optimal regret in both settings has been well understood. In the stochastic setting, Wu et al. (2015) develop the instance-optimal regret of order $\Theta(\log T)$. And in the adversarial setting, Alon et al. (2015); Chen et al. (2021)

	Regret bound (stochastic)	Regret bound (adversarial)
Wu et al. (2015)	$O(D \log T / \Delta^2), \Omega(D \log T / \Delta^2)$	-
Alon et al. (2015)	-	$O\left(D \log K)^{\frac{1}{3}} T^{\frac{2}{3}}\right), \Omega\left(D / \log^2 K)^{\frac{1}{3}} T^{\frac{2}{3}}\right)$
Chen et al. (2021)	-	$O\left((\delta^* \log K)^{\frac{1}{3}} T^{\frac{2}{3}}\right), \Omega\left((\delta^* / \alpha)^{\frac{1}{3}} T^{\frac{2}{3}}\right)$
Ours	$O\left(D ^2 (\log T / \Delta^2)^{\frac{3}{2}}\right)$	$O\left(D ^{\frac{1}{3}} K^{\frac{2}{3}} T^{\frac{2}{3}} \sqrt{\log T}\right)$

Table 1. Comparisons of regret bounds with most related works in different environments for online learning with general graph feedback. T is the horizon, K is the number of arms, Δ is the minimum reward gap between the optimal and sub-optimal arms, D is the dominating set of the feedback graph, δ^* is the fractional weak domination number with $\delta^* \leq |D|$ and $\alpha = O(K/\delta^*)$ is the integrality gap of the linear program for vertex packing. The regret bounds in Wu et al. (2015) do not show explicit dependence on D and Δ . We compute the order of them in the bar graph which consists of $K/2$ disjoint undirected edges which connect K arms.

show the minimax-optimal regret is $\Theta(T^{2/3})$.

However, all these works treat different environment types separately. Once the prior knowledge of the reward type is unknown or wrong, the algorithms can suffer great loss. So a natural question is whether there is an algorithm that can simultaneously achieve poly $\log T$ regret in the stochastic setting and $\tilde{O}(T^{2/3})$ regret in the adversarial setting.

Similar best-of-both-worlds (BoBW) questions have been extensively studied under the bandit feedback (Bubeck & Slivkins, 2012; Seldin & Slivkins, 2014; Auer & Chiang, 2016; Seldin & Lugosi, 2017; Zimmert & Seldin, 2019) and a more general undirected feedback graph with self-loops (Erez & Koren, 2021). But as shown in the literature (Wu et al., 2015; Alon et al., 2015), even in a known single-type environment, learning with general feedback graphs (possibly without self-loops) is much harder. It is because the observations on an arm may not be available by simply selecting it, which makes the exploration of arms more expensive and the exploration-exploitation trade-off more challenging. This problem becomes especially difficult when trying to optimize for two different environments. Specifically, to achieve instance-optimality, the algorithm tends to first select arms with more observations to identify the optimal one. But these arms may have low rewards, which destroys the minimax-optimality over a specified horizon. So despite the importance of the above BoBW results under special feedback graphs, it is still open whether such type of results can be obtained with general graph feedback.

In this paper, we make progress on this problem. We develop a new mechanism to trade off exploration and exploitation with carefully-designed proportions for these two parts. Equipped with this mechanism and the carefully-designed testing for the environment type, the algorithm is proved to achieve poly $\log T$ regret in the stochastic setting and $\tilde{O}(T^{2/3})$ regret in the adversarial setting. Table 1 compares the regret bounds with most related works. It is shown that, compared with the state-of-the-art algorithms in both

settings, our algorithm only suffers from additional logarithmic and constant factors in the regret while without prior knowledge of the environment type.

2. Related Work

The study of online learning with bandit feedback has a long history (Thompson, 1933; Lai & Robbins, 1985). This problem has been well studied in both stochastic (Auer et al., 2002a) and adversarial setting (Auer et al., 2002b). To generalize the feedback type, Mannor & Shamir (2011) first introduce the framework of graph feedback with self-loops in the adversarial setting, which problem has been extensively studied by the following works (Kocák et al., 2014; Cohen et al., 2016). Alon et al. (2015) successfully remove the assumption of self-loops. They develop the Exp3.G algorithm and establish theoretical guarantees for graphs with different observabilities. To be specific, the optimal regret is $\Theta(T)$ for non-observable graphs, $\tilde{\Theta}(\sqrt{T})$ for strongly-observable graphs, and $\tilde{\Theta}(T^{2/3})$ for general observable graphs, where a strongly-observable graph requires each arm to have a self-loop or have in-edges from all of the other arms. The $\tilde{\Theta}(T^{2/3})$ regret also indicates the hardness to learn with a general feedback graph, which has been recently improved by Chen et al. (2021) in terms of the dependence on graph structure. In the stochastic setting, Wu et al. (2015) first consider general observable graphs and show the optimal regret is of order poly $\log T$. Li et al. (2020) later generalize this result by considering probabilistic graph feedback.

All of the above works separately study different environment types. A natural question is whether an algorithm can be derived to achieve the BoBW results. The BoBW problem was first studied by Bubeck & Slivkins (2012) in the bandit setting. They propose an algorithm and show it achieves not only poly $\log T$ regret in the stochastic environment but also $\tilde{O}(\sqrt{T})$ regret in the adversarial environment. Their algorithm is based on several consistency conditions to detect the optimality of the arms and the type of the envi-

ronment, which is later refined by Auer & Chiang (2016). A similar mechanism has also been adopted to obtain the BoBW result for linear bandit (Lee et al., 2021). The other line of BoBW works are variants of online mirror descent (OMD) (Zimmert & Seldin, 2019), which achieves better results than Bubeck & Slivkins (2012); Auer & Chiang (2016) in terms of logarithmic factors. Apart from bandit feedback, the BoBW problem has been recently studied with more general feedback graphs, though the graphs are assumed to be undirected and contain a self-loop on each arm (Erez & Koren, 2021). The algorithm in this work is a variant of OMD, which achieves $\text{poly} \log T$ regret in the stochastic setting and $O(\sqrt{T})$ regret in the adversarial setting. Recently, Lu et al. (2021) study a model between stochastic and adversarial, where the rewards are generated stochastically with a small fraction being corrupted adversarially. Their work is also only applicable to undirected feedback graphs with self-loops. To the best of our knowledge, we are the first to give the BoBW results for online learning with general graph feedback.

3. Setting

The problem can be characterized by a graph $G = (V, E)$, where $V = \{1, 2, \dots, K\}$ is the set of K arms and $E = \{(i, j)\}$ is the set of directed edges between arms. For each arm i , denote $N^{\text{in}}(i) = \{j \in V : (j, i) \in E\}$ as the set of in-neighbors and $N^{\text{out}}(i) = \{j \in V : (i, j) \in E\}$ as the set of out-neighbors of i . In each round $t = 1, 2, \dots$, the learner selects an arm $I_t \in V$ and the environment provides a reward vector $r_t := (r_t(1), r_t(2), \dots, r_t(K)) \in [0, 1]^K$. The learner is then rewarded with $r_t(I_t)$ and observes $r_t(j)$ for each arm $j \in N^{\text{out}}(I_t)$.

It is not hard to see that if some arm i is unobservable, i.e., $N^{\text{in}}(i) = \emptyset$, the learner could not determine which arm is optimal. Therefore, we consider a general observable graph, i.e., $N^{\text{in}}(i) \neq \emptyset$ for each arm $i \in V$. It is worth noting that in such a general graph, there may exist arm i with $(i, i) \notin E$. In this case, the learner though obtains the reward $r_t(i)$ by selecting arm i , has no chance to get the observation on it.

Stochastic and adversarial are two standard environments studied in the literature (Wu et al., 2015; Alon et al., 2015; Li et al., 2020; Chen et al., 2021). In the stochastic setting, the reward vector r_t at time t is independently drawn from a fixed reward distribution with expectation $\mathbb{E}[r_t] = \mu := (\mu_1, \mu_2, \dots, \mu_K)$. Denote $i^* = \text{argmax}_{i \in V} \mu_i$ as the optimal arm, which we assume is unique as previous works (Wu et al., 2015; Li et al., 2020; Erez & Koren, 2021). And for each arm $i \in V$, define $\Delta_i = \mu_{i^*} - \mu_i$ as the expected reward gap of i compared with i^* . In this setting, the pseudo

regret is defined as

$$\text{Reg}(T) = \sum_{t=1}^T \mu_{i^*} - \mu_{I_t}. \quad (1)$$

The instance-optimal expected regret attracts more interest in this setting, which is shown to be of order $\Theta(C(\mu) \log T)$ (Wu et al., 2015), where $C(\mu)$ is a constant related to the specific problem instance. The adversarial setting is more general than the stochastic setting, where the reward vector r_t can be chosen arbitrarily. Following previous works (Alon et al., 2015; Chen et al., 2021), we consider the oblivious adversary where the rewards are chosen at the start of the game. The pseudo regret in this setting is defined as

$$\text{Reg}(T) = \max_{i \in V} \sum_{t=1}^T r_t(i) - \sum_{t=1}^T r_t(I_t). \quad (2)$$

Researchers are interested in the minimax-optimal expected regret in this setting, which is of order $\Theta(T^{2/3})$ (Alon et al., 2015; Chen et al., 2021).

4. Algorithm

In this section, we present our BoBW algorithm for online learning with general graph feedback (Algorithm 1).

The algorithm takes the feedback graph G and a dominating set D of the graph as input (Line 1). A dominating set contains arms whose out-neighbors can cover the whole arm set, i.e., $\cup_{j \in D} N^{\text{out}}(j) = V$. Thus any arm can be observed by selecting arms in D .

At a high level, the algorithm starts with the assumption that the environment is stochastic and continuously monitors whether the assumption is satisfied. Typical optimal stochastic algorithms (Wu et al., 2015; Li et al., 2020) first explore arms in the dominating set to collect more observations until the optimal arm is identified, and then focuses on the estimated optimal one. But such type of algorithms would fail if the underlying environment is adversarial, since the regret caused by exploring arms cannot be controlled before the environment is detected to be adversarial.

As in the classical multi-armed bandits (MAB) problem, to minimize the cumulative regret, the learner needs to trade off exploitation and exploration. The former focuses on arms with high observed rewards so far to maintain high profits, while the latter aims to select unfamiliar arms to collect more observations (Auer et al., 2002a; Lattimore & Szepesvári, 2020). However, with general graph feedback, the observation of an arm may not be available by selecting it. Thus, to collect observations on an arm i , the learner needs to select some arms from $N^{\text{in}}(i)$. To this end, the algorithm maintains the active arm set A and the active dominating set D_A , which are initialized as V and

Algorithm 1 BoBW with General Graph Feedback

- 1: Input: graph $G = (V, E)$, dominating set D , hyper-parameter $\{\gamma_t\}_t$
- 2: Initialize: active arm set $A \leftarrow V; \forall i \in V, \tau_i \leftarrow \infty, \tau'_i \leftarrow \infty;$
active dominating set $D_A \leftarrow D; \forall i \in D, \tau_i^D \leftarrow \infty, u_i \leftarrow \frac{1}{|D|}$
- 3: **for** $t = 1, 2, \dots$ **do**
- 4: **for** $i \in V$ **do**
- 5: $p_{t,A}(i) \leftarrow \mathbb{1}\{i \in A\} \cdot \frac{1}{|A|}$
- 6: $p_{t,D}(i) \leftarrow \mathbb{1}\{i \in D_A\} \cdot \left(1 - \sum_{j \in D \setminus D_A} \frac{u_j \tau_j^D}{t}\right) \frac{1}{|D_A|} + \mathbb{1}\{i \in D \setminus D_A\} \cdot \frac{u_i \tau_i^D}{t}$
- 7: $p_t(i) \leftarrow (1 - \gamma_t) \cdot p_{t,A}(i) + \gamma_t \cdot p_{t,D}(i)$
- 8: **end for**
- 9: Select $I_t \sim p_t$ and be rewarded with $r_t(I_t)$; observe the reward $\{(i, r_t(i)), \forall i \in N^{\text{out}}(I_t)\}$
- 10: For each $i \in V$, compute

$$\tilde{H}_t(i) \leftarrow \frac{1}{t} \sum_{s=1}^t \frac{r_s(i) \cdot \mathbb{1}\{i \in N^{\text{out}}(I_s)\}}{\sum_{j \in N^{\text{in}}(i)} p_s(j)}. \quad (3)$$

- 11: **for** $i \in A$ such that

$$\tilde{H}_t(j') - \tilde{H}_t(i) > 3\text{Radius}_t(j') + 3\text{Radius}_t(i) \text{ for } j' \in \underset{j \in A}{\text{argmax}} \tilde{H}_t(j) \quad (4)$$

$$\text{where } \text{Radius}_t(j) = \sqrt{4 \left(\frac{|D|}{t^2} \sum_{s=1}^{\min\{t, \tau'_j\}} \frac{1}{\gamma_s} + \frac{|D| \cdot \max\{t - \tau'_j, 0\}}{\gamma_t \tau'_j t} \right) \log \frac{t}{\delta} + \frac{5|D|^2}{\gamma_t^2 \min\{t^2, \tau_j'^2\}} \log^2 \frac{t}{\delta}}, \forall j \quad (5)$$

do

- 12: $A \leftarrow A \setminus \{i\}, \tau_i \leftarrow t$
- 13: **end for**
- 14: **for** $j \in D_A$ such that $i \notin A, \forall i \in N^{\text{out}}(j)$ or $|A| = 1$ **do**
- 15: $D_A \leftarrow D_A \setminus \{j\}, \tau_j^D \leftarrow t, u_j \leftarrow p_{t,D}(j)$
- 16: **end for**
- 17: **for** $i \in V$ such that $\tau'_i > t$ and $\tau_i^D \leq t, \forall j \in D$ with $i \in N^{\text{out}}(j)$ **do**
- 18: $\tau'_i \leftarrow t$
- 19: **end for**
- 20: **if** $\exists i \notin A$ such that

$$\tilde{H}_t(j') - \tilde{H}_t(i) \leq \text{Radius}_t(i) + \text{Radius}_t(j') \text{ for } j' \in \underset{j \in A}{\text{argmax}} \tilde{H}_t(j) \quad (6)$$

then

- 21: Start Exp3.G and set $\tau \leftarrow t$
- 22: **end if**
- 23: **end for**

D , respectively (Line 2). Intuitively, A is the candidate set of arms with the potential to be optimal and D_A is a dominating set of A . The learner can then explore arms in D_A to get more observations and exploit arms in A to collect more rewards. How to balance these two parts is the key to achieving optimal regret in both stochastic and adversarial environments. The hyper-parameter $(\gamma_t)_t$ is taken as input

to control this trade-off (Line 1).

For the proceeding of the above exploration-exploitation trade-off, the algorithm maintains several indicators for arms. Specifically, for each arm $i \in V$, denote τ_i as the time when i is considered as sub-optimal and deleted from active set A . Then if the environment is truly stochastic, there is no need to exploit i anymore after τ_i . Similarly,

for each arm $j \in D$, let τ_j^D represent the time when all of j 's out-neighbors have been considered as sub-optimal and j is deleted from D_A . Based on this definition, the learner has no need to explore j for more observations on $N^{\text{out}}(j)$ after τ_j^D . Besides, for each arm $i \in V$, denote $\tau_i' = \operatorname{argmin}_{j \in D, j \in N^{\text{in}}(i)} \tau_j^D$ as the time when all of i 's dominating arms have been deleted from D_A . All of above values are initialized as ∞ at the beginning (Line 2).

To deal with the potential adversary, the learner needs to adopt a randomized strategy. Let p_t be the action distribution at t . The probability $p_t(i)$ of selecting i at t comes from two parts: the exploitation probability $p_{t,A}(i)$ (Line 5) and the exploration probability $p_{t,D}(i)$ (Line 6). Since arms in $V \setminus A$ are considered to be sub-optimal, the learner only needs to exploit arms in A . Here the algorithm conducts exploitation through a uniform distribution over active arm set A as Line 5. And for arms in the active dominating set $i \in D_A$, the learner also needs to explore i to get more observations on its out-neighbors, which is also implemented by a uniform distribution over D_A (Line 6). Besides, since the environment still has the potential to be adversarial and the deleted arms may become better, the observations on the deleted arms are still required for monitoring. For this reason, the algorithm would re-sample arms $j \in D \setminus D_A$ with a certain probability (Line 6). Note that the re-sampling probability should be small enough to avoid more regret since j may be sub-optimal arms. Here we follow [Bubeck & Slivkins \(2012\)](#) for the re-sampling schedule and continuously decrease the exploration probability for arms $j \in D \setminus D_A$ after τ_j^D as shown in Line 6, where u_j records the exploration probability of dominating arm j at time τ_j^D . The probability $p_t(i)$ is then defined as the weighted sum of these two parts and the weight is decided by the hyper-parameter γ_t (Line 7).

In each round t , the learner selects an arm $I_t \sim p_t$ and observes the reward $r_t(i)$ for each arm $i \in N^{\text{out}}(I_t)$ (Line 9). Based on the collected observations, it constantly detects the optimality of arms, which is implemented relying on the estimator \hat{H}_t defined as Eq.(3). The detection condition Eq.(4) is based on the concentration of the constructed estimators, where the concentration radius is shown in Eq.(5) and the detailed analysis is provided in Section 6. Specifically, if some arm i satisfies the detection condition in Eq.(4), then it is considered to be sub-optimal and would be eliminated from the active arm set (Line 12). The active dominating set D_A and the indicator τ_j' for each arm j would be updated accordingly (Line 14-19). It is worth noting that if the optimal arm is identified, i.e., $|A| = 1$, D_A should be emptied since there is no need to collect more observations on A (Line 14).

Recall that the learner also needs to re-sample arms $j \in D \setminus D_A$ to collect observations on deleted arms $i \notin A$ to

detect whether they become better with the influence of the potential adversary. This detection condition should be weak enough to guarantee to be satisfied in the stochastic setting and also strong enough such that before the condition is satisfied, the learner does not pay too much regret in the actual adversarial setting. Here we define Eq.(6) to check whether a previous sub-optimal arm i gets better. We show that in the stochastic setting, with high probability, this condition does not hold. And once it is satisfied, the learner believes the environment is actually adversarial and starts to run the optimal Exp3.G algorithm in the adversarial setting ([Alon et al., 2015](#)).

5. Theoretical Guarantees and Discussions

The regret upper bounds of Algorithm 1 in both stochastic and adversarial environments are presented in Theorem 5.1.

Theorem 5.1 (Main). *Let $\gamma_t = K^{\frac{2}{3}}|D|^{\frac{1}{3}}t^{-\frac{1}{3}}$. In the stochastic setting, Algorithm 1 guarantees that with probability at least $1 - \delta$, the regret is at most $O\left(|D|^2 \cdot \max_{i:i \neq i^*} \left(\frac{\log T/\delta}{\Delta_i^2}\right)^{\frac{3}{2}}\right)$. Moreover, in the adversarial setting, with probability at least $1 - \delta$, the expected regret is at most $O\left(K^{\frac{2}{3}}|D|^{\frac{1}{3}}T^{\frac{2}{3}}\sqrt{\log \frac{T}{\delta}}\right)$.*

5.1. Discussions

Theorem 5.1 is the first theoretical result for general graph feedback that simultaneously achieves near-optimal regret in both stochastic and adversarial environments. As shown in Table 1, our regret upper bound is only $O(|D|\sqrt{\log T}/\Delta)$ worse than the optimal regret in the stochastic setting ([Wu et al., 2015](#)) and $O(K^{2/3}\sqrt{\log T})$ worse than the optimal regret in the adversarial setting ([Alon et al., 2015](#); [Chen et al., 2021](#)). Above all, to simultaneously learn in both environments, our algorithm only suffers from additional logarithmic and constant factors in the regret.

Though our adopted framework belongs to the same line as [Bubeck & Slivkins \(2012\)](#), new techniques are needed to deal with the unique challenge brought by general graph feedback. With a general feedback graph, the observations and rewards of an arm cannot be simultaneously obtained. Thus an important modification of our algorithm is introducing the exploration on dominating arms and the challenge accompanying it is to trade off between selecting arms in dominating set and active set, which mechanism is controlled by the parameter γ_t . It is also worth noting that setting γ_t as previous works in one of the separate environments would cause the algorithm to fail in the other one. Specifically, if we directly follow the stochastic algorithm ([Wu et al., 2015](#)) and set $\gamma = 1$ before the optimal arm is identified and $\gamma = 0$ otherwise before starting Exp3.G, the regret in the adversarial setting would be $O(T)$ since the

regret before starting Exp3.G cannot be controlled. On the other hand, if we directly follow the adversarial algorithm (Alon et al., 2015) and set $\gamma = T^{-1/3}$, the regret in the stochastic setting would be $O(T^{1/3} \log T)$. Both these two results are not desirable.

Intuitively, to achieve the minimax-optimality in the adversarial setting, the learner needs to simultaneously conduct exploration on the dominating set and exploitation on the active set. However, if the environment is truly stochastic, the exploitation before an optimal arm is identified would produce additional regret. To bound this part of regret, we adjust the exploitation budget to let it polynomial depend on the exploration budget by setting γ_t to vary with time t . Thus when the optimal arm is identified after $O(\log T)$ exploration rounds, the additional cost of exploitation is also polynomial on $O(\log T)$. And based on the consideration to minimize the regret in the adversarial setting, the exact value of γ_t is set as $O(t^{-1/3})$. The additional $\sqrt{\log T}/\Delta$ regret in the stochastic setting compared with the lower bound is just the cost for preserving the minimax-optimality in a harder adversarial environment.

6. Proof of Theorem 5.1

Before the main proof, we first introduce some useful notations and inequalities. At any time t , let A_t be the active arm set A at the start of round t and denote $H_t(i) := \frac{1}{t} \sum_{s=1}^t r_s(i)$ as the averaged reward of i at t .

The proof of Theorem 5.1 highly depends on the concentration of the estimator $\tilde{H}_t(i)$. In the stochastic setting, it is an estimator of the unknown expected reward μ_i . And in the adversarial setting, it is the estimator of unknown averaged reward $H_t(i)$. In these two environments, we want to separately bound the difference between $\tilde{H}_t(i)$ and μ_i ($H_t(i)$, respectively).

Lemma 6.1. *For any arm $i \in V$ and time $t < \tau$, with probability at least $1 - \delta$, $|\tilde{H}_t(i) - \mu_i| \leq \text{Radius}_t(i)$ in the stochastic setting and $|\tilde{H}_t(i) - H_t(i)| \leq \text{Radius}_t(i)$ in the adversarial setting.*

Proof. Since the proofs of the two inequalities are similar, here we mainly focus on the adversarial setting.

In each round t , let \mathcal{F}_t be the history of observations. Define

$$X_t(i) = \tilde{r}_t(i) - r_t(i) = r_t(i) \left(\frac{\mathbb{1}\{i \in N^{\text{out}}(I_t)\}}{\sum_{j \in N^{\text{in}}(i)} p_s(j)} - 1 \right).$$

Note that $(X_s(i))_{1 \leq s \leq t}$ is a martingale difference sequences. According to the algorithm, at each time $t < \tau$,

$$\sum_{j \in N^{\text{in}}(i)} p_t(j) \geq \frac{\gamma_t}{|D|} \mathbb{1}\{t \leq \tau'_i\} + \frac{\gamma_t \tau'_i}{|D|t} \mathbb{1}\{t > \tau'_i\}.$$

Thus $|X_t(i)| \leq \max \left\{ \frac{|D|}{\gamma_t}, \frac{|D|t}{\gamma_t \tau'_i} \right\} \leq \frac{|D|}{\gamma_t} \max \left\{ 1, \frac{t}{\tau'_i} \right\}$ and

$$\begin{aligned} \sum_{s=1}^t \mathbb{E} [X_s(i)^2 | \mathcal{F}_{s-1}] &\leq \sum_{s=1}^{\min\{t, \tau'_i\}} \frac{|D|}{\gamma_s} + \sum_{s=\tau'_i}^t \frac{|D|s}{\gamma_s \tau'_i} \\ &\leq \sum_{s=1}^{\min\{t, \tau'_i\}} \frac{|D|}{\gamma_s} + \frac{t|D|}{\gamma_t \tau'_i} \cdot \max\{t - \tau'_i, 0\}, \end{aligned}$$

where the last inequality holds since γ_t decreases with the increase of t by choosing $\gamma_t = (K^2|D|)^{1/3} t^{-1/3}$. According to Lemma 6.7 and the definition of $\text{Radius}_t(i)$, we have with probability at least $1 - \delta$

$$\sum_{s=1}^t X_s(i) \leq t \cdot \text{Radius}_t(i),$$

and further

$$\left| \tilde{H}_t(i) - H_t(i) \right| \leq \text{Radius}_t(i),$$

since $\tilde{H}_t(i) - H_t(i) = \frac{1}{t} \sum_{s=1}^t X_s(i)$. \square

In the following, we provide the analysis of Algorithm 1 in both settings.

6.1. Regret Analysis in the Stochastic Setting

We first simply show that the optimal arm i^* is never deleted from the active arm set A .

Lemma 6.2. *At any time t , $i^* \in A_t$.*

Proof. At any time t , for any arm $j \in A_t$,

$$\begin{aligned} \tilde{H}_t(j) - \tilde{H}_t(i^*) &\leq \mu_j - \mu_{i^*} + \text{Radius}_t(j) + \text{Radius}_t(i^*) \\ &= -\Delta_j + \text{Radius}_t(j) + \text{Radius}_t(i^*) \\ &\leq 3\text{Radius}_t(j) + 3\text{Radius}_t(i^*), \end{aligned}$$

where the first inequality comes from Lemma 6.1. Thus the detection condition Eq.(4) is never satisfied for i^* . \square

We then prove that the Exp3.G algorithm is never started in the stochastic setting.

Lemma 6.3. *In the stochastic setting, Eq.(6) in Algorithm 1 is never satisfied and thus Exp3.G is never started.*

Proof. For each arm i , let $j \in \arg\max_{j' \in A_{\tau_i}} \tilde{H}_{\tau_i}(j')$. Then according to the definition of τ_i ,

$$\tilde{H}_{\tau_i}(j) - \tilde{H}_{\tau_i}(i) > 3\text{Radius}_{\tau_i}(j) + 3\text{Radius}_{\tau_i}(i).$$

Besides, according to Lemma 6.1,

$$\tilde{H}_{\tau_i}(j) - \tilde{H}_{\tau_i}(i) \leq \mu_j - \mu_i + \text{Radius}_{\tau_i}(j) + \text{Radius}_{\tau_i}(i)$$

$$\leq \Delta_i + \text{Radius}_{\tau_i}(j) + \text{Radius}_{\tau_i}(i).$$

Combining the above two inequalities, we have

$$\begin{aligned} \Delta_i &> 2\text{Radius}_{\tau_i}(j) + 2\text{Radius}_{\tau_i}(i) \\ &= 4\sqrt{\frac{4|D|}{\tau_i^2} \sum_{s=1}^{\tau_i} \frac{1}{\gamma_s} \log \frac{T}{\delta} + \frac{5|D|^2}{\gamma_{\tau_i}^2 \tau_i^2} \log^2 \frac{T}{\delta}}, \end{aligned} \quad (7)$$

where the last equality holds by computing $\text{Radius}_{\tau_i}(j)$ and $\text{Radius}_{\tau_i}(i)$ with $\tau_i' \geq \tau_i$ and $\tau_j' \geq \tau_i$.

Further, at any time t , according to Lemma 6.1 and the fact that $i^* \in A_t$, we have for each arm $i \notin A_t$,

$$\begin{aligned} \max_{j \in A_t} \tilde{H}_t(j) - \tilde{H}_t(i) &\geq \tilde{H}_t(i^*) - \tilde{H}_t(i) \\ &\geq \Delta_i - \text{Radius}_t(i^*) - \text{Radius}_t(i) \\ &> \text{Radius}_t(i^*) + \text{Radius}_t(i), \end{aligned} \quad (8)$$

where Eq.(8) holds by computing $\text{Radius}_t(i^*)$ and $\text{Radius}_t(i)$ and comparing them with Δ_i which is lower bounded by Eq.(7).

Let $j_t \in \arg\max_{j \in A_t} \tilde{H}_t(j)$. If $j_t = i^*$, then we have proved Eq.(6) in Algorithm 1 is not satisfied. Otherwise, there must be $\text{Radius}_t(j_t) = \text{Radius}_t(i^*)$, since both $\tau_{j_t}' > t$ and $\tau_{i^*}' > t$ when the optimal arm is not identified. Above all, no matter in either case, Exp3.G is not started. \square

In the following, we try to bound the number of arms' selections. For convenience, denote $T_i(t) := \sum_{s=1}^t \mathbb{1}\{I_s = i\}$ as the number of i 's selections at the end of t .

Lemma 6.4. *For consistency, define $\tau_i^D = 0$ for each arm $i \notin D$. Then at any time $t = 1, 2, \dots$, with probability at least $1 - \delta$, for each arm $i \in V$,*

$$\begin{aligned} T_i(t) &\leq \tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=\tau_i^D}^T \gamma_s \frac{\tau_i^D}{s} \\ &\quad + \sqrt{4 \left(\tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=\tau_i^D}^T \gamma_s \frac{\tau_i^D}{s} \right) \log \frac{T}{\delta} + 5 \log^2 \frac{T}{\delta}}. \end{aligned}$$

Proof. Let $Z_{s,i}(p) \sim \text{Bernoulli}(p)$ for time $s = 1, 2, \dots, t$, $i \in V$. Then $\forall i, t$, $T_i(t) = \sum_{s=1}^t Z_{s,i}(p_s(i))$. Define $X_s(i) = Z_{s,i}(p_s(i)) - p_s(i)$. We have $|X_s(i)| \leq 1$ and

$$\sum_{s=1}^t \mathbb{E} [X_s(i)^2 \mid \mathcal{F}_{s-1}] \leq \tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=1}^{\tau_i^D} \gamma_s \frac{\tau_i^D}{s}.$$

According to Lemma 6.7, with probability at least $1 - \delta$,

$$T_i(t) \leq \tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=1}^{\tau_i^D} \gamma_s \frac{\tau_i^D}{s} + \sum_{s=1}^t X_s(i)$$

$$\begin{aligned} &\leq \tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=1}^{\tau_i^D} \gamma_s \frac{\tau_i^D}{s} \\ &\quad + \sqrt{4 \left(\tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=1}^{\tau_i^D} \gamma_s \frac{\tau_i^D}{s} \right) \log \frac{T}{\delta} + 5 \log^2 \frac{T}{\delta}}. \end{aligned}$$

\square

Based on above results, the pseudo regret $\text{Reg}(T) = \sum_{i \neq i^*} \Delta_i T_i(T)$ in the stochastic setting can be upper bounded by

$$\begin{aligned} &\sum_{i \notin D, i \neq i^*} \Delta_i \left(\tau_i + \sqrt{4\tau_i \log \frac{T}{\delta} + 5 \log^2 \frac{T}{\delta}} \right) \\ &\quad + \sum_{i \in D, i \neq i^*} \Delta_i \left(\tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=\tau_i^D}^T \gamma_s \frac{\tau_i^D}{s} \right. \\ &\quad \left. + \sqrt{4 \left(\tau_i + \sum_{s=1}^{\tau_i^D} \gamma_s + \sum_{s=\tau_i^D}^T \gamma_s \frac{\tau_i^D}{s} \right) \log \frac{T}{\delta} + 5 \log^2 \frac{T}{\delta}} \right). \end{aligned}$$

In the following, we separately bound each part.

Bound τ_i : At time $\tau_i - 1$, since $i^* \in A_{\tau_i-1}$, it holds that

$$\begin{aligned} \Delta_i - 2\sqrt{\frac{4|D|}{(\tau_i-1)^2} \sum_{s=1}^{\tau_i-1} \frac{1}{\gamma_s} \log \frac{\tau_i}{\delta} + \frac{5|D|^2 \log^2 \frac{\tau_i}{\delta}}{\gamma_{\tau_i-1}^2 (\tau_i-1)^2}} \\ &= \mu_{i^*} - \mu_i - \text{Radius}_{\tau_i-1}(i^*) - \text{Radius}_{\tau_i-1}(i) \\ &\leq \tilde{H}_{\tau_i-1}(i^*) - \tilde{H}_{\tau_i-1}(i) \\ &\leq \max_{j \in A_{\tau_i-1}} \tilde{H}_{\tau_i-1}(j) - \tilde{H}_{\tau_i-1}(i) \\ &\leq 6\sqrt{\frac{4|D|}{(\tau_i-1)^2} \sum_{s=1}^{\tau_i-1} \frac{1}{\gamma_s} \log \frac{\tau_i}{\delta} + \frac{5|D|^2 \log^2 \frac{\tau_i}{\delta}}{\gamma_{\tau_i-1}^2 (\tau_i-1)^2}}, \end{aligned}$$

where the first equality and the last inequality hold according to the algorithm and the values of $\text{Radius}_{\tau_i-1}(i)$ for these different arms, the first inequality comes from Lemma 6.1.

We can then conclude $\tau_i \leq \frac{|D|}{K} \left(\frac{1024 \log T / \delta}{\Delta_i^2} \right)^{3/2}$.

Bound τ_i^D for $i \in D$: According to Algorithm 1, for each arm $i \in D$, if $i^* \notin N^{\text{out}}(i)$, $\tau_i^D = \max_{j: j \in N^{\text{out}}(i)} \tau_j$. And if $i^* \in N^{\text{out}}(i)$, $\tau_i^D = \max_{j: j \neq i^*} \tau_j$.

Bound $\sum_{s=1}^{\tau_i^D} \gamma_s$ for $i \in D$:

$$\sum_{s=1}^{\tau_i^D} \gamma_s = \sum_{s=1}^{\tau_i^D} K^{\frac{2}{3}} |D|^{\frac{1}{3}} s^{-\frac{1}{3}}$$

$$\begin{aligned} &\leq K^{\frac{2}{3}}|D|^{\frac{1}{3}} \left(1 + \int_1^{\tau_i^D} s^{-\frac{1}{3}} ds\right) \\ &\leq K^{\frac{2}{3}}|D|^{\frac{1}{3}} \left(1 + (\tau_i^D)^{2/3}\right) \leq O\left(K^{\frac{2}{3}}|D|^{\frac{1}{3}}(\tau_i^D)^{2/3}\right). \end{aligned}$$

Bound $\tau_i^D \sum_{s=\tau_i^D}^T \frac{\gamma_s}{s}$ for $i \in D$:

$$\begin{aligned} \tau_i^D \sum_{s=\tau_i^D}^T \frac{\gamma_s}{s} &\leq K^{2/3}|D|^{1/3} \tau_i^D \sum_{s=\tau_i^D}^T s^{-4/3} \\ &\leq 4K^{2/3}|D|^{1/3} \tau_i^D. \end{aligned}$$

Combining the above upper bound for each part, the regret can be bounded by

$$\begin{aligned} &O\left(\sum_{i \neq i^*} \frac{\Delta_i |D|}{K} \left(\frac{\log T/\delta}{\Delta_i^2}\right)^{\frac{3}{2}} + \sum_{i \in D, i \neq i^*} \Delta_i \cdot \max_{j: j \neq i^*} \right. \\ &\quad \left. \left(\frac{|D| \log T/\delta}{\Delta_j^2} + \frac{|D|^{\frac{4}{3}}}{K^{\frac{1}{3}}} \left(\frac{\log T/\delta}{\Delta_j^2}\right)^{\frac{3}{2}}\right)\right) \\ &\leq O\left(\frac{|D|^{\frac{7}{3}}}{K^{\frac{1}{3}}} \max_{j: j \neq i^*} \left(\frac{\log T/\delta}{\Delta_j^2}\right)^{\frac{3}{2}}\right) \\ &\leq O\left(|D|^2 \max_{j: j \neq i^*} \left(\frac{\log T/\delta}{\Delta_j^2}\right)^{\frac{3}{2}}\right). \end{aligned}$$

Thus we obtain the upper bound for the pseudo regret in Theorem 5.1.

6.2. Regret Analysis in the Adversarial Setting

Recall that τ is the time when Exp3.G is started. In the adversarial setting, let $i^* \in \operatorname{argmax}_{i \in V} \sum_{t=1}^{\tau} r_t(i)$ be one of the optimal arms in the first τ rounds. We first show that i^* would not be deleted from A before τ .

Lemma 6.5. *At time $\tau - 1$, $i^* \in A_{\tau-1}$.*

Proof. Denote $I^* \in \operatorname{argmax}_{i \in A_{\tau-1}} \sum_{t=1}^{\tau-1} r_t(i)$ as one of the optimal arms in $A_{\tau-1}$. Then for any arm $i \notin A_{\tau-1}$,

$$\begin{aligned} &\sum_{t=1}^{\tau-1} r_t(I^*) - \sum_{t=1}^{\tau-1} r_t(i) \\ &= (\tau - 1) \left(H_{\tau-1}(I^*) - \tilde{H}_{\tau-1}(I^*) + \tilde{H}_{\tau-1}(I^*) \right. \\ &\quad \left. - \tilde{H}_{\tau-1}(i) + \tilde{H}_{\tau-1}(i) - H_{\tau-1}(i) \right) \\ &\geq (\tau - 1) \left(-\operatorname{Radius}_{\tau-1}(I^*) - \operatorname{Radius}_{\tau-1}(i) \right. \\ &\quad \left. + \tilde{H}_{\tau-1}(I^*) - \tilde{H}_{\tau-1}(i) \right) > 0 \end{aligned}$$

where the last inequality holds based on Lemma 6.1 and the fact that Eq.(6) in Algorithm 1 is not satisfied at $\tau - 1$.

Above all, any previously deleted arm $i \notin A_{\tau-1}$ is worse than I^* , thus we can conclude that $i^* \in A_{\tau-1}$ holds. \square

In the following, we bound the averaged regret before τ compared with i^* caused by selecting different arms.

Lemma 6.6. *For each arm i such that $i \notin A_{\tau-1}$, at $\tau_i - 1$,*

$$H_{\tau_i-1}(i^*) - H_{\tau_i-1}(i) \leq 8\operatorname{Radius}_{\tau_i-1}(i).$$

Similarly, for each arm $i \in A_{\tau-1}$, at time $\tau - 1$,

$$H_{\tau-1}(i^*) - H_{\tau-1}(i) \leq 8\operatorname{Radius}_{\tau-1}(i).$$

Proof. For arm $i \notin A_{\tau-1}$, Eq.(4) in Algorithm 1 does not hold at time $\tau_i - 1$, and since $i^* \in A_{\tau_i-1}$,

$$\begin{aligned} &\tilde{H}_{\tau_i-1}(i^*) - \tilde{H}_{\tau_i-1}(i) \\ &\leq \max_{j \in A_{\tau_i-1}} \tilde{H}_{\tau_i-1}(j) - \tilde{H}_{\tau_i-1}(i) \\ &\leq 3\operatorname{Radius}_{\tau_i-1}(j) + 3\operatorname{Radius}_{\tau_i-1}(i) \\ &= 6\operatorname{Radius}_{\tau_i-1}(i). \end{aligned} \tag{9}$$

where Eq.(9) holds since at $\tau_i - 1$, $\operatorname{Radius}_{\tau_i-1}(j) = \operatorname{Radius}_{\tau_i-1}(i)$. Besides, at $\tau_i - 1$, $\operatorname{Radius}_{\tau_i-1}(i^*) = \operatorname{Radius}_{\tau_i-1}(i)$ also holds according to Algorithm 1. Thus based on Lemma 6.1,

$$\begin{aligned} &\tilde{H}_{\tau_i-1}(i^*) - \tilde{H}_{\tau_i-1}(i) \\ &\geq H_{\tau_i-1}(i^*) - H_{\tau_i-1}(i) - \operatorname{Radius}_{\tau_i-1}(i^*) - \operatorname{Radius}_{\tau_i-1}(i) \\ &= H_{\tau_i-1}(i^*) - H_{\tau_i-1}(i) - 2\operatorname{Radius}_{\tau_i-1}(i). \end{aligned} \tag{10}$$

Combining Eq.(9) and (10), we can obtain the first inequality in Lemma 6.6. Similar arguments hold at time $\tau - 1$ for arm $i \in A_{\tau-1}$ to get the second inequality. \square

Based on the above results, taking the expectation over the randomness of I_t , we are able to bound the expected regret during the first τ rounds as follows.

$$\begin{aligned} &\mathbb{E} \left[\sum_{t=1}^{\tau} r_t(i^*) - r_t(I_t) \right] \\ &= \sum_{t=1}^{\tau} \left(r_t(i^*) - \sum_{i \in V} p_t(i) r_t(i) \right) \\ &= \sum_{t=1}^{\tau} \left(r_t(i^*) - \sum_{i \in V} (1 - \gamma_t) p_{t,A}(i) r_t(i) \right. \\ &\quad \left. - \sum_{i \in V} \gamma_t p_{t,D}(i) r_t(i) \right) \\ &= \sum_{t=1}^{\tau} \gamma_t \sum_{i \in V} p_{t,D}(i) (r_t(i^*) - r_t(i)) \\ &\quad + \sum_{t=1}^{\tau} (1 - \gamma_t) \sum_{i \in V} p_{t,A}(i) (r_t(i^*) - r_t(i)) \end{aligned}$$

$$\begin{aligned} &\leq \sum_{t=1}^{\tau} \gamma_t + \sum_{t=1}^{\tau} \sum_{i \in V} p_{t,A}(i) (r_t(i^*) - r_t(i)) \\ &\leq \sum_{t=1}^{\tau} \gamma_t + \sum_{i \in V} \sum_{t=1}^{\min\{\tau_i, \tau\}} (r_t(i^*) - r_t(i)) \end{aligned} \quad (11)$$

$$\begin{aligned} &\leq \sum_{t=1}^{\tau} \gamma_t + \sum_{i \in V} (\mathbb{1}\{i \notin A_{\tau}\} \cdot \tau_i (H_{\tau_i}(i^*) - H_{\tau_i}(i)) \\ &\quad + \mathbb{1}\{i \in A_{\tau}\} \cdot \tau (H_{\tau}(i^*) - H_{\tau}(i))) \\ &\leq \sum_{t=1}^{\tau} \gamma_t + O \left(\sqrt{K^2 |D| \sum_{t=1}^{\tau} \frac{1}{\gamma_t} \log \frac{T}{\delta}} \right) \end{aligned} \quad (12)$$

$$\leq O \left(K^{2/3} |D|^{1/3} \tau^{2/3} \sqrt{\log \frac{T}{\delta}} \right), \quad (13)$$

where Eq.(11) holds according to Algorithm 1, Eq.(12) comes from Lemma 6.6 and the exact value of $\text{Radius}_t(i)$ for $t = \tau_i - 1, \tau - 1$. The last inequality holds by replacing γ_t with $K^{2/3} |D|^{1/3} t^{-1/3}$.

We can divide the total horizon into two phases and bound the cumulative expected regret over T rounds by $\mathbb{E}[\text{Reg}(T)] \leq \mathbb{E}[\sum_{t=1}^{\tau} r_t(i^*) - r_t(I_t)] + \mathbb{E}[\max_{i \in V} \sum_{t=\tau+1}^T r_t(i) - r_t(I_t)]$. The regret in the first phase is upper bounded by Eq.(13). For the second phase, as shown in Alon et al. (2015, Theorem 2), the regret is of order $O((|D| \log K)^{1/3} T^{2/3})$. Above all, the regret upper bound in Theorem 5.1 is obtained.

6.3. Technical Lemma

Lemma 6.7 (Lemma 4.4 in Bubeck & Slivkins (2012)). *Let $\mathcal{F}_1 \subseteq \mathcal{F}_2 \cdots \mathcal{F}_n$ be filtrations, and X_1, X_2, \dots, X_n be real random variables such that X_t is \mathcal{F}_t measurable, $\mathbb{E}[X_t | \mathcal{F}_t] = 0$ and $|X_t| \leq b$ for some $b > 0$. Denote $V_n = \sum_{t=1}^n \mathbb{E}[X_t^2 | \mathcal{F}_{t-1}]$ and $\delta > 0$. Then with probability at least $1 - \delta$,*

$$\sum_{t=1}^n X_t \leq \sqrt{4V_n \log \frac{n}{\delta} + 5b^2 \log^2 \frac{n}{\delta}}.$$

7. Conclusion

In this paper, we propose the first algorithm for online learning with general graph feedback that achieves near-optimal regret in both stochastic and adversarial settings. Compared with previous BoBW works, we face a more challenging problem where only partial observations are available. To simultaneously preserve the minimax-optimality of $O(T^{2/3})$ in a much harder adversarial setting and the instance-optimality in the stochastic setting, we introduce a new trade-off mechanism between exploration and exploitation with carefully designed proportions. Without knowing

the environment type in advance, the regret upper bound of our proposed algorithm matches the near-optimal results in both settings only up to some logarithmic and constant factors.

Since the current upper bound in the adversarial setting is on the expected regret, one of the possible directions is to derive a high-probability pseudo regret bound as Bubeck & Slivkins (2012). It is more difficult as the pseudo regret in this problem with general graph feedback can only be decomposed into unobservable quantities, and similar techniques in Bubeck & Slivkins (2012) could not be adopted. The other interesting future direction is to investigate possible regularizers for OMD-type algorithms as Zimmert & Seldin (2019); Erez & Koren (2021), which are hoped to achieve tighter regret bounds in both settings.

Acknowledgement

The corresponding author Shuai Li is supported by National Natural Science Foundation of China (62006151, 62076161). This work is sponsored by Shanghai Sailing Program.

References

- Alon, N., Cesa-Bianchi, N., Dekel, O., and Koren, T. Online learning with feedback graphs: Beyond bandits. In *Conference on Learning Theory*, pp. 23–35. PMLR, 2015.
- Auer, P. and Chiang, C.-K. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Conference on Learning Theory*, pp. 116–120. PMLR, 2016.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002a.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002b.
- Bubeck, S. and Slivkins, A. The best of both worlds: Stochastic and adversarial bandits. In *Conference on Learning Theory*, pp. 42–1. JMLR Workshop and Conference Proceedings, 2012.
- Buccapatnam, S., Eryilmaz, A., and Shroff, N. B. Stochastic bandits with side observations on networks. In *The 2014 ACM international conference on Measurement and modeling of computer systems*, pp. 289–300, 2014.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge university press, 2006.

- Chen, H., Huang, Z., Li, S., and Zhang, C. Understanding bandits with graph feedback. In *Advances in Neural Information Processing Systems*, 2021.
- Cohen, A., Hazan, T., and Koren, T. Online learning with feedback graphs without the graphs. In *International Conference on Machine Learning*, pp. 811–819. PMLR, 2016.
- Eban, E., Birnbaum, A., Shalev-Shwartz, S., and Globerson, A. Learning the experts for online sequence prediction. In *International Conference on Machine Learning*, pp. 259–266, 2012.
- Erez, L. and Koren, T. Towards best-of-all-worlds online learning with feedback graphs. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- Kocák, T., Neu, G., Valko, M., and Munos, R. Efficient learning by implicit exploration in bandit problems with side observations. In *Advances in Neural Information Processing Systems*, pp. 613–621, 2014.
- Lai, T. L. and Robbins, H. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1): 4–22, 1985.
- Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2020.
- Lee, C.-W., Luo, H., Wei, C.-Y., Zhang, M., and Zhang, X. Achieving near instance-optimality and minimax-optimality in stochastic and adversarial linear bandits simultaneously. In *International Conference on Machine Learning*, 2021.
- Li, S., Chen, W., Wen, Z., and Leung, K.-S. Stochastic online learning with probabilistic graph feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 4675–4682, 2020.
- Lu, S., Wang, G., and Zhang, L. Stochastic graphical bandits with adversarial corruptions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI), to appear*, 2021.
- Mannor, S. and Shamir, O. From bandits to experts: On the value of side-observations. In *Advances in Neural Information Processing Systems*, volume 24, pp. 684–692, 2011.
- Seldin, Y. and Lugosi, G. An improved parametrization and analysis of the exp3++ algorithm for stochastic and adversarial bandits. In *Conference on Learning Theory*, pp. 1743–1759. PMLR, 2017.
- Seldin, Y. and Slivkins, A. One practical algorithm for both stochastic and adversarial bandits. In *International Conference on Machine Learning*, pp. 1287–1295. PMLR, 2014.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Tossou, A. C., Dimitrakakis, C., and Dubhashi, D. Thompson sampling for stochastic bandits with graph feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017.
- Wu, Y., György, A., and Szepesvári, C. Online learning with gaussian payoffs and side observations. In *Advances in Neural Information Processing Systems*, pp. 1360–1368, 2015.
- Zimmert, J. and Seldin, Y. An optimal algorithm for stochastic and adversarial bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 467–475. PMLR, 2019.