

A Near-Optimal Change-Detection Based Algorithm for Piecewise-Stationary Combinatorial Semi-Bandits

Huozhi Zhou,^{1*} Lingda Wang,^{1*} Lav R. Varshney,¹ Ee-Peng Lim²

¹Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign

²School of Information Systems, Singapore Management University
{hzhou35, lingdaw2, varshney}@illinois.edu, eplim@smu.edu.sg

Abstract

We investigate the piecewise-stationary combinatorial semi-bandit problem. Compared to the original combinatorial semi-bandit problem, our setting assumes the reward distributions of base arms may change in a piecewise-stationary manner at unknown time steps. We propose an algorithm, GLR-CUCB, which incorporates an efficient combinatorial semi-bandit algorithm, CUCB, with an almost parameter-free change-point detector, the *Generalized Likelihood Ratio Test* (GLRT). Our analysis shows that the regret of GLR-CUCB is upper bounded by $\mathcal{O}(\sqrt{NKT \log T})$, where N is the number of piecewise-stationary segments, K is the number of base arms, and T is the number of time steps. As a complement, we also derive a nearly matching regret lower bound on the order of $\Omega(\sqrt{NKT})$, for both piecewise-stationary multi-armed bandits and combinatorial semi-bandits, using information-theoretic techniques and judiciously constructed piecewise-stationary bandit instances. Our lower bound is tighter than the best available regret lower bound, which is $\Omega(\sqrt{T})$. Numerical experiments on both synthetic and real-world datasets demonstrate the superiority of GLR-CUCB compared to other state-of-the-art algorithms.

1 Introduction

The multi-armed bandit (MAB) problem, first proposed by Thompson (1933), has been studied extensively in the statistics and machine learning communities, as it models many online decision making problems such as online recommendation (Li, Karatzoglou, and Gentile 2016), computational advertising (Tang et al. 2013), and crowdsourcing task allocation (Hassan and Curry 2014). The classical MAB is modeled as an agent repeatedly pulling one of K arms and observing the reward generated, with the goal of minimizing the *regret* which is the difference between the reward of the optimal arm in hindsight and the reward of the arm chosen by the agent. This classical problem is well understood for both stochastic (Lai and Robbins 1985) and adversarial settings (Auer et al. 2002). The stochastic setting is when the reward of each arm is generated from a fixed distribution, and it is well known that the problem-dependent

lower bound is of order $\Omega(\log T)$ (Lai and Robbins 1985), where T is the number of time steps. Several algorithms have been proposed and proven to achieve $\mathcal{O}(\log T)$ regret (Agrawal and Goyal 2012; Auer, Cesa-Bianchi, and Fischer 2002). The adversarial setting is when at each time step the environment generates the reward in an adversarial manner, whose minimax regret lower bound is of order $\Omega(\sqrt{T})$. Several algorithms achieving order-optimal (up to poly-logarithm factors) regret have been proposed in recent years (Hazan and Kale 2011; Bubeck, Cohen, and Li 2018; Li, Chen, and Giannakis 2019).

Many real-world applications, however, have a combinatorial nature that cannot be fully characterized by the classical MAB model. For example, online movie sites aim to recommend multiple movies to the users to maximize their utility under some constraints (e.g., recommend at most one movie for each category). This phenomenon motivates the study of combinatorial semi-bandits (CMAB), which aims to identify the best super arm, a set of base arms with the highest aggregated reward. Several algorithms for stochastic CMAB with provable guarantees based on optimism principle have been proposed recently (Chen, Wang, and Yuan 2013; Kveton et al. 2015; Combes et al. 2015). All of these algorithms use some oracle to overcome the curse of dimensionality of the action space for solving some combinatorial optimization problem at each iteration. In addition, these algorithms achieve the optimal $\mathcal{O}(C \log T)$ regret upper bound, where C is an instance-dependent parameter. Adversarial CMAB is also well studied. Several algorithms achieve optimal regret of order $\mathcal{O}(\sqrt{T})$ based on either *Follow-the-Regularized-Leader* or *Follow-the-Perturbed-Leader*, both of which are general frameworks for adversarial online learning algorithm design (Hazan 2016). Moreover, one recent study has developed an algorithm that is order-optimal for both stochastic and adversarial CMAB (Zimmert, Luo, and Wei 2019).

Although both stochastic and adversarial CMAB are well-studied, understanding of the scenario lying in the “middle” of these two settings is still limited. Such a “middle” setting where the reward distributions of base arms slowly change over time may be a more realistic model in many applications. For instance, in online recommendation sys-

*The first two authors contributed equally to this work
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tems, users' preference is unlikely to be either time-invariant or to change significantly and frequently over time. Thus in this case it would be too ideal to assume the stochastic CMAB model and too conservative to assume the adversarial CMAB model. Similar situations appear in web search, online advertisement, and crowdsourcing (Yu and Mannor 2009; Pereira et al. 2018; Vempaty, Varshney, and Varshney 2014). As such, we investigate a setting lying between these two standard CMAB models, namely piecewise-stationary combinatorial semi-bandit, which we will define formally in Section 2. Piecewise-stationary CMAB is a natural generalization of the piecewise-stationary MAB model (Hartland et al. 2007; Kocsis and Szepesvári 2006; Garivier and Moulines 2011), and can be interpreted as an approximation to the slow-varying CMAB problem. Roughly, compared to the stochastic CMAB, we assume reward distributions of base arms remain fixed for certain time periods called *piecewise-stationary segments*, but can change abruptly at some unknown time steps, called *change-points*.

Previous works on piecewise-stationary MAB may be divided into two categories: *passively adaptive approaches* (Garivier and Moulines 2011; Besbes, Gur, and Zeevi 2014; Wei and Srivatsva 2018) and *actively adaptive approaches* (Cao et al. 2019; Liu, Lee, and Shroff 2018; Besson and Kaufmann 2019; Auer, Gajane, and Ortner 2019). Passively adaptive approaches make decisions based on the most recent observations and are unaware of the underlying distribution changes. On the contrary, actively adaptive approaches incorporate a change-point detector subroutine to monitor the reward distributions, and restart the algorithm once a change-point is detected. Numerous empirical experiments have shown that actively adaptive approaches outperform passively adaptive approaches (Mellor and Shapiro 2013), which motivates us to adopt an actively adaptive approach.

Our main contributions include the following:

- We propose a simple and general algorithm for piecewise-stationary CMAB, named GLR-CUCB, which is based on CUCB (Chen, Wang, and Yuan 2013) with a novel change-point detector, the *generalized likelihood ratio test* (GLRT) (Besson and Kaufmann 2019). The advantage of GLRT change-point detector is that it is almost parameter-free and thus easy to tune compared to previously proposed change-point detection methods used in nonstationary MAB, such as CUSUM (Liu, Lee, and Shroff 2018) and SW (Cao et al. 2019).
- For any combinatorial action set, we derive the problem-dependent regret bound for GLR-CUCB under mild assumptions (see Section 4). When the number of change-points N is known beforehand, the regret of GLR-CUCB is upper bounded by $\mathcal{O}(C_1NK^2 \log T + C_2\sqrt{NKT} \log T)$ (nearly order-optimal within polylogarithm factor in T), where K is number of base arms. When N is unknown, the algorithm achieves $\mathcal{O}(C_1NK^2 \log T + C_2N\sqrt{KT} \log T)$. Here, C_1 and C_2 are problem-dependent constants which do not depend on T , N , or K .
- We derive a tighter minimax lower bound for both

piecewise-stationary MAB and piecewise-stationary CMAB on the order of $\Omega(\sqrt{NKT})$. Since piecewise-stationary MAB is a special instance of piecewise-stationary CMAB in which every super arm is a single arm, thus any minimax lower bound holds for piecewise-stationary MAB also holds for piecewise-stationary CMAB. To the best of our knowledge, this is the best existing minimax lower bound for piecewise-stationary CMAB. Previously, the best available lower bound is $\Omega(\sqrt{T})$ (Garivier and Moulines 2011), which does not depend on N or K .

- We demonstrate that GLR-CUCB performs significantly better than state-of-the-art algorithms through experiments on both synthetic and real-world datasets.

The remainder of this paper is organized as follows: the formal problem formulation and some preliminaries are introduced in Section 2, then the proposed GLR-CUCB algorithm in Section 3. We derive the upper bound on the regret of our algorithm in Section 4, and the minimax regret lower bound in Section 5. Section 6 gives our experiment results. Finally, we conclude the paper. Due to the page limitation, we postpone proofs and additional experimental results to the appendix¹.

2 Problem Formulation and Background

In this section, we start with the formal definition of piecewise-stationary combinatorial semi-bandit as well as some technical assumptions in Section 2.1. Then, we introduce the GLR change-point detector used in our algorithm design, and its advantage in Section 2.2.

2.1 Piecewise-Stationary Combinatorial Semi-Bandits

A piecewise-stationary combinatorial semi-bandit is characterized by a tuple $(\mathcal{K}, \mathcal{F}, \mathcal{T}, \{f_{k,t}\}_{k \in \mathcal{K}, t \in \mathcal{T}}, r_{\mu_t}(S_t))$. Here, $\mathcal{K} = \{1, \dots, K\}$ is the set of K base arms; $\mathcal{F} \subseteq 2^{\mathcal{K}}$ is the set of all super arms; $\mathcal{T} = \{1, \dots, T\}$ is a sequence of T time steps; $f_{k,t}$ is the reward distribution of arm k at time t with mean $\mu_{k,t}$ and bounded support within $[0, 1]$; $r_{\mu_t}(S_t) : \mathcal{F} \times [0, 1]^K \mapsto \mathbb{R}$ is the expected reward function defined on the super arm S_t and mean vector of all base arms $\mu_t := [\mu_{1,t}, \mu_{2,t}, \dots, \mu_{K,t}]^T$ at time t . Like Chen, Wang, and Yuan (2013), we assume the expected reward function $r_{\mu}(S)$ satisfies the following two properties:

Assumption 2.1 (Monotonicity). Given two arbitrary mean vectors μ and μ' , if $\mu_k \geq \mu'_k, \forall k \in \mathcal{K}$, then $r_{\mu}(S) \geq r_{\mu'}(S)$.

Assumption 2.2 (L -Lipschitz). Given two arbitrary mean vectors μ and μ' , there exists an $L < \infty$ such that $|r_{\mu}(S) - r_{\mu'}(S)| \leq L \|\mathcal{P}_S(\mu - \mu')\|_2, \forall S \in \mathcal{F}$, where $\mathcal{P}_S(\cdot)$ is the projection operator specified as $\mathcal{P}_S(\mu) = [\mu_1 \mathbb{I}\{1 \in S\}, \dots, \mu_k \mathbb{I}\{k \in S\}, \dots, \mu_K \mathbb{I}\{K \in S\}]^T$ in terms of the indicator function $\mathbb{I}\{\cdot\}$.

¹<https://arxiv.org/abs/1908.10402>

In the piecewise i.i.d. model, we define N , the number of piecewise-stationary segments in the reward process, to be

$$N = 1 + \sum_{t=1}^{T-1} \mathbb{I}\{\exists k \in \mathcal{K} \text{ s.t. } f_{k,t} \neq f_{k,t+1}\}.$$

We denote these $N - 1$ change-points as $\nu_1, \nu_2, \dots, \nu_{N-1}$ respectively, and we let $\nu_0 = 0$ and $\nu_N = T$. For each piecewise-stationary segment $t \in [\nu_{i-1} + 1, \nu_i]$, we use f_k^i and μ_k^i to denote the reward distribution and the expected reward of arm k on the i th piecewise-stationary segment, respectively. The vector encoding the expected rewards of all base arms at the i th segment is denoted as $\boldsymbol{\mu}^i = [\mu_1^i, \dots, \mu_K^i]^\top$, $\forall i = 1, \dots, N$. Note that when a change-point occurs, there must be at least one arm whose reward distribution has changed, however, the rewards distributions of all base arms do not necessarily change.

For a piecewise-stationary combinatorial semi-bandit problem, at each time step t , the learning agent chooses a super arm $S_t \in \mathcal{F}$ to play based on the rewards observed up to time t . When the agent plays a super arm S_t , the reward $\{X_{I_t}\}_{I_t \in S_t}$ of base arms contained in super arm S_t are revealed to the agent and the reward of super arm $R_t(S_t)$ as well. We assume that the agent has access to an α -approximation oracle, to carry out combinatorial optimization, defined as follows.

Assumption 2.3 (α -approximation oracle). Given a mean vector $\boldsymbol{\mu}$, the α -approximation oracle $\text{Oracle}_\alpha(\boldsymbol{\mu})$ outputs an α -suboptimal super arm S such that $r_\mu(S) \geq \alpha \max_{S \in \mathcal{F}} r_\mu(S)$.

Remark 2.1. The approximation oracle assumption was first proposed in Chen, Wang, and Yuan (2013) for the combinatorial semi-bandit setting. This assumption is reasonable since many combinatorial NP-hard problems admit approximation algorithms, which can be solved efficiently in polynomial time (Ausiello, Crescenzi, and Protasi 1995). There are also many combinatorial problems which are not NP hard and can be solved efficiently. One example is the top- m arm identification problem in the bandit setting (Cao et al. 2015), where any efficient sorting algorithm suffices.

As only an α -approximation oracle is used for optimization, it is reasonable to use expected α -approximation cumulative regret to measure the performance of the learning agent, defined as follows.

Definition 2.4 (Expected α -approximation cumulative regret). The agent's policy is evaluated by its expected α -approximation cumulative regret,

$$\mathcal{R}(T) = \mathbb{E} \left[\alpha \sum_{t=1}^T \max_{S \in \mathcal{F}} r_{\boldsymbol{\mu}^t}(S) - \sum_{t=1}^T r_{\boldsymbol{\mu}^t}(S_t) \right],$$

where the expectation $\mathbb{E}[\cdot]$ is taken with respect to the selection of $\{S_t | S_t \in \mathcal{F}\}$.

2.2 Generalized Likelihood Ratio Change-Point Detector for Sub-Bernoulli Distribution

Sequential change-point detection is a classical problem in statistical sequential analysis, but most existing works

make additional assumptions on the pre-change and post-change distributions which might not hold in the bandit setting (Siegmund 2013; Basseville and Nikiforov 1993). In general, designing algorithms with provable guarantees for change-point detection with little assumption on pre-change and post-change distributions is very challenging. In our algorithm design, we will use the GLR change-point detector (Besson and Kaufmann 2019), which works for any sub-Bernoulli distribution. Compared to other existing change-point detectors used in piecewise-stationary MAB, the GLR change-point detector has less parameters to be tuned and needs less prior knowledge for the bandit instance. Specifically, GLR change-point detector only needs to tune the threshold δ , and does not require the smallest change in expectation among all change-points. On the contrary, CUSUM (Liu, Lee, and Shroff 2018) and SW (Cao et al. 2019) both need more parameters to be tuned and need to know the smallest magnitude among all change-points beforehand, which limits their practicality.

To define the GLR change-point detector, we need some more definitions for clarity. A distribution f is said to be sub-Bernoulli if $\mathbb{E}_{X \sim f}[e^{\lambda X}] \leq e^{\phi_\mu(\lambda)}$, where $\mu = \mathbb{E}_{X \sim f}[X]$; $\phi_\mu(\lambda) = \log(1 - \mu - \mu e^\lambda)$ is the log moment generating function of a Bernoulli distribution with mean μ . Notice that the support of reward distribution $f_{k,t}$, $\forall k \in \mathcal{K}$, $t \in \mathcal{T}$ is a subset of the interval $[0, 1]$, thus all $\{f_{k,t}\}$ are sub-Bernoulli distributions with mean $\{\mu_{k,t}\}$, due to the following lemma.

Lemma 2.5 (Lemma 1 in Cappé et al. (2013)). Any distribution f with bounded support within the interval $[0, 1]$ is a sub-Bernoulli distribution that satisfies:

$$\mathbb{E}_{X \sim f}[e^{\lambda X}] \leq e^{\phi_\mu(\lambda)}.$$

Suppose we have either a time sequence $\{X_t\}_{t=1}^n$ drawn from a sub-Bernoulli distribution for any $t \leq n$ or two sub-Bernoulli distributions with an unknown change-point $s \in [1, n - 1]$. This problem can be formulated as a parametric sequential test:

$$\begin{aligned} \mathcal{H}_0 &: \exists f_0 : X_1, \dots, X_n \stackrel{i.i.d.}{\sim} f_0, \\ \mathcal{H}_1 &: \exists f_0 \neq f_1, \tau \in [1, n - 1] : X_1, \dots, X_\tau \stackrel{i.i.d.}{\sim} f_0 \\ &\text{and } X_{\tau+1}, \dots, X_n \stackrel{i.i.d.}{\sim} f_1. \end{aligned}$$

The GLR statistic for sub-Bernoulli distributions is:

$$\text{GLR}(n) = \sup_{s \in [1, n-1]} [s \times \text{kl}(\hat{\mu}_{1:s}, \hat{\mu}_{1:n}) + (n - s) \times \text{kl}(\hat{\mu}_{s+1:n}, \hat{\mu}_{1:n})], \quad (1)$$

where $\hat{\mu}_{s:s'}$ is the mean of the observations collected between s and s' , and $\text{kl}(x, y)$ is the binary relative entropy between Bernoulli distributions,

$$\text{kl}(x, y) = x \log \left(\frac{x}{y} \right) + (1 - x) \log \left(\frac{1 - x}{1 - y} \right).$$

If the GLR in Eq. (1) is large, it indicates that hypothesis \mathcal{H}_1 is more likely. Now, we are ready to define the sub-Bernoulli GLR change-point detector with confidence level $\delta \in (0, 1)$.

Definition 2.6. The sub-Bernoulli GLR change-point detector with threshold function $\beta(n, \delta)$ is

$$\tau := \inf\{n \in \mathbb{N} : \sup_{s \in [1, n-1]} [s \times \text{kl}(\hat{\mu}_{1:s}, \hat{\mu}_{1:n}) + (n-s) \times \text{kl}(\hat{\mu}_{s+1:n}, \hat{\mu}_{1:n})] \geq \beta(n, \delta)\},$$

where $\beta(n, \delta) = 2\mathcal{Q}\left(\frac{\log(3n\sqrt{n}/\delta)}{2}\right) + 6\log(1 + \log(n))$, and $\mathcal{Q}(\cdot)$ is the same as in Eq. (13) in Kaufmann and Koolen (2018).

The pseudo-code of sub-Bernoulli GLR change-point detector is summarized in Algorithm 1 for completeness.

Algorithm 1 Sub-Bernoulli GLR Change-Point Detector: $\text{GLR}(X_1, \dots, X_n; \delta)$

Require: observations X_1, \dots, X_n and confidence level δ .

- 1: **if** $\sup_{s \in [1, n-1]} [s \times \text{kl}(\hat{\mu}_{1:s}, \hat{\mu}_{1:n}) + (n-s) \times \text{kl}(\hat{\mu}_{s+1:n}, \hat{\mu}_{1:n})] \geq \beta(n, \delta)$ **then**
- 2: Return **True**
- 3: **else**
- 4: Return **False**
- 5: **end if**

3 The GLR-CUCB Algorithm

Our proposed algorithm, GLR-CUCB, incorporates an efficient combinatorial semi-bandit algorithm CUCB (Chen, Wang, and Yuan 2013) with a change-point detector running on each base arm (See Algorithm 2). The GLR-CUCB requires the number of time steps T , the number of base arms K , uniform exploration probability p , and the confidence level δ as inputs. Let τ denote the last change-point detection time, and n_k denote the number of observations of base arm k after τ , which are both initialized as zero at the beginning of the algorithm.

At each time step, the GLR-CUCB first determines if it will enter forced uniform exploration (to ensure each base arm collects sufficient samples for the change-point detection) according to the condition in line 3. If it is in a forced exploration, a random super arm S_t that contains a (line 4) is played, to ensure sufficient number of samples are collected for each base arm. Otherwise, the next super arm S_t to be played is determined by the α -approximation oracle $\text{Oracle}_\alpha(\cdot)$ (line 6) given the UCB indices (line 19). Then, the learning agent plays the super arm S_t , and gets the reward $R_t(S_t)$ of the super arm S_t and the rewards $\{X_{I_t}\}_{I_t \in S_t}$ of the base arm I_t 's that are contained in the super arm S_t (line 8). In the next step, the algorithm updates the statistics for each base arm (lines 10-11) in order to run the GLR change-point detector (Algorithm 1) with confidence level δ (line 12). If the GLR change-point detector detects a change in distribution for any of the base arms, the algorithm sets τ to be the current time step and all n_k 's to be 0 (line 13-14) before going into time step $t+1$. Lastly, the UCB indices of all base arms are updated (line 19).

Remark 3.1. The uniform exploration is necessary for this algorithm, and similar strategy has been adopted in Liu, Lee, and Shroff (2018); Cao et al. (2019). Intuitively, uniform

Algorithm 2 The GLR-CUCB Algorithm

Require: $T \in \mathbb{N}$, $K \in \mathbb{N}$, exploration probability $p \in (0, 1)$, confidence level $\delta \in (0, 1)$.

- 1: **Initialization:** $\forall k \in \mathcal{K}$, $n_k \leftarrow 0$; $\tau \leftarrow 0$.
- 2: **for all** $t = 1, 2, \dots, T$ **do**
- 3: **if** $p > 0$ and $a \leftarrow (t - \tau) \bmod \lfloor \frac{K}{p} \rfloor \in \mathcal{K}$ **then**
- 4: Randomly choose S_t with $a \in S_t$.
- 5: **else**
- 6: $S_t = \text{Oracle}_\alpha(\text{UCB})$.
- 7: **end if**
- 8: Play super arm S_t and get the reward $R_t(S_t)$ and $X_{I_t, t}$, $\forall I_t \in S_t$.
- 9: **for all** $I_t \in S_t$ **do**
- 10: $Z_{I_t, n_{I_t}} \leftarrow X_{I_t, t}$.
- 11: $n_{I_t} \leftarrow n_{I_t} + 1$.
- 12: **if** $\text{GLR}(Z_{I_t, 1}, \dots, Z_{I_t, n_{I_t}}; \delta) = \text{True}$ **then**
- 13: $n_k \leftarrow 0, \forall k \in \mathcal{K}$.
- 14: $\tau \leftarrow t$.
- 15: **end if**
- 16: **end for**
- 17: **for all** $k = 1, \dots, K$ **do**
- 18: **if** $n_k \neq 0$ **then**
- 19: $\text{UCB}(k) \leftarrow \frac{1}{n_k} \sum_{n=1}^{n_k} Z_{k, n} + \sqrt{\frac{3 \log(t-\tau)}{2n_k}}$.
- 20: **end if**
- 21: **end for**
- 22: **end for**

exploration ensures each base arm gathers sufficient samples to guarantee quick change detection whereas pure UCB exploration is incapable of this. One more rigorous argument is given in Garivier and Moulines (2011), which shows that theoretically pure UCB exploration performs badly on piecewise-stationary MAB.

Remark 3.2. Thompson sampling (TS) often performs better than UCB policy in empirical simulations, but it has been shown that one cannot incorporate an approximate oracle in TS for even MAB problems (Wang and Chen 2018). Thus our algorithm adopts UCB policy for the bandit component to ensure compatibility with approximation oracle.

4 Regret Upper Bound

In this section, we analyze the T -step regret of our proposed algorithm GLR-CUCB. Recall T is the time horizon, N is the number of piecewise-stationary segments, ν_1, \dots, ν_{N-1} are the change-points, and for each segment $i \in [N]$, $\mu^i \in \mathbb{R}^K$ is the vector encoding the expected rewards of all base arms. A super arm S is *bad* with respect to the i th piecewise-stationary distributions if $r_{\mu^i}(S) \leq \alpha \max_{\tilde{S} \in \mathcal{F}} r_{\mu^i}(\tilde{S})$. We define $S_B^i = \{S | r_{\mu^i}(S) \leq \alpha \max_{\tilde{S} \in \mathcal{F}} r_{\mu^i}(\tilde{S})\}$ to be the set of bad super arms with respect to the i th piecewise-stationary segment. We define the *suboptimality gap* in the i th stationary segment as follows:

$$\Delta_{\text{opt}}^{\min, i} = \alpha \max_{\tilde{S} \in \mathcal{F}} r_{\mu^i}(\tilde{S}) - \max\{r_{\mu^i}(S) | S \in S_B^i\},$$

$$\Delta_{\text{opt}}^{\max, i} = \alpha \max_{\tilde{S} \in \mathcal{F}} r_{\mu^i}(\tilde{S}) - \min\{r_{\mu^i}(S) | S \in S_B^i\}.$$

Furthermore, let $\Delta_{\text{opt}}^{\max} = \max_{i \in [N]} \Delta_{\text{opt}}^{\max, i}$ and $\Delta_{\text{opt}}^{\min} = \min_{i \in [N]} \Delta_{\text{opt}}^{\min, i}$ be the maximum and minimum sub-optimal gaps for the whole time horizon, respectively. Lastly, denote the largest gap at change-point ν_i as $\Delta_{\text{change}}^i = \max_{k \in \mathcal{K}} |\mu_k^{i+1} - \mu_k^i|$, $\forall 1 \leq i \leq N-1$, and $\Delta_{\text{change}}^0 = \max_{k \in \mathcal{K}} |\mu_k^1|$. We need the following assumption for our theoretical analysis.

Assumption 4.1. Define $d_i = d_i(p, \delta) = \left[\{4K/p (\Delta_{\text{change}}^i)^2\} \beta(T, \delta) + \frac{K}{p} \right]$ and assume $\nu_i - \nu_{i-1} \geq 2 \max\{d_i, d_{i-1}\}$, $\forall i = 1, \dots, N-1$, where $\nu_N - \nu_{N-1} \geq 2d_{N-1}$.

Tuning δ and p properly (See Corollary 4.3), and applying the upper bound on $\mathcal{Q}(x)$ by Kaufmann and Koolen (2018) with $x \geq 5$,

$$\mathcal{Q}(x) \leq x + 4 \log(1 + x + \sqrt{2x}),$$

the length of each piecewise-stationary segment is $\Omega(\sqrt{T \log T})$. Roughly, we assume the length of each stationary segment to be sufficiently long, in order to let the GLR change-point detector detect the change in distribution within a reasonable delay with high probability. Similar assumption on the length of stationary segments also appears in other literature on piecewise stationary MAB (Liu, Lee, and Shroff 2018; Cao et al. 2019; Besson and Kaufmann 2019). Note that Assumption 4.1 is only required for the theoretical analysis; Algorithm 2 can be implemented regardless of this assumption. Now we are ready to state the regret upper bound for Algorithm 2.

Theorem 4.2. Running GLR-CUCB with Assumptions 2.1, 2.2, and 4.1, the expected α -approximation cumulative regret of GLR-CUCB with exploration probability p and confidence level δ satisfies

$$\mathcal{R}(T) \leq \underbrace{\sum_{i=1}^N \tilde{C}_i}_{(a)} + \underbrace{\Delta_{\text{opt}}^{\max} T p}_{(b)} + \underbrace{\sum_{i=1}^{N-1} \Delta_{\text{opt}}^{\max, i+1} d_i}_{(c)} + \underbrace{3NT \Delta_{\text{opt}}^{\max} K \delta}_{(d)},$$

where $\tilde{C}_i = \left(6L^2 K^2 \log T / (\Delta_{\text{opt}}^{\min, i})^2 + \pi^2/6 + K \right) \Delta_{\text{opt}}^{\max, i}$.

Theorem 4.2 indicates that the regret comes from four sources. Terms (a) and (b) correspond to the cost of exploration, while terms (c) and (d) correspond to the cost of change-point detection. More specifically, term (a) is due to the UCB exploration, term (b) is due to the uniform exploration, term (c) is due to the expected delay of GLR change-point detector, and term (d) is due to the false alarms of GLR change-point detector. We need to carefully tune the exploration probability p and confidence level δ of GLR change-point detector to balance the trade-off.

The following corollary directly follows Theorem 4.2 by properly tuning the parameters in the algorithm.

Corollary 4.3. Let $\Delta_{\text{change}}^{\min} = \min_{i \in [N-1]} \Delta_{\text{change}}^i$, we have

1. (N is known) Choosing $\delta = \frac{1}{T}$, $p = \sqrt{\frac{NK \log T}{T}}$, gives

$$\mathcal{R}(T) = \mathcal{O} \left(\frac{NK^2 \log T \Delta_{\text{opt}}^{\max}}{(\Delta_{\text{opt}}^{\min})^2} + \frac{\sqrt{NK T \log T} \Delta_{\text{opt}}^{\max}}{(\Delta_{\text{change}}^{\min})^2} \right).$$

2. (N is unknown) Choosing $\delta = \frac{1}{T}$, $p = \sqrt{\frac{K \log T}{T}}$, gives

$$\mathcal{R}(T) = \mathcal{O} \left(\frac{NK^2 \log T \Delta_{\text{opt}}^{\max}}{(\Delta_{\text{opt}}^{\min})^2} + \frac{N \sqrt{KT \log T} \Delta_{\text{opt}}^{\max}}{(\Delta_{\text{change}}^{\min})^2} \right).$$

Remark 4.1. The effect of oracle is reducing the dependency on number of base arms K from $|\mathcal{F}|$ to K^2 during exploration (first term in the regret appeared in Corollary 4.3). In the worst case, $|\mathcal{F}|$ can be exponential with respect to K . Recall that if we use standard MAB algorithms for exploration, the dependency on K is $|\mathcal{F}|$.

Remark 4.2. As T becomes larger, the regret is dominated by the cost of change-point detection, which has similar order compared to the regret bound of piecewise-stationary MAB algorithms. This is reasonable since our setting assumes that we have access to the rewards of the base arms contained in the super arm played by the agent.

Remark 4.3. When T is large, the order of the regret bound is similar to the regret bound of adversarial bandit. But note that regret definition for adversarial bandit and piecewise-stationary bandit is different. For the first case, the regret is evaluated with respect to one fixed arm which is optimal for the whole horizon. But for the second case, the regret is evaluated with respect to point-wise optimal arm, which is much more challenging.

We can use Corollary 4.3 as a guide for parameter tuning. The above corollary indicates that without knowledge of the number of change-points N , we pay a penalty of a factor of \sqrt{N} in the long run.

For the detailed proof of Theorem 4.2, please refer to Appendix A. Here we sketch the proof with some additional lemmas. We start by proving the regret of GLR-CUCB under the stationary scenario.

Lemma 4.4. Under the stationary scenario, i.e. $N = 1$, the α -approximation cumulative regret of GLR-CUCB is upper bounded as:

$$\mathcal{R}(T) \leq \Delta_{\text{opt}}^{\max, 1} T \mathbb{P}(\tau_1 \leq T) + \Delta_{\text{opt}}^{\max, 1} T p + \tilde{C}_1.$$

The first term is due to the possible false alarms of the change-point detection subroutine, the second term is due to the uniform exploration, and the last term is due to the UCB exploration. We upper bound the false alarm probability in Lemma 4.5 as follows.

Lemma 4.5 (False alarm probability in the stationary scenario). Consider the stationary scenario, i.e. $N = 1$, with confidence level $\delta > 0$; we have that

$$\mathbb{P}(\tau_1 \leq T) \leq K \delta.$$

Remark 4.4. By setting $\delta = \frac{1}{T}$, we will have $\mathbb{P}(\tau_1 \leq T) \leq \frac{K}{T}$. Asymptotically, the false alarm probability will go to 0.

In the next lemma, we show the GLR change-point detector is able to detect change in distribution reasonably well with high probability, given all previous change-points were detected reasonably well. The formal statement is as follows.

Lemma 4.6. (Lemma 12 in Besson and Kaufmann (2019)) Define the event $\mathcal{C}^{(i)}$ that all the change-points up to i th one have been detected successfully within a small delay:

$$\mathcal{C}^{(i)} = \{\forall j \leq i, \tau_j \in \{\nu_j + 1, \dots, \nu_j + d_j\}\}. \quad (2)$$

Then, $\mathbb{P}[\tau_i \leq \nu_i | \mathcal{C}^{(i-1)}] \leq K\delta$, and $\mathbb{P}[\tau_i \geq \nu_i + d_i | \mathcal{C}^{(i-1)}] \leq \delta$, where τ_i is the detection time of the i th change-point.

Lemma 4.6 provides an upper bound for the conditional expected detection delay, given the good events $\{\mathcal{C}^{(i)}\}$.

Corollary 4.7 (Bounded conditional expected delay). $\mathbb{E}[\tau_i - \nu_i | \mathcal{C}^{(i)}] \leq d_i$.

Given these lemmas, we can derive the regret upper bound for GLR-CUCB in a recursive manner. Specifically, we prove Theorem 4.2 by recursively decomposing the regret into a collection of good events and bad events. The good events contain all sample paths that GLR-CUCB reinitialize the UCB indices of base arms after all change-points correctly within a small delay. On the other hand, the bad events contain all sample paths where either GLR change detector fails to detect the change in distribution or detects the change with a large delay. The cost incurred given the good events can be upper bounded by Lemma 4.4 and Lemma 4.7. By upper bounding the probabilities of bad events via Lemma 4.5 and Lemma 4.6, the cost incurred given the bad events is analyzable. Detailed proofs are presented in Appendix A.

5 Regret Lower Bound

The lower bound for MAB problems has been studied extensively. Previously, the best available minimax lower bound for piecewise-stationary MAB was $\Omega(\sqrt{T})$ by Garivier and Moulines (2011). Note that piecewise-stationary MAB is a special instance for piecewise-stationary CMAB in which each super arm is a base arm, thus this lower bound still holds for piecewise-stationary CMAB. We derive a tighter lower bound by characterizing the dependency on N and K .

Theorem 5.1. If $K \geq 3$ and $T \geq M_1 N \frac{(K-1)^2}{K}$, then the worst-case regret for any policy is lower bounded by

$$\mathcal{R}(T) \geq M_2 \sqrt{NKT},$$

where $M_1 = 1/\log \frac{4}{3}$, $M_2 = 1/24\sqrt{\log \frac{4}{3}}$.

Proof. (Sketch) The high level idea is to construct a randomized ‘hard’ instance which is appropriate to our setting (Bubeck and Cesa-Bianchi 2012; Besbes, Gur, and Zeevi 2014; Lattimore and Szepesvári 2018), then analyze its regret lower bound which holds for any exploration policy. The construction of this ‘hard’ instance is as follows.

We partition the time horizon into N segments with equal length except for the last segment. In each segment, assume the rewards of all arms are Bernoulli distributions and stay unchanged. At each time step there is an optimal arm with expected reward of $\frac{1}{2} + \epsilon$ and the remaining arms have the same expected reward of $\frac{1}{2}$. The optimal arm will change in two consecutive segments by sampling uniformly at random from the remaining $K - 1$ arms.

We then use Lemma A.1 in Auer et al. (2002) to upper bound the expected number of pulls to any arm being optimal under change of distributions, from the sub-optimal reward distribution to the optimal reward distribution (Bern($\frac{1}{2}$) to Bern($\frac{1}{2} + \epsilon$)). Given the upper bound of expected number of pulls to the optimal arm, we can lower bound the regret for

any exploration policy. By properly tuning ϵ and after some additional steps, we can derive the minimax regret lower bound. The condition $K \geq 3$ comes from the fact that the lower bound needs to be non-trivial, and $T \geq M_1 N \frac{(K-1)^2}{K}$ comes from the tuning of ϵ .

For the detailed proof, please refer to Appendix B. \square

The conditions for this minimax lower bound are mild, since in practice the number of base arms K is usually much larger and we care about the long-term regret, in other words, large T regime.

Our minimax lower bound shows that GLR-CUCB is nearly order-optimal with respect to all parameters. On the other hand, as a byproduct, this bound also indicates that EXP3S (Auer et al. 2002) and MUCB (Cao et al. 2019) are nearly order-optimal for piecewise-stationary MAB, up to poly-logarithm factors. To be more specific, EXP3S and MUCB achieve regret $\mathcal{O}(\sqrt{NKT \log KT})$ and $\mathcal{O}(\sqrt{NKT \log T})$ respectively.

6 Experiments

We compare GLR-CUCB with five baselines from the literature, one variant of GLR-CUCB, and one oracle algorithm. Specifically, DUCB (Kocsis and Szepesvári 2006) and MUCB (Cao et al. 2019) are selected from piecewise-stationary MAB literature; CUCB (Chen, Wang, and Yuan 2013), CTS (Wang and Chen 2018), and Hybrid (Zimmer, Luo, and Wei 2019) are selected from stochastic combinatorial semi-bandit literature. The variant of GLR-CUCB, termed LR-GLR-CUCB, uses different restart strategy. Instead of restarting the estimation of all base arms once a change-point is detected, LR-GLR-CUCB uses local restart strategy (only restarts the estimation of the base arms that are detected to have changes in reward distributions). For the oracle algorithm, termed Oracle-CUCB, we assume the algorithm knows when the optimal super arm changes and restarts CUCB at these change-points. Note that this is stronger than knowing the change-points, since change in distribution does not imply change in optimal super arm. Experiments are conducted on both synthetic and real-world dataset for the m -set bandit problems, which aims to identify the m arms with highest expected reward at each time step. Equivalently, the reward function $r_{\mu_t}(S_t)$ is the summation of the expected rewards of m base arms. Since DUCB and MUCB are originally designed for piecewise-stationary MAB, to adapt them to the piecewise-stationary CMAB setting, we treat every super arm as a single arm when we run these two algorithms. Reward distributions of base arms along time are postponed to Appendix C.1. The details about parameter tuning for all of these algorithms for different experiments are included in Appendix C.2.

6.1 Synthetic Dataset

In this case we design a synthetic piecewise-stationary combinatorial semi-bandit instance as follows:

- Each base arm follows Bernoulli distribution.
- Only one base arm changes its distribution between two consecutive piecewise-stationary segments.

- Every piecewise-stationary segment is of equal length.

We let $T = 5000$, $K = 6$, $m = 2$, and $N = 5$. The average regret of all algorithms are summarized in Figure 1. Note that the optimal super arm does not change for the last




Figure 1: Expected cumulative regret for different algorithms on synthetic dataset.

three piecewise-stationary segments. Observe that the well-tuned GLR change-point detector is insensitive to change with small magnitude, which implicitly avoids unnecessary and costly global restart, since small change is less likely to affect the optimal super arm. Surprisingly, GLR-CUCB and LR-GLR-CUCB perform nearly as well as Oracle-CUCB and significantly better than other algorithms in regrets. In general, algorithms designed for stochastic CMAB outperform algorithms designed for piecewise-stationary MAB. The reason is when the horizon is small, the dimension of the action space dominates the regret, and this effect becomes more obvious when m is larger. Although order-wise, the cost incurred by the change-point detection is much higher than the cost incurred by exploration.

Note that our experiment on this synthetic dataset does not satisfy Assumption 4.1. For example, the gap between the first segment and second segment is 0.6, and we choose $\delta = \frac{20}{T}$ and $p = 0.05\sqrt{(N-1)\log T/T}$ for GLR-CUCB, which means the length of the second segment should be at least 9874. However, the actual length of the second segment is only 1000. Thus our algorithm performs very well compared to other algorithms even if Assumption 4.1 is violated. If Assumption 4.1 is satisfied, GLR-CUCB can only perform better since it is easier to detect the change in distribution.

6.2 Yahoo! Dataset

We adopt the benchmark dataset for the real-world evaluation of bandit algorithms from Yahoo!². This dataset contains user click log for news articles displayed in the Featured Tab of the Today Module (Li et al. 2011). Every base arm corresponds to the click rate of one article. Upon arrival of a user, our goal is to maximize the expected number of clicked articles by presenting m out of K articles to the users.

²Yahoo! Front Page Today Module User Click Log Dataset on <https://webscope.sandbox.yahoo.com>

Yahoo! Experiment 1 ($K = 6$, $m = 2$, $N = 9$). We pre-process the dataset following Cao et al. (2019). To make the experiment nontrivial, we modify the dataset by: 1) the click rate of each base arm is enlarged by 10 times; 2) Reducing the time horizon to $T = 22500$. Results are in Figure 2.




Figure 2: Expected cumulative regret for different algorithms on Yahoo! experiment 1.

Yahoo! experiment 1 is much harder than the synthetic problem, since it is much more non-stationary. Our experiments show GLR-CUCB still significantly outperforms other algorithms and only has a small gap with respect to Oracle-CUCB. Again, Assumption 4.1 does not hold for these two instances, thus we believe it is fair to compare GLR-CUCB with other algorithms. Unexpectedly, LR-GLR-CUCB performs even better than oracle-CUCB, which suggests there is still much to exploit in the piecewise-stationary bandits, since global restart has inferior performance in some cases, especially when the change in distribution is not significant. Additional experiments on Yahoo! dataset can be found in Appendix C.1.

7 Conclusion and Future Work

We have developed the first efficient and general algorithm for piecewise-stationary CMAB, termed GLR-UCB, which extends CUCB (Chen, Wang, and Yuan 2013), by incorporating a GLR change-point detector. We derive the regret upper bound of GLR-CUCB on the order of $\mathcal{O}(\sqrt{NKT} \log T)$, and prove the minimax lower bound for piecewise-stationary MAB and CMAB on the order of $\Omega(\sqrt{NKT})$, which shows our algorithm is nearly order-optimal within poly-logarithm factors. Experimental results show our proposed algorithm outperforms other state-of-the-art algorithms.

Future work includes designing algorithms for piecewise-stationary CMAB with better restart strategy. Another very challenging unsolved problem is whether one can close the gap between the regret upper bound and the minimax regret lower bound. Specifically, develop algorithm which is order-optimal for piecewise-stationary CMAB.

Acknowledgement

The authors are grateful to Julian Zimmert for providing the code of Hybrid algorithm. This work was funded in part by the IBM-Illinois Center for Cognitive Computing Systems Research (C3SR), a research collaboration as part of the IBM AI Horizons Network.

References

- Agrawal, S., and Goyal, N. 2012. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proc. 25th Conf. on Learn. Theory (COLT '12)*, 1–26.
- Auer, P.; Cesa-Bianchi, N.; Freund, Y.; and Schapire, R. E. 2002. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.* 32(1):48–77.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.* 47(2-3):235–256.
- Auer, P.; Gajane, P.; and Ortner, R. 2019. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Proc. 32nd Conf. on Learn. Theory (COLT '19)*, 138–158.
- Ausiello, G.; Crescenzi, P.; and Protasi, M. 1995. Approximate solution of NP optimization problems. *Theor. Comput. Sci.* 150(1):1–55.
- Basseville, M., and Nikiforov, I. V. 1993. *Detection of abrupt changes: theory and application*, volume 104. Prentice Hall Englewood Cliffs.
- Besbes, O.; Gur, Y.; and Zeevi, A. 2014. Stochastic multi-armed-bandit problem with non-stationary rewards. In *Proc. 24th Annu. Conf. Neural Inf. Process. Syst. (NeurIPS '14)*, 199–207.
- Besson, L., and Kaufmann, E. 2019. The generalized likelihood ratio test meets klucb: an improved algorithm for piecewise non-stationary bandits. *arXiv preprint arXiv:1902.01575*.
- Bubeck, S., and Cesa-Bianchi, N. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5(1):1–122.
- Bubeck, S.; Cohen, M.; and Li, Y. 2018. Sparsity, variance and curvature in multi-armed bandits. In *Proc. 29th Int. Conf. Algorithmic Learning Theory (ALT '18)*, 111–127.
- Cao, W.; Li, J.; Tao, Y.; and Li, Z. 2015. On top-k selection in multi-armed bandits and hidden bipartite graphs. In *Proc. 29th Annu. Conf. Neural Inf. Process. Syst. (NeurIPS '15)*, 1036–1044.
- Cao, Y.; Wen, Z.; Kveton, B.; and Xie, Y. 2019. Nearly optimal adaptive procedure with change detection for piecewise-stationary bandit. In *Proc. 22nd Int. Conf. Artif. Intell. Stat. (AISTATS 2019)*, 418–427.
- Cappé, O.; Garivier, A.; Maillard, O.-A.; Munos, R.; Stoltz, G.; et al. 2013. Kullback–Leibler upper confidence bounds for optimal sequential allocation. *Ann. Stat.* 41(3):1516–1541.
- Chen, W.; Wang, Y.; and Yuan, Y. 2013. Combinatorial multi-armed bandit: General framework and applications. In *Proc. 30th Int. Conf. Mach. Learn. (ICML 2013)*, 151–159.
- Combes, R.; Shahi, M. S. T. M.; Proutiere, A.; et al. 2015. Combinatorial bandits revisited. In *Proc. 29th Annu. Conf. Neural Inf. Process. Syst. (NeurIPS '15)*, 2116–2124.
- Garivier, A., and Moulines, E. 2011. On upper-confidence bound policies for switching bandit problems. In *Proc. 22th Int. Conf. Algorithmic Learning Theory (ALT '11)*, 174–188. Springer.
- Hartland, C.; Baskiotis, N.; Gelly, S.; Sebag, M.; and Teytaud, O. 2007. Change point detection and meta-bandits for online learning in dynamic environments. *CAP* 237–250.
- Hassan, U. U., and Curry, E. 2014. A multi-armed bandit approach to online spatial task assignment. In *IEEE 11th Intl Conf Ubiquitous Intell. Comput. (UIC '14)*, 212–219. IEEE.
- Hazan, E., and Kale, S. 2011. Better algorithms for benign bandits. *J. Mach. Learn. Res.* 12:1287–1311.
- Hazan, E. 2016. Introduction to online convex optimization. *Foundations and Trends in Optimization* 2(3-4):157–325.
- Kaufmann, E., and Koolen, W. 2018. Mixture martingales revisited with applications to sequential tests and confidence intervals. *arXiv preprint arXiv:1811.11419*.
- Kocsis, L., and Szepesvári, C. 2006. Discounted UCB. In *2nd PASCAL Challenges Workshop*, volume 2.
- Kveton, B.; Wen, Z.; Ashkan, A.; and Szepesvári, C. 2015. Tight regret bounds for stochastic combinatorial semi-bandits. In *Proc. 18th Int. Conf. Artif. Intell. Stat. (AISTATS 2015)*, 535–543.
- Lai, T. L., and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* 6(1):4–22.
- Lattimore, T., and Szepesvári, C. 2018. Bandit algorithms. *preprint*.
- Li, L.; Chu, W.; Langford, J.; and Wang, X. 2011. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proc. 4th ACM Int. Conf. Web Search Data Min. (WSDM '11)*, 297–306. ACM.
- Li, B.; Chen, T.; and Giannakis, G. B. 2019. Bandit online learning with unknown delays. In *Proc. 22th Int. Conf. Artif. Intell. Stat. (AISTATS 2019)*, 993–1002.
- Li, S.; Karatzoglou, A.; and Gentile, C. 2016. Collaborative filtering bandits. In *Proc. 39th Annu. Int. ACM SIGIR Conf. Res. Dev. Inf. Retrieval (SIGIR '16)*, 539–548. ACM.
- Liu, F.; Lee, J.; and Shroff, N. 2018. A change-detection based framework for piecewise-stationary multi-armed bandit problem. In *Proc. 32nd AAAI Conf. Artif. Intell.*
- Mellor, J., and Shapiro, J. 2013. Thompson sampling in switching environments with bayesian online change detection. In *Proc. 16th Int. Conf. Artif. Intell. Stat. (AISTATS 2013)*, 442–450.
- Pereira, F. S.; Gama, J.; de Amo, S.; and Oliveira, G. M. 2018. On analyzing user preference dynamics with temporal social networks. *Mach. Learn.* 107(11):1745–1773.
- Siegmund, D. 2013. *Sequential analysis: tests and confidence intervals*. Springer Science & Business Media.
- Tang, L.; Rosales, R.; Singh, A.; and Agarwal, D. 2013. Automatic ad format selection via contextual bandits. In *Proc. 22nd ACM Int. Conf. Inf. Knowl. Manage. (CIKM'13)*, 1587–1594. ACM.
- Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25(3/4):285–294.
- Vempaty, A.; Varshney, L. R.; and Varshney, P. K. 2014. Reliable crowdsourcing for multi-class labeling using coding theory. *IEEE J. Sel. Topics Signal Process.* 8(4):667–679.
- Wang, S., and Chen, W. 2018. Thompson sampling for combinatorial semi-bandits. In *Proc. 35th Int. Conf. Mach. Learn. (ICML 2018)*, 5114–5122.
- Wei, L., and Srivatsva, V. 2018. On abruptly-changing and slowly-varying multiarmed bandit problems. In *Proc. Am. Contr. Conf. (ACC 2018)*, 6291–6296. IEEE.
- Yu, J. Y., and Mannor, S. 2009. Piecewise-stationary bandit problems with side observations. In *Proc. 26th Int. Conf. Mach. Learn. (ICML 2009)*, 1177–1184. ACM.
- Zimmert, J.; Luo, H.; and Wei, C.-Y. 2019. Beating stochastic and adversarial semi-bandits optimally and simultaneously. In *Proc. 36th Int. Conf. Mach. Learn. (ICML 2019)*, 7683–7692.