

# Stochastic Online Learning with Probabilistic Graph Feedback

Shuai Li,<sup>1</sup> Wei Chen,<sup>2</sup> Zheng Wen,<sup>3</sup> Kwong-Sak Leung<sup>4</sup>

<sup>1</sup>Shanghai Jiao Tong University, <sup>2</sup>Microsoft Research, <sup>3</sup>DeepMind, <sup>4</sup>The Chinese University of Hong Kong  
<sup>1</sup>shuaili8@sjtu.edu.cn, <sup>2</sup>weic@microsoft.com, <sup>3</sup>zhengwen@google.com, <sup>4</sup>ksleung@cse.cuhk.edu.hk

## Abstract

We consider a problem of stochastic online learning with general probabilistic graph feedback, where each directed edge in the feedback graph has probability  $p_{ij}$ . Two cases are covered. (a) The one-step case, where after playing arm  $i$  the learner observes a sample reward feedback of arm  $j$  with independent probability  $p_{ij}$ . (b) The cascade case where after playing arm  $i$  the learner observes feedback of all arms  $j$  in a probabilistic cascade starting from  $i$  – for each  $(i, j)$  with probability  $p_{ij}$ , if arm  $i$  is played or observed, then a reward sample of arm  $j$  would be observed with independent probability  $p_{ij}$ . Previous works mainly focus on deterministic graphs which corresponds to one-step case with  $p_{ij} \in \{0, 1\}$ , an adversarial sequence of graphs with certain topology guarantees, or a specific type of random graphs. We analyze the asymptotic lower bounds and design algorithms in both cases. The regret upper bounds of the algorithms match the lower bounds with high probability.

## 1 Introduction

Stochastic online learning is a general framework of sequential decision problem. At each time, the learner selects (or plays) an action from a given finite action set, receives some random reward and observes some random feedback. One simplest, though often unrealistic, feedback model is full-information feedback where the learning agent can observe the random rewards of all actions no matter which action is selected. Another popular feedback model is bandit feedback where only the random reward of the selected action is revealed to the learner (Auer, Cesa-Bianchi, and Fischer 2002). Recent studies further generalize them to graph feedback where the feedback model is characterized by a (directed) graph (Mannor and Shamir 2011). Each edge  $(i, j)$  means the learner will observe the random reward of action  $j$  if playing action  $i$ . This problem is motivated by advertisements where the response for a vacation advertisement could provide side-information for a similar vacation place and social networks where the response from a user to a promotion could infer her neighbors to similar offers.

The problem of online learning with graph feedback has been extensively studied in both adversarial (Mannor

and Shamir 2011; Alon et al. 2015a; Kocák et al. 2014; Cohen, Hazan, and Koren 2016; Kocák, Neu, and Valko 2016b) and stochastic settings (Caron et al. 2012; Buccapatnam, Eryilmaz, and Shroff 2014; Tossou, Dimitrakakis, and Dubhashi 2017; Wu, György, and Szepesvári 2015). While many of them assume self-loops on the feedback graphs, some succeed to remove this assumption (Alon et al. 2015a; Wu, György, and Szepesvári 2015) where the reward of the selected action might be invisible. This general setting would fit into the partial monitoring framework (Bartók et al. 2014; Komiyama, Honda, and Nakagawa 2015), but the literature on the latter mainly focus on finite case where the possible outcomes are finite. We also consider general feedback graphs that do not assume self-loops.

Though some studies assume feedback graphs could vary over time or even invisible to the learner before selecting actions (Kocák et al. 2014; Tossou, Dimitrakakis, and Dubhashi 2017), most works focus on deterministic graphs or an adversarial list of graphs with certain topology guarantees. To the best of our knowledge, only a few of them work on probabilistic graphs with (Kocák, Neu, and Valko 2016a; Alon et al. 2017) on adversarial case and (Liu, Buccapatnam, and Shroff 2018) on stochastic case and they only discuss about Erdős-Rényi random graphs (Erdős and Rényi 1960). Recall that an Erdős-Rényi graph with parameter  $p$  is by random sampling the edge of every pair of nodes with probability  $p$  independently.

We consider general probabilistic feedback graphs in both the one-step case and the cascade case. The one-step case is the usual one where the learner observes reward of  $j$  if edge  $(i, j)$  exists in the random graph and  $i$  is selected. The cascade case assumes the learner observes reward of  $j$  if there is a (directed) path from  $i$  to  $j$  in the random graph and  $i$  is selected. The observations of the cascade case, in other words, follow a probabilistic cascading starting from the selected action — for each edge  $(i, j)$  with probability  $p_{ij}$ , if action  $i$  is either played or observed, then with an independent probability  $p_{ij}$  a random reward sample of action  $j$  will be observed. As a motivating example, consider the information propagation in social networks. If selecting a user in a social network causes an information cascade in the social network, one may be able to observe further feedback from

the cascade users.

This paper makes three major contributions.

1. We formalize the setting of stochastic online learning with general probabilistic graph feedback and consider both the one-step and the cascade cases.
2. We derive asymptotic lower bounds for both the one-step and the cascade cases.
3. We design algorithms for both the one-step and the cascade cases and analyze their regrets. Their asymptotic upper regret bounds match the asymptotic lower bounds with high probability.

**Related work** The studies on online learning with graph feedback started from adversarial online learning with side observations where a decision maker can observe rewards of other actions as well as observe the reward of the selected action (Mannor and Shamir 2011). The observation structure can be encoded as a directed graph where there is an edge  $(i, j)$  if the reward of action  $j$  is observed when  $i$  is selected. Their setting assumes that self-loops exist on every node. Alon et al. (2015a) then generalize to arbitrary directed graphs as long as each action is observable by selecting some action. They show the structure of feedback graph controls the inherent difficulty of the learning problem and present a classification over graphs. These works assume the feedback graph is fixed over time and known to the learner. A follow-up (Alon et al. 2015b) extends to time-varying feedback graphs where the graphs are revealed either at the beginning of the round or at the end of the round but assumes good topology properties on the graphs. Kocak et al. (2014) also allow the feedback graph to vary over time and can be revealed to the learner at the end of the round. The results of (Kocák, Neu, and Valko 2016b) depend on the topological properties of the feedback graphs. Cohen et al. (2016) assume the graph is not revealed in both adversarial and stochastic cases. All these works focus on the adversarial case.

Besides (Cohen, Hazan, and Koren 2016), there are also other works on the stochastic case with deterministic feedback graphs. Caron et al. (2012) first study the stochastic case with side observations and design UCB-like algorithms with improved regret bound over the standard UCB without additional feedback. Baccapatnam et al. (2014) derive an asymptotic lower bound and design two algorithms that are near-optimal. Tossou et al. (2017) apply Thompson sampling and allow the feedback graph to be unknown and/or changing. They bound the Bayesian regret in terms of the size of minimum clique covering. Wu et al. (2015) consider general feedback graphs but assume different observation variance from different choices of actions. They provide non-asymptotic problem-dependent regret lower bound and also design algorithms that achieve the problem-dependent lower bound and the minimax lower bounds. They are the first to remove the self-loop assumption in stochastic case.

There are several works on specific Erdős-Rényi random feedback graphs where the feedback graph at each time is randomly generated by Erdős-Rényi model. Kocak et al.

(2016a) consider adversarial case with the unknown generating probability of the feedback graphs. Liu et al. (2018) consider stochastic case and design a randomized policy with Bayesian regret guarantee. Also both of them assume self-observability. An updated version (Alon et al. 2017) of Alon et al. (2015b) extends one result to Erdős-Rényi model in the adversarial case. We consider general probabilistic feedback graphs and provide gap-dependent regret bounds, which are also new in the setting of Erdős-Rényi random feedback graphs.

The setting of graph feedback can be fit into a more general setting of partial monitoring (Rustichini 1999; Cesa-Bianchi and Lugosi 2006) where feedback matrix and reward matrix are given for each pair of the chosen action and the environment. Bartok et al. (2014) make a significant progress on classifying finite adversarial partial monitoring games which is completed by Lattimore and Szepesvari (2019). Komiyama et al. (2015) derive a problem-dependent regret lower bound and design an algorithm with asymptotically optimal regret upper bound in the stochastic case. Most studies on general partial monitoring framework focus on finite case where the number of actions and possible outcomes are finite. The algorithms for general partial monitoring games are not efficient in our case since the feedback matrix might be infinite or exponentially large.

The cascade observation feedback resembles the independent cascade model in the context of influence maximization studies (Kempe, Kleinberg, and Tardos 2003; Chen, Lakshmanan, and Castillo 2013), but the goal is different: influence maximization aims at finding a set of  $k$  seeds that generates the largest expected cascade size, while our goal is to find the best action (arm) utilizing the cascade feedback. Influence maximization has been combined with online learning in several studies (Vaswani et al. 2015; Chen et al. 2016; Wen et al. 2017; Wang and Chen 2017; Saritaç and Tekin 2017), but again their goal is to maximize influence cascade size while using online learning to gradually learn edge probabilities.

## 2 Settings

Our considered problem is characterized by a quadruple  $(V, E, p, \mu)$ , where  $V = [K]$  is a set of  $K$  actions,  $E \subseteq V \times V$  is a set of directed edges between actions,  $p : E \rightarrow (0, 1]$  maps edges to their *triggering probabilities*, and  $\mu = \{\mu_i\}_{i \in V}$  encodes the reward distributions of all actions. The set of all possible reward distributions is denoted as  $\mathcal{C}$ . Without loss of generality, we assume that each distribution candidate is 1-sub-Gaussian. The set of all feasible vectors of reward distributions is denoted as  $\mathcal{S}$ . The (directed) *probabilistic feedback graph* is also denoted as  $G = (V, E, p)$ . We assume that the learner knows  $G$  and the fact that  $\mu_i$ 's have 1-sub-Gaussian tail, but does not know the reward mean  $\theta_i$ 's.

At each time step  $t = 1, 2, \dots$ , the environment first draws a reward vector  $r_t = (r_t(i) : i \in V)$  by independently sampling  $r_t(i) \sim \mu_i$ , and a random graph  $G_t = (V, E_t)$  based on  $G$ . Specifically,  $E_t = \{(i, j) \in E : o_{t ij} = 1\} \subseteq E$ , where  $o_{t ij}$  is an independent Bernoulli random variable with mean  $p_{ij}$ . Simultaneously, the learner adaptively chooses an

action  $i_t \in V$  based on its past observations, without observing  $r_t$  or  $G_t$ . Then, the learner receives an instantaneous reward  $r_t(i_t)$ , and depending on the specific feedback model, it might also observe part of  $r_t$ . In this paper, we consider the following two feedback models:

**One-Step Triggering** The learner will receive feedback  $(j, r_t(j))$  if and only if  $(i_t, j) \in E_t$ .

**Cascade Triggering** The learner will receive feedback  $(j, r_t(j))$  if and only if there is a directed path from  $i_t$  to  $j$  in  $G_t$ .

It is worth pointing out that though the learner receives the reward  $r_t(i_t)$ , however, if  $(i_t, i_t)$  is not in  $G_t$  in the one-step triggering case, or there is no directed circle from  $i_t$  to  $i_t$  in the cascade triggering case,  $(i_t, r_t(i_t))$  is not observed. In other words, the learner might not observe the reward of its chosen action. Also note that existing works with graph feedback (Caron et al. 2012; Buccapatnam, Eryilmaz, and Shroff 2014; Tossou, Dimitrakakis, and Dubhashi 2017; Alon et al. 2015a; Wu, György, and Szepesvári 2015) are special cases of the one-step triggering case discussed above, with  $p_{ij} = 1$  for all  $(i, j) \in E$ . The work (Liu, Buccapatnam, and Shroff 2018) is also a special case of the one-step triggering case but with  $p_{ij}$  having the same value.

We assume the feedback graph is observable, that is each action has the chance to be observed by pulling some action.

**Assumption 1 (observability)** For each action  $j$ , there is an edge  $(i, j) \in E$  for some  $i$ .

Next assumption states each feasible distribution vector is composed of distributions of “same type”. For example, distributions over a bounded interval will not be put together with Gaussian distributions.

**Assumption 2 (same type)** For each  $\mu \in \mathcal{S}$ ,  $\text{KL}(\mu_i, \mu_j)$  is well-defined for any  $i, j \in V$ . For each  $\mu \in \mathcal{C}^V$ , if  $\text{KL}(\mu_i, \mu_j)$  is well-defined for any  $i, j \in V$ , then  $\mu \in \mathcal{S}$ .

The last assumption says the KL divergence of the reward distributions is continuous with respect to the their means.

**Assumption 3 (continuity)** There exists some universal constant  $B > 0$  such that for each  $\mu_i, \mu_j \in \mathcal{C}$  and any  $\epsilon \in (0, 1)$ , there exists  $\mu'_i \in \mathcal{C}$  satisfying  $\text{KL}(\mu_j, \mu'_i)$  is well-defined,  $\theta(\mu_i) + \epsilon \leq \theta(\mu'_i) < \theta(\mu_i) + 2\epsilon$  and  $|\text{KL}(\mu_j, \mu'_i) - \text{KL}(\mu_j, \mu_i)| \leq B\epsilon$ .

The learner’s objective is to maximize its expected cumulative reward, or equivalently, to minimize its expected cumulative regret

$$R_\mu(T; G) = T \max_{i \in V} \theta(\mu_i) - \mathbb{E} \left[ \sum_{t=1}^T \theta(\mu_{i_t}) \right],$$

where the expectation is over the randomness of  $r_t$  and  $G_t$ . Here  $\theta : \mathcal{C} \rightarrow \mathbb{R}$  is the mapping from the distributions to their means.

We will omit  $G$  in the regret expression and write  $\theta(\mu_i)$  as  $\theta_i$  if the context is clear. For simplicity, we assume there is only one best action and  $\theta_1 > \theta_2 \geq \theta_3 \geq \dots \geq \theta_K$ . Denote  $\theta = (\theta_i : i \in V)$ . Let  $\Delta_i(\mu) = \theta_1 - \theta_i$  be the reward

gap between the best action and action  $i$ . Denote  $\Delta(\mu) = (\Delta_i(\mu) : i \in V)$ . We will omit  $\mu$  in the above notations if the context is clear.

Let  $V^{\text{in}}(j) = \{i \in [K] : (i, j) \in E\}$  be the set of incoming neighbors of action  $j$ . Let  $N_i(t)$  be the number of times the learner selects an action  $i$  and  $N(t) = (N_i(t) : i \in V)$  by the end of time  $t$ .

For general  $\mu$ , let  $i_k(\mu)$  be the  $k$ -th best action index for the distributions  $\mu$ , which has the  $k$ -th largest mean. We will write  $i_k$  for simplicity when the context is clear. Then  $\theta_{i_1(\mu)} > \theta_i, \forall i \neq i_1(\mu)$ .

### 3 Asymptotic Lower Bounds

#### 3.1 Lower Bound for One-Step Triggering

Define

$$C(\mu) = \left\{ c \in [0, \infty)^V : \sum_{i \in V^{\text{in}}(1)} p_{i1} c_i \geq \frac{1}{\text{KL}(\mu_2, \mu_1)} \right. \\ \left. \sum_{i \in V^{\text{in}}(j)} p_{ij} c_i \geq \frac{1}{\text{KL}(\mu_j, \mu_1)}, \forall j \neq 1; \right\}. \quad (1)$$

Each element in the set represents an asymptotic pulling “fraction” of arms that can be used to distinguish these arms from the best arm.

Recall that an algorithm is *consistent* if  $R_\mu(T) = o(T^a)$  for any  $a > 0$  and any feasible  $\mu \in \mathcal{S}$ . Then the asymptotic lower bound for any consistent algorithm is provided in the following theorem.

**Theorem 1** For any consistent algorithm, the regret satisfies

$$\liminf_{T \rightarrow \infty} \frac{R_\mu(T)}{\log T} \geq \inf_{c \in C(\mu)} \langle c, \Delta(\mu) \rangle. \quad (2)$$

Note this lower bound can easily recover the lower bound in (Wang and Chen 2017, Theorem 3) where they only consider a special probabilistic graph  $G$ .

**Proof.** Fix any consistent algorithm and any distribution vector  $\mu$ .

For any  $j \neq 1$  and  $n \geq 1$ , by Assumption 3, there exists a  $\mu_j^{(n)} \in \mathcal{C}$  such that  $\theta_1 + \frac{1}{2^n} \leq \theta(\mu_j^{(n)}) < \theta_1 + \frac{1}{2^{n-1}}$  and  $|\text{KL}(\mu_j, \mu_j^{(n)}) - \text{KL}(\mu_j, \mu_1)| \leq \frac{B}{2^n}$ . Define  $\mu^{(n)} = \mu$  by setting  $\mu_i^{(n)} = \mu_i$  for any  $i \neq j$ . Then by Assumption 2,  $\mu^{(n)} \in \mathcal{S}$ .

Let

$$H = \{i_1, \{r_1(j) : (i_1, j) \in E_1\}; \\ i_2, \{r_2(j) : (i_2, j) \in E_2\}; \dots\}$$

be the random variable of all outcomes, which is based on  $\mu$ , the algorithm and the graph realizations. Let  $\mathbb{P}$  and  $\mathbb{P}^{(n)}$  be the probability distribution over all possible realisations of outcomes when the distribution vector is  $\mu$  and  $\mu^{(n)}$  respectively.

By high-dimensional Pinsker's inequality (Lattimore and Szepesvari 2017, Lemma 5),

$$\begin{aligned} & \mathbb{P}[N_1(T) < T/2] + \mathbb{P}^{(n)}[N_1(T) \geq T/2] \\ & \geq \frac{1}{2} \exp\left(-\text{KL}\left(\mathbb{P}, \mathbb{P}^{(n)}\right)\right). \end{aligned}$$

Note that

$$\begin{aligned} \text{KL}\left(\mathbb{P}, \mathbb{P}^{(n)}\right) &= \sum_{i \in V^{\text{in}}(j)} p_{ij} \mathbb{E}[N_i(T)] \text{KL}\left(\mu_j, \mu_j^{(n)}\right) \quad (3) \\ &\leq \text{KL}\left(\mu_j, \mu_j^{(n)}\right) \sum_{i \in V^{\text{in}}(j)} p_{ij} \mathbb{E}[N_i(T)]. \end{aligned}$$

Then

$$\begin{aligned} & \sum_{i \in V^{\text{in}}(j)} p_{ij} \mathbb{E}[N_i(T)] \\ & \geq \frac{1}{\text{KL}\left(\mu_j, \mu_j^{(n)}\right)} \\ & \quad \cdot \log \frac{1/2}{\mathbb{P}[N_1(T) < T/2] + \mathbb{P}^{(n)}[N_1(T) \geq T/2]} \\ & \geq \frac{1}{\text{KL}\left(\mu_j, \mu_j^{(n)}\right)} \\ & \quad \cdot \log \frac{1/2}{R_\mu(T)/(\Delta_2 \cdot T/2) + R_{\mu^{(n)}}(T)/\left(\frac{1}{2^n} \cdot T/2\right)} \\ & = \frac{1}{\text{KL}\left(\mu_j, \mu_j^{(n)}\right)} \log \frac{T/4}{R_\mu(T)/\Delta_2 + R_{\mu^{(n)}}(T)/\frac{1}{2^n}}, \end{aligned}$$

where the second inequality is due to

$$\begin{aligned} R_\mu(T) &\geq \mathbb{P}[N_1(T) < T/2] \Delta_2 \cdot T/2, \\ R_{\mu^{(n)}}(T) &\geq \mathbb{P}[N_1(T) \geq T/2] \left(\theta\left(\mu_j^{(n)}\right) - \theta_1\right) T/2. \end{aligned}$$

Since the algorithm is consistent,  $R_\mu(T) = o(T^a)$  and  $R_{\mu^{(n)}}(T) = o(T^a)$  for any  $a > 0$ , or equivalently

$$\limsup_{T \rightarrow \infty} \frac{\log R_\mu(T)}{\log T} = 0, \quad \limsup_{T \rightarrow \infty} \frac{\log R_{\mu^{(n)}}(T)}{\log T} = 0.$$

Thus

$$\sum_{i \in V^{\text{in}}(j)} p_{ij} \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[N_i(T)]}{\log T} \geq \frac{1}{\text{KL}\left(\mu_j, \mu_j^{(n)}\right)}.$$

Next take  $n \rightarrow \infty$ ,

$$\sum_{i \in V^{\text{in}}(j)} p_{ij} \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[N_i(T)]}{\log T} \geq \frac{1}{\text{KL}(\mu_j, \mu_1)}.$$

For  $j = 1$  and  $n \geq 1$ , take  $\mu^{(n)} = \mu$  except  $\mu_2^{(n)} \neq \mu_2$  with  $\theta_1 + \frac{1}{2^n} \leq \theta\left(\mu_2^{(n)}\right) < \theta_1 + \frac{1}{2^{n-1}}$  and  $\left|\text{KL}\left(\mu_2, \mu_2^{(n)}\right) - \text{KL}(\mu_2, \mu_1)\right| \leq \frac{B}{2^n}$ . Similar result follows

$$\sum_{i \in V^{\text{in}}(1)} p_{i1} \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[N_i(T)]}{\log T} \geq \frac{1}{\text{KL}(\mu_2, \mu_1)}.$$

Thus the vector  $\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[N(T)]}{\log T} \in C(\mu)$ . Recall the regret is  $R_\mu(T) = \sum_{i=1}^K \mathbb{E}[N_i(T)] \Delta_i(\mu)$ . The result follows.  $\square$

### 3.2 Lower Bound for Cascade Triggering

Let  $p'_{ij}$  be the probability that there is a directed path from  $i$  to  $j$  in a random realization of  $G$ . Define

$$C'(\mu) = \left\{ c \in [0, \infty)^V : \sum_i p'_{i1} c_i \geq \frac{1}{\text{KL}(\mu_2, \mu_1)}, \right. \\ \left. \sum_i p'_{ij} c_i \geq \frac{1}{\text{KL}(\mu_j, \mu_1)}, \forall j \neq 1 \right\}.$$

**Theorem 2** For any consistent algorithm, the regret satisfies

$$\liminf_{T \rightarrow \infty} \frac{R_\mu(T)}{\log T} \geq \inf_{c \in C'(\mu)} \langle c, \Delta(\mu) \rangle.$$

This proof is similar to the above one by replacing (3) with the following formula

$$\text{KL}\left(\mathbb{P}, \mathbb{P}^{(n)}\right) = \sum_i p'_{ij} \mathbb{E}[N_i(T)] \text{KL}\left(\mu_j, \mu_j^{(n)}\right).$$

Note that the computation of  $p'_{ij}$  is #P-hard for general graphs (Valiant 1979; Wang, Chen, and Wang 2012). Thus the lower bound is not efficiently computable even when  $\mu$  is known.

## 4 Algorithm and Analysis

In this section, we design algorithms that can match the lower bounds with high probability asymptotically. The lower bounds in the last section are stated in terms of KL-divergence of distributions. Since the KL-divergence of a real distribution and its estimated empirical distribution might be undefined, we assume the KL-divergence of distributions could be represented by their corresponding means and is also continuous in means, which is also a tradition in bandit area. For example, a previous work (Wu, György, and Szepesvári 2015) assumes distributions to be Gaussian to make statement simpler. We will give more discussions in Section 4.4. In the following, we use mean vector  $\theta$  to represent the vector of distributions  $\mu$  for simplicity.

Let  $\hat{\theta}_t$  be the sample-mean estimates of  $\theta$  by the end of time  $t$ . Let  $n_{ij}(t)$  be the number of times that action  $i$  is selected and reward for action  $j$  is observed by the end of time  $t$ . Then  $\mathbb{E}[n_{ij}(t) | N_i(t)] = N_i(t)p_{ij}$ . Let  $m_j(t) = \sum_i n_{ij}(t)$  be the number of observations for action  $j$  by the end of time  $t$ .

### 4.1 One-Step Uniform Case

The uniform case in which all  $p_{ij}$ 's have the same value  $p$  is first considered in this section. When  $E$  contains edges between every pair of actions, this graph reduces to Erdős-Rényi random graph with parameter  $p$ .

Let  $M_j(t) = \sum_{i \in V^{\text{in}}(j)} N_i(t)p$  be the expected number of observations for action  $j$  at the end of time  $t$ . Then  $\mathbb{E}[m_j(t) | M_j(t)] = M_j(t)$ .

---

**Algorithm 1** One-Step Uniform Case

---

```
1: Set  $N^e(0) = 0$  and  $\hat{\theta}_0 = (1, 1, \dots, 1)$ .
2: for  $t = 1, 2, \dots$  do
3:   if  $m_j(t-1) < M_j(t-1)/2$  for some  $j$  then
4:     Play  $i_t \in V^{\text{in}}(j)$ ;
5:      $N^e(t) = N^e(t-1)$ ;
6:   else if  $\frac{N^e(t-1)}{16 \log(t-1)} \in C(\hat{\theta}_{t-1})$  then
7:     Play  $i_t = i_1(\hat{\theta}_{t-1})$ ;
8:      $N^e(t) = N^e(t-1)$ ;
9:   else if  $M_j(t-1) < 2\beta(N^e(t-1))/K$  for some  $j$ 
   then
10:    Play  $i_t \in V^{\text{in}}(j)$ ;
11:     $N^e(t) = N^e(t-1) + 1$ ;
12:   else
13:    Play  $i_t$  such that  $N_i(t-1) < 16 c_i(\hat{\theta}_{t-1}) \log(t-1)$ ;
14:     $N^e(t) = N^e(t-1) + 1$ ;
15:   end if
16: end for
```

---

The pseudocode of the algorithm is provided in Algorithm 1. It starts with the initialization of  $N^e$  and the estimates of  $\theta$  (line 1). Here  $N^e$  is the number of exploration rounds for the learner to know more about unknown  $\theta$  which will be clearer later. At each time  $t$ , if for some  $j$  the real observation times of action  $j$  is less than half the expected observation times (line 3), then the learner selects a parent of  $j$  to try to observe reward of  $j$  once more (line 4) and keeps  $N^e$  unchanged (line 5). Note that  $\mathbb{E}[m_j(t) \mid M_j(t) = m] = m$  and  $m_j(t)$  will concentrate at  $m$  as  $m$  goes to infinity. The condition  $m_j(t) < M_j(t)/2$  means part of the realizations of graph  $G$  is far from the expectation and 2 can be changed to other larger-than-1 constant. This is one of the key differences from deterministic graph feedback (Wu, György, and Szepesvári 2015) where the number of observations is well controlled by just selecting actions. While under the probabilistic graph feedback, there is a gap between the number of real observations and expected number of observations.

When  $m_j(t) \geq M_j(t)/2$  for all  $j$ , then the realizations of  $G$  are good enough and the learner can rely on the quantities of selections to control the accuracy of the estimates. If the selection vector is good enough for current  $\hat{\theta}$  under current accuracy level (line 6), then the learner will exploit the current best action (line 7) and keep  $N^e$  unchanged. Here  $C(\cdot)$  is defined as in (1) and represents the set of good selected “fractions” of actions that are able to identify the reward gaps between actions.

If the current selection vector  $N$  is not good enough, then the learner will first check if  $\hat{\theta}$  is close enough to  $\theta$  (line 9-11) and if yes, will explore according to current  $\hat{\theta}$ . The number  $N^e$  of exploration rounds for the learner to know more about  $\theta$  will increase in this part (line 11&14). The condition of line 9 has an auxiliary function  $\beta : \mathbb{N} \rightarrow [0, \infty)$  to guide the exploration such that  $\hat{\theta}$  will be close to  $\theta$  in the long run. This auxiliary function is also crucial in previous work (Wu, György, and Szepesvári 2015) to control the regret bound in the asymptotic sense. The auxiliary function  $\beta$  can be any

non-decreasing function satisfying  $0 \leq \beta(n) \leq n/2$  and the subadditivity  $\beta(m+n) \leq \beta(m) + \beta(n)$ . If some component of  $\hat{\theta}$  has not been explored enough (line 9), then the learner selects a parent to try to get one more observation (line 10) and increases  $N^e$  (line 11).

When all components of  $\hat{\theta}$  are close to  $\theta$ , the learner selects an action according to the current  $\hat{\theta}$  with minimal cost on the regret instructed by the asymptotic lower bound (2). Here  $c_i(\theta')$  denotes any optimal solution of the linear programming problem that minimizes  $\langle c, \theta' \rangle$  among all  $c \in C(\theta')$ . Since  $\hat{\theta}$  is close enough to  $\theta$  under current accuracy level, the vector  $c_i(\hat{\theta}_{t-1})$  is close enough to  $c_i(\theta)$  (which is part of the proof for the following theorem). There must be at least an  $i$  such that  $N_i(t-1) < 16 c_i(\hat{\theta}_{t-1}) \log(t-1)$  or else the condition of line 6 holds.

The regret bound for the algorithm is stated as follows.

**Theorem 3** *The regret of Algorithm 1 for one-step uniform case satisfies for any  $\epsilon > 0$ ,*

$$\begin{aligned} R_\theta(T) \leq & 4 \log(T) \sum_{i=1}^K c_i(\theta, \epsilon) \Delta_i(\theta) \\ & + 10 \log(KT^2) \sum_{i=1}^K \frac{\Delta_i(\theta)}{p} + 4 \sum_{s=0}^T \exp\left(-\frac{\beta(s)\epsilon^2}{2K}\right) \\ & + 2\beta\left(4 \sum_{i=1}^K c_i(\theta, \epsilon) \log(T) + K\right) + 15K, \end{aligned} \quad (4)$$

where  $c_i(\theta, \epsilon) = \sup\{c_i(\theta') : |\theta'_j - \theta_j| \leq \epsilon, \forall j \in [K]\}$ .

Assume  $\beta(n) = o(n)$  and  $\sum_{s=0}^{\infty} \exp\left(-\frac{\beta(s)\epsilon^2}{2K}\right) < \infty$  for any  $\epsilon > 0$ . Then for any  $\theta$  such that  $c(\theta)$  is unique,

$$\limsup_{T \rightarrow \infty} R_\theta(T) / \log(T) \leq 4 \inf_{c \in C(\theta)} \langle c, \Delta(\theta) \rangle \quad (5)$$

holds with probability at least  $1 - \delta$  for any  $\delta > 0$ .

Note that any  $\beta(n) = an^b$  with  $a \in (0, \frac{1}{2}]$ ,  $b \in (0, 1)$  meets the requirements. The proof is by bounding the forced exploration (line 9-11), the exploration by LP solutions (line 13-14) and the exploitation (line 6-8). The main difference with previous works is to bound the difference of realized random graphs and the expected graph (line 3-5). The detailed proof is provided in (Li et al. 2019).

## 4.2 One-Step General Case

In the general case where  $p_{ij}$  can be different,  $M_j(t) = \sum_{i \in V^{\text{in}}(j)} N_i(t) p_{ij}$ . The algorithm follows as in Algorithm 1 by only replacing line 4 with

(4') Play  $i_t \in \text{argmax}_{i \in V^{\text{in}}(j)} p_{ij}$ .

Let

$$V^e = \left\{ i \in [K] : i \in \text{argmax}_{i' \in V^{\text{in}}(j)} p_{i'j} \text{ for some } j \right\} \quad (6)$$

be the set of exploration nodes that have the largest live probability among all incoming edges to some  $j$ . Let

$$p_i^e = \min \left\{ p_{ij} : i \in \text{argmax}_{i' \in V^{\text{in}}(j)} p_{i'j} \text{ for some } j \right\} \quad (7)$$

be the minimal exploration probability for any  $i \in V^e$ . With a modified proof to the uniform case, the theoretical guarantee for the general case follows.

**Theorem 4** *The regret of the modified Algorithm 1' for one-step general case satisfies for any  $\epsilon > 0$ ,*

$$\begin{aligned}
R_\theta(T) &\leq 4 \log(T) \sum_{i=1}^K c_i(\theta, \epsilon) \Delta_i(\theta) \\
&\quad + 10 \log(KT^2) \sum_{i \in V^e} \frac{\Delta_i(\theta)}{p_i^e} \\
&\quad + 4 \sum_{s=0}^T \exp\left(-\frac{\beta(s)\epsilon^2}{2K}\right) \\
&\quad + 2\beta \left( 4 \sum_{i=1}^K c_i(\theta, \epsilon) \log(T) \right) + 15K.
\end{aligned} \tag{8}$$

Assume  $\beta(n) = o(n)$  and  $\sum_{s=0}^{\infty} \exp\left(-\frac{\beta(s)\epsilon^2}{2K}\right) < \infty$  for any  $\epsilon > 0$ . Then for any  $\theta$  such that  $c(\theta)$  is unique,

$$\limsup_{T \rightarrow \infty} R_\theta(T) / \log(T) \leq 4 \inf_{c \in C(\theta)} \langle c, \Delta(\theta) \rangle \tag{9}$$

holds with probability at least  $1 - \delta$  for any  $\delta > 0$ .

### 4.3 Cascade Case

---

#### Algorithm 2 Cascade Case

---

```

1: Set  $N^e(0) = 0$  and  $\hat{\theta}_0 = (1, 1, \dots, 1)$ .  $\eta : \mathbb{N}_+ \rightarrow [0, 1)$ .
2: for  $t = 1, 2, \dots$  do
3:   if  $m_j(t-1) < M_j'(t-1) / 2$  for some  $j$  then
4:     Play  $i_t = i$  if  $(P_t)_{ij} \geq \frac{1}{2} \max_{i'} (P_t)_{i'j}$ ;
5:      $N^e(t) = N^e(t-1)$ ;
6:   else if  $\frac{N(t-1)}{16 \log(t-1)} \in S_{\text{feas}}(\text{LP}_t)$  then
7:     Play  $i_t = i_1(\hat{\theta}_{t-1})$ ;
8:      $N^e(t) = N^e(t-1)$ ;
9:   else if  $M_j'(t-1) < 2\beta(N^e(t-1))/K$  for some  $j$  then
10:    Play  $i_t \in V^{\text{in}}(j)$ ;
11:     $N^e(t) = N^e(t-1) + 1$ ;
12:   else
13:    Play  $i_t = i$  such that  $N_i(t-1) < 16 c_{t,i} \log(t-1)$ 
    where  $c_t \in S_{\text{opt}}(\text{LP}_t)$ ;
14:     $N^e(t) = N^e(t-1) + 1$ ;
15:   end if
16: end for

```

---

For the deterministic graphs, there is no essential difference between one-step case and cascade case — the cascade case on a deterministic graph would be equivalent to constructing a new graph where an edge exists if and only if there is a path on the original graph. For a probabilistic graph, one might try a similar solution for the cascade case by constructing a new graph  $G'$  where the probability of an edge  $(i, j)$  is just the probability  $p'_{ij}$  of  $i$  connecting to  $j$  in a random realization of the original graph. However

the computation of  $p'_{ij}$  is #P-hard for general graphs, and thus the accurate graph  $G'$  is unattainable, though it can be approximated within any accuracy by Monte Carlo simulations. Therefore, during the running of the algorithm, a reasonable approximation of  $G'$  is needed.

Define  $V^{e'}$  and  $p_i^{e'}$  similarly with (6) and (7) by replacing  $p_{ij}$  with  $p'_{ij}$ . Since the computation of  $p'_{ij}$  is #P-hard, we define an estimated version of  $V^{e'}$  and  $p_i^{e'}$  respectively:

$$\begin{aligned}
\hat{V}^e &= \left\{ i \in [K] : p'_{ij} \geq \frac{1}{2} \max_{i'} p'_{i'j} \text{ for some } j \right\} \\
\hat{p}_i^e &= \min \left\{ p'_{ij} : p'_{ij} \geq \frac{1}{2} \max_{i'} p'_{i'j} \text{ for some } j \right\}
\end{aligned}$$

for any  $i \in \hat{V}^e$ . Then  $\hat{p}_i^e \geq p_i^{e'}/2$  for some  $i'$ .

To overcome the stated challenge, we need an auxiliary functions  $\eta : \mathbb{N}_+ \rightarrow [0, 1)$  to set up the tolerance of the approximation. At each time  $t$ , the path from  $i$  to  $j$  with probability  $p'_{ij} \leq \eta(t)$  can be treated as nonexistent (with probability 0) and the estimation of  $p'_{ij}$  has noise at most  $\eta(t)/2$  if the real value  $p'_{ij} > \eta(t)$ . Any non-increasing function with limit 0 can be chosen as  $\eta$ . The choice of  $\eta$  is to control the complexity of the graph with only focusing the path of a reasonable length.

Let  $\text{LP}(\theta', \eta)$  be the following linear programming problem

$$\begin{aligned}
&\min \langle \Delta(\theta'), c \rangle \\
&\text{over all } c \in \mathbb{R}^K \text{ satisfying } P^\top c \geq b(\theta') \text{ and } c \geq 0
\end{aligned} \tag{10}$$

where  $P \in [0, 1]^{K \times K}$  satisfies  $P_{ij} = 0$  if  $p'_{ij} \leq \eta$  and  $|P_{ij} - p'_{ij}| \leq \eta/2$  if  $p'_{ij} > \eta$  and  $b_i(\theta') = \frac{1}{\Delta_i^2(\theta')}$  for  $i \neq i_1(\theta')$  and  $b_{i_1(\theta')}(\theta') = \frac{1}{\Delta_{i_2(\theta')}^2(\theta')}$ .

With the approximation  $G_t$  and the estimated value for reward vector  $\hat{\theta}_{t-1}$ , the linear programming problem considered in time  $t$  is  $\text{LP}_t = \text{LP}(\hat{\theta}_{t-1}, \eta(t))$  and the corresponding  $P$  in (10) is denoted as  $P_t$ . Then the algorithm runs with  $\text{LP}_t$  accordingly. The complete pseudocode is presented in Algorithm 2. In particular, the examination on the realization is performed on approximated graph  $G_t$  with probability matrix  $P_t$  (line 3). The exploitation condition is on the  $\text{LP}_t$  (line 6). Here  $S_{\text{feas}}(\text{LP}_t)$  is the feasible solution set of the linear programming problem  $\text{LP}_t$  which is the set of all  $c \in \mathbb{R}^K$  satisfying  $P_t^\top c \geq b, c \geq 0$ . The exploration when all components of estimated  $\hat{\theta}$  are accurate enough with minimal cost instructed by linear programming solutions is also related to  $\text{LP}_t$  (line 13). Here  $S_{\text{opt}}(\text{LP}_t)$  is the optimal solution set of  $\text{LP}_t$ .

Also  $M_j'(t) = \sum_i N_i(t)(P_t)_{ij}$  is changed accordingly.

The regret of the Algorithm 2 is upper bounded in the following theorem.

**Theorem 5** *The regret of the Algorithm 2 for cascade case*

satisfies for any  $\epsilon > 0$ ,

$$\begin{aligned}
R_\theta(T) &\leq 4 \sum_{i=1}^K \Delta_i(\theta) \max_{t \in [T]} \{c_i(\theta, \epsilon, \eta(t)) \log(t)\} \\
&\quad + 10 \log(KT^2) \sum_{i \in \hat{V}^\epsilon} \frac{\Delta_i(\theta)}{\hat{p}_i^\epsilon} \\
&\quad + 2\beta \left( 4 \sum_{i=1}^K \max_{t \in [T]} \{c_i(\theta, \epsilon, \eta(t)) \log(t)\} + K \right) \\
&\quad + 4 \sum_{s=0}^T \exp\left(-\frac{\beta(s)\epsilon^2}{2K}\right) + 15K, \tag{11}
\end{aligned}$$

where

$$\begin{aligned}
c_i(\theta, \epsilon, \eta) &= \\
&\sup \{c_i : c \in S_{\text{opt}}(\text{LP}(\theta', \eta)) \text{ and } |\theta'_j - \theta_j| \leq \epsilon, \forall j \in [K]\}.
\end{aligned}$$

Assume  $\beta(n) = o(n)$  and  $\sum_{s=0}^\infty \exp\left(-\frac{\beta(s)\epsilon^2}{2K}\right) < \infty$  for any  $\epsilon > 0$ . Then for any  $\theta$  such that  $c(\theta)$  is unique,

$$\limsup_{T \rightarrow \infty} R_\theta(T) / \log(T) \leq 4 \inf_{c \in C'(\theta)} \langle c, \Delta(\theta) \rangle \tag{12}$$

holds with probability at least  $1 - \delta$  for any  $\delta > 0$ .

The result depends on the robustness of the linear programming problems. The  $P$  matrix in the LP problem (10) is noisy, which is much different from one-step case and the case of deterministic graphs where the noise is only on  $\theta'$ . See discussions in the next section. The full proof is put in (Li et al. 2019).

#### 4.4 Discussions

The assumptions on the reward distributions are mainly used to ensure that the learning algorithms are able to differentiate them in the worst case (or the regret lower bound). The Gaussian distribution, Bernoulli distribution and common continuous random distribution on a common bounded interval like Beta distribution all satisfy the requirements.

The assumption that the reward distribution can be represented by its mean is commonly adopted in bandit literature. Since there is always gap between a continuous distribution with its discrete empirical estimate and the reward only cares about the mean, previous works hardly choose to estimate the real distribution but mainly choose to estimate the mean. The real mean can be well analysed by constructing a confidence interval around the sample mean.

The term  $O\left(\log(T) \sum_{i=1}^K \frac{\Delta_i(\theta)}{p}\right)$  in the regret bound for one-step uniform case (same for other two cases) is due to the gap between the realizations and the expectations of the probabilistic graphs. Such a term can be removed in the asymptotic sense with high probability based on a different proof. With high probability, the connection between the realizations and the expectations of the probabilistic graphs can be guaranteed for large enough  $T$ , so the realizations of the probabilistic graphs are good enough and no regret would be caused from line 3 - 5 of Algorithm 1 for large

enough  $T$ . If we remove the high probability condition, such a  $1/p$  term remains in the asymptotic sense. Such  $1/p$  term also appears in the regret  $O(\sqrt{T/p})$  of (Kocák, Neu, and Valko 2016a) on Erdős-Rényi random graphs in adversarial setting, as compared with adversarial case on deterministic graphs. It is not clear whether this  $1/p$  term represents hindsight difficulty between the probabilistic graphs and deterministic graphs. This would be an interesting future direction.

The terms  $\{p_i^e : i \in [K]\}$  in the one-step general case describes the minimal exploration probabilities to observe every action. For each  $i \in [K]$ ,  $p_i^e = \max_{i'} p_{i'j}$  for some  $j$ , that is  $p_{ij}$  is the largest live probabilities among all incoming edges for some  $j$ . These terms represent the problem complexities for the underlying probabilistic graph. When all  $p_{ij}$  are equal to  $p$ ,  $p_i^e = p$ .

The term  $p_i^{e'}$  in the cascade case is usually larger than  $p_i^e$  since it takes the same operations on the connection probabilities of incoming paths which are larger than live probabilities of incoming edges. The term  $\hat{p}_i^e$  is an estimation satisfying  $\hat{p}_i^e \geq p_i^{e'}/2$  for some  $i'$ .

Next we discuss the difference in proof of the cascade case. If the noise of the linear programming problems is on the  $b$  vector in (10), then by the standard results in statistics (Dontchev and Rockafellar 2009, §3C.5), the resulting optimal solution sets are Lipschitz continuous. The property of Lipschitz continuity is essential since actions are selected according to the optimal solution of a noisy LP problem (line 13) and we need to guarantee this kind of selections is safe. The noise on  $\Delta$  vector in (10) is also easy to deal with by considering the dual problem. However, it is much different if the noise is on the  $P$  matrix. For example, consider the LP problem that minimizes  $x$  over all  $ax \geq 1$  and  $x \geq 0$  with parameter  $a > 0$ . The optimal solution  $x^* = 1/a$  is not Lipschitz continuous with respect to  $a$ . So the standard statistical tools could not apply here. We derive a novel property of the Lipschitz continuity when there is noise on  $P$  for our specific  $P$  matrix.

Last we would like to stress that our regret bounds are the first gap-dependent bounds even under the one-step uniform case, which contains the simple case of Erdős-Rényi random graph feedback. The previous works on Erdős-Rényi random graphs study gap free bound, no matter in the stochastic setting or the adversarial setting.

## 5 Conclusion and Future Work

We are the first to formalize the setting of stochastic online learning with probabilistic feedback graph. We derive asymptotic lower bounds for both one-step and cascade cases. The regret bounds of our designed algorithms match the lower bounds with high probability.

This framework is new and we only provide asymptotic lower bounds and finite-time problem-dependent upper bounds. Finite-time lower bounds and minimax upper/lower bounds are all interesting future directions. Deriving Bayesian regret bounds is also an interesting topic.

## References

- Alon, N.; Cesa-Bianchi, N.; Dekel, O.; and Koren, T. 2015a. Online learning with feedback graphs: Beyond bandits. In *Conference on Learning Theory*, 23–35.
- Alon, N.; Cesa-Bianchi, N.; Dekel, O.; and Koren, T. 2015b. Online learning with feedback graphs: Beyond bandits. *arXiv preprint arXiv:1502.07617*.
- Alon, N.; Cesa-Bianchi, N.; Gentile, C.; Mannor, S.; Mansour, Y.; and Shamir, O. 2017. Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing* 46(6):1785–1826.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47(2-3):235–256.
- Bartók, G.; Foster, D. P.; Pál, D.; Rakhlin, A.; and Szepesvári, C. 2014. Partial monitoring—classification, regret bounds, and algorithms. *Mathematics of Operations Research* 39(4):967–997.
- Buccapatnam, S.; Eryilmaz, A.; and Shroff, N. B. 2014. Stochastic bandits with side observations on networks. *ACM SIGMETRICS Performance Evaluation Review* 42(1):289–300.
- Caron, S.; Kveton, B.; Lelarge, M.; and Bhagat, S. 2012. Leveraging side observations in stochastic bandits. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI)*, 142–151. AUAI Press.
- Cesa-Bianchi, N., and Lugosi, G. 2006. *Prediction, learning, and games*. Cambridge university press.
- Chen, W.; Wang, Y.; Yuan, Y.; and Wang, Q. 2016. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *The Journal of Machine Learning Research (JMLR)* 17(1):1746–1778.
- Chen, W.; Lakshmanan, L. V. S.; and Castillo, C. 2013. *Information and Influence Propagation in Social Networks*. Morgan & Claypool Publishers.
- Cohen, A.; Hazan, T.; and Koren, T. 2016. Online learning with feedback graphs without the graphs. In *International Conference on Machine Learning (ICML)*, 811–819.
- Dontchev, A. L., and Rockafellar, R. T. 2009. Implicit functions and solution mappings. *Springer Monogr. Math*.
- Erdős, P., and Rényi, A. 1960. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences* 5:17–61.
- Kempe, D.; Kleinberg, J. M.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 137–146.
- Kocák, T.; Neu, G.; Valko, M.; and Munos, R. 2014. Efficient learning by implicit exploration in bandit problems with side observations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 613–621.
- Kocák, T.; Neu, G.; and Valko, M. 2016a. Online learning with erdős-rényi side-observation graphs. In *Uncertainty in Artificial Intelligence (UAI)*.
- Kocák, T.; Neu, G.; and Valko, M. 2016b. Online learning with noisy side observations. In *Artificial Intelligence and Statistics (AISTATS)*, 1186–1194.
- Komiyama, J.; Honda, J.; and Nakagawa, H. 2015. Regret lower bound and optimal algorithm in finite stochastic partial monitoring. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1792–1800.
- Lattimore, T., and Szepesvári, C. 2017. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics (AISTATS)*, 728–737.
- Lattimore, T., and Szepesvári, C. 2019. Cleaning up the neighborhood: A full classification for adversarial partial monitoring. In *Algorithmic Learning Theory (ALT)*, 529–556.
- Li, S.; Chen, W.; Wen, Z.; and Leung, K.-S. 2019. Stochastic online learning with probabilistic graph feedback. *arXiv preprint arXiv:1903.01083*.
- Liu, F.; Buccapatnam, S.; and Shroff, N. 2018. Information directed sampling for stochastic bandits with graph feedback. In *Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*.
- Mannor, S., and Shamir, O. 2011. From bandits to experts: On the value of side-observations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 684–692.
- Rustichini, A. 1999. Minimizing regret: The general case. *Games and Economic Behavior* 29(1-2):224–243.
- Sarıtaç, A. Ö., and Tekin, C. 2017. Combinatorial multi-armed bandit problem with probabilistically triggered arms: A case with bounded regret. In *2017 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 111–115. IEEE.
- Tossou, A. C.; Dimitrakakis, C.; and Dubhashi, D. 2017. Thompson sampling for stochastic bandits with graph feedback. In *Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*.
- Valiant, L. G. 1979. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8(3):410–421.
- Vaswani, S.; Lakshmanan, L.; Schmidt, M.; et al. 2015. Influence maximization with bandits. *arXiv preprint arXiv:1503.00024*.
- Wang, Q., and Chen, W. 2017. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1161–1171.
- Wang, C.; Chen, W.; and Wang, Y. 2012. Scalable influence maximization for independent cascade model in large-scale social networks. *DMKD*.
- Wen, Z.; Kveton, B.; Valko, M.; and Vaswani, S. 2017. Online influence maximization under independent cascade model with semi-bandit feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*, 3025–3035.
- Wu, Y.; György, A.; and Szepesvári, C. 2015. Online learning with gaussian payoffs and side observations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1360–1368.