# Bandit Algorithms for Precision Medicine

Yangyi Lu\*, Ziping Xu\*, Ambuj Tewari {yylu,zipingxu,tewaria}@umich.edu

> Department of Statistics University of Michigan

> > October 2, 2024

## 1 Introduction

The Oxford English Dictionary defines precision medicine as "medical care designed to optimize efficiency or therapeutic benefit for particular groups of patients, especially by using genetic or molecular profiling." It is not an entirely new idea: physicians from ancient times have recognized that medical treatment needs to consider individual variations in patient characteristics (Konstantinidou et al., 2017). However, the modern precision medicine movement has been enabled by a confluence of events: scientific advances in fields such as genetics and pharmacology, technological advances in mobile devices and wearable sensors, and methodological advances in computing and data sciences.

This chapter is about bandit algorithms: an area of data science of special relevance to precision medicine. With their roots in the seminal work of Bellman, Robbins, Lai and others, bandit algorithms have come to occupy a central place in modern data science (see the book by Lattimore and Szepesvári (2020) for an up-to-date treatment). Bandit algorithms can be used in any situation where treatment decisions need to be made to optimize some health outcome. Since precision medicine focuses on the use of patient characteristics to guide treatment, *contextual* bandit algorithms are especially useful since they are designed to take such information into account.

The role of bandit algorithms in areas of precision medicine such as mobile health and digital phenotyping has been reviewed before (Tewari and Murphy, 2017; Rabbi et al., 2019). Since these reviews were published, bandit algorithms have continued to find uses in mobile health and several new topics have emerged in the research on bandit algorithms. This chapter is written for quantitative researchers in fields such as statistics, machine learning, and operations research who might be interested in knowing more about the algorithmic and mathematical details of bandit algorithms that have been used in mobile health.

We have organized this chapter to meet two goals. First, we want to provide a concise exposition of basic topics in bandit algorithms. Section 2 will help the reader become familiar with basic problem setups and algorithms that appear frequently in applied work in precision medicine and mobile health (see, for example, Paredes et al. (2014); Piette et al. (2015); Rabbi et al. (2015); Piette et al. (2016); Yom-Tov et al. (2017); Rindtorff et al. (2019); Forman et al. (2019); Liao et al. (2020); Ameko et al. (2020); Aguilera et al. (2020); Tomkins et al. (2021)). Second, we want to highlight a few advanced topics that are important for mobile health and precision medicine applications but whose full potential remains to be realized. Section 3 will provide the reader with helpful entry points into the bandit literature on non-stationarity, robustness to corrupted rewards, satisfying additional constraints, algorithmic fairness, and causality.

<sup>\*</sup>These authors made equal contributions to the writing of this chapter.

## 2 Basic Topics

In this section, we begin by introducing the most simple of all bandit problems: the multi-armed bandit. Then we discuss a more advanced variant called contextual bandit that is especially suitable for precision medicine applications. The last topic we discuss in this section is offline learning which deals with algorithms that can use already collected data. The offline learning setting is to be contrasted with the online learning setting where the bandit algorithm has control over the data it collects.

## 2.1 Multi-armed Bandit

In recent years, the multi-armed bandit (MAB) framework has attracted a lot of attention in many application areas such as healthcare, marketing, and recommendation systems. MAB is a simple model that describes the interaction between an  $agent^1$  and an *environment*. At every time step, the agent makes a choice from an action<sup>2</sup> set and receives a reward. The agent may have different goals, such as maximizing the (discounted) cumulative reward within a time horizon, identifying the best arm, or competing with the arm with the best risk-return trade-off etc. In this section, we focus on maximizing the cumulative rewards for simplicity. An important observation is that the agent needs to balance between *exploration* and *exploitation* to achieve its goal of receiving high cumulative reward. That is, both under-explored arms as well as tested-and-tried arms with high rewards should be selected often but for different reasons: the former have the potential to achieve high rewards and the latter are already confirmed to be good based on the past experience.

To formally define the bandit framework, we start with introducing some notation. Suppose the agent interacts with the environment for T time steps, where T is called the horizon. In each round  $t \in [T]$ , the learner chooses an action  $A_t$  from the action set  $\mathcal{A}$  and receives a corresponding reward  $R_t \in \mathbb{R}$ . We denote the cardinality of  $\mathcal{A}$  by K. The choice of  $A_t$  depends on the action/reward history up to time t-1:  $H_{t-1} = (A_1, R_1, \ldots, A_{t-1}, R_{t-1})$ . A policy  $\pi_t$  is defined as a mapping from the history up to time t-1 to the actions. For short, we use  $\pi$  as the sequence of policies  $(\pi_0, \ldots, \pi_{T-1})$ .

In a healthcare setting, the fundamental pattern that often occurs is the following. Of course, this simple pattern fails to capture the full complexity of decision making in healthcare, but it is a reasonable starting point, especially for theoretical analysis.

Algorithm 1 Bandit	Framework in	h Healthcare
--------------------	--------------	--------------

<b>Input:</b> Available treatment options, treatment period length $T$ .
for $t = 1, \ldots, T$ do
A treatment option (action) is selected and delivered to the patient.
Patient's health outcome (reward) following the treatment is recorded
end for

In the remainder of this section, we will review bandit algorithms that learn good decision policies over time. We focus on the two key settings: *stochastic bandit* and *adversarial bandit*. In both settings, the algorithms aim at minimizing their regret, which measures the difference between the maximal reward one can get, and the reward obtained by the algorithm. We will formally define regret in each setting.

#### 2.1.1 Stochastic Multi-armed Bandit

A stochastic bandit is a set of distributions  $\nu = (P_a : a \in \mathcal{A})$  and we define the environment class  $\mathcal{E}$  as a set of such distributions

$$\mathcal{E} = \{ \nu = (P_a : a \in \mathcal{A}) : P_a \in \mathcal{M}_a \text{ for all } a \in \mathcal{A} \},\$$

<sup>&</sup>lt;sup>1</sup>Also referred to as a *learner*, *statistician*, or *decision maker*.

 $<sup>^{2}</sup>$ Since the historical roots of probability theory lie in gambling and casinos, it is not surprising that the MAB terminology comes from imagining a slot machine in a casino. A slot machine is also called a "one-armed bandit" as it robs you of your money. Therefore, we will use *actions* and *arms* interchangeably.

where for each a,  $\mathcal{M}_a$  is a set of distributions. For unstructured bandits, playing one action cannot help the agent deduce anything about other actions. Environment classes that are not unstructured are called structured, such as linear bandits (Abbasi-Yadkori et al., 2011), low-rank bandits (Lu et al., 2021b) and combinatorial bandits (Cesa-Bianchi and Lugosi, 2012) etc. Throughout this chapter, we assume all bandit instances are  $\mathcal{E}_{SG}^{K}(1)$ , which means the reward distribution for all K arms is 1-subgaussian.

**Definition 1** (Subgaussianity). A random variable X is  $\sigma$ -subgaussian if for all  $\lambda \in \mathbb{R}$ ,  $\mathbb{E}\left[e^{\lambda X}\right] \leq e^{\lambda^2 \sigma^2/2}$ .

It is not hard to see from the definition that many well-known distributions are subgaussian, e.g., any bounded-domain distribution, Bernoulli distribution and Gaussian distributions. Intuitively, a subgaussian distribution has tails no heavier than a Gaussian distribution. Many nice concentration inequalities have been developed for subgaussian variables and are widely used in the proofs of bandit algorithms.

In the process of interactions, once the agent performs action  $a_t$  following a particular policy, the environment samples a reward  $R_t$  from the distribution  $P_{a_t}$ . The combination of an environment and agent policy induces a probability measure on the sequence of outcomes  $a_1, R_1, \ldots, a_T, R_T$ . A standard stochastic MAB protocol is following. At every time step  $t = 1, \ldots, T$ , the learning agent

- 1. picks an action  $a_t \in \mathcal{A}$  following policy  $\pi_{t-1}$ ,
- 2. receives reward  $R_t \sim P_{a_t}$ ,
- 3. updates its policy to  $\pi_t$ .

We note that  $P_{a_t}$  is the conditional reward distribution of  $R_t$  given  $\{H_{t-1}, a_t\}$  and  $\pi_t$  is a function from  $H_{t-1}$  to  $\mathcal{A}$ . The expected reward of action a is defined by  $\mu_a(\nu) \stackrel{\text{def}}{=} \mathbb{E}_{\nu}[R|a] = \int_{-\infty}^{\infty} r dP_a(r)$ , where R is used as the reward variable. Then the maximum expected reward and the optimal action are given by

$$\mu^*(\nu) = \max_{a \in \mathcal{A}} \mu_a(\nu) \text{ and } a^*(\nu) \in \operatorname{argmax}_{a \in \mathcal{A}} \mu_a(\nu).$$

According to above definition, more than one optimal actions can exist and the optimal policy is to select an optimal action at every round. For actions whose expected rewards are less than optimal actions, we call them sub-optimal actions and define the reward gap between action a and  $a^*(\nu)$  by  $\Delta_a := \mu^*(\nu) - \mu_a(\nu)$ .

As mentioned earlier, the learner's goal is to maximize the cumulative reward  $S_T = \sum_{t=1}^T R_t$ . We now define a performance metric called *regret* which is the difference between the expected reward that  $\pi^*(\nu)$  can obtain and  $\mathbb{E}[S_T]$ . Minimizing the regret is equivalent to maximizing the reward. The reason why we do not directly optimize  $S_T$  is that the cumulative rewards depends on the environment and it is hard to tell whether a policy is good or not by merely looking at the cumulative rewards unless it is compared to a good policy. So we define the problem-dependent regret of a policy  $\pi$  on bandit instance  $\nu$  by

$$\operatorname{Reg}_{T}(\pi,\nu) = T\mu^{*}(\nu) - \mathbb{E}[S_{T}],$$

where the expectation is taken over actions and rewards up to time T. The worst-case regret of a policy  $\pi$  is defined by

$$\operatorname{Reg}_T(\pi) = \sup_{\nu \in \mathcal{E}} \operatorname{Reg}_T(\pi, \nu).$$

We will drop  $\pi$  and  $\nu$  from the regret when they are clear from the context.

**Remark about Pure Exploration.** Throughout this chapter, we focus on minimizing regret by balancing exploration and exploitation. We also want to point out that in a different setting, the exploration cost may not be a concern and the agent just wants to output a final recommendation for the best arm after an exploration phase. Problems of this type are called pure exploration problems. In such cases, algorithms are usually evaluated by sample complexity or simple regret (Bubeck et al., 2009). Pure exploration is also related to randomized controlled trials (RCTs) including modern variants that involve sequential

Algorithm 2 Explore-then-Commit (ETC)

1: Input: action space  $\mathcal{A}$ , where  $K = |\mathcal{A}|$ , number of exploration steps, m, horizon T. 2: for t = 1, ..., T do 3: if  $t \leq mK$  then 4: play  $A_t = (t \mod K) + 1$ . 5: else 6: play  $A_t = \operatorname{argmax}_a \hat{\mu}_a(mK)$ . 7: end if 8: receive  $R_t$ . 9: end for

randomization such as sequential multiple assignment randomized trials (SMARTs) (Lei et al., 2012) and micro-randomized trials (MRTs) (Klasnja et al., 2015). Randomized trials are typically designed to enable estimation of treatment effects with sufficient statistical power. Since the concerns of pure exploration and randomized controlled trials are different from those of bandit algorithms, we do not discuss them further in this chapter. However, note that in an actual application, methodology from bandits and randomized trials may need to be integrated. Researchers may start off with a randomized trial and follow it up with bandit algorithm in the next iteration of their health app. They can also decide to run a randomized trial for one health outcome while simultaneously running a bandit for a different outcome (e.g., an outcome related to user engagement with the health app) in the same study. There is also ongoing work (Yao et al., 2020; Zhang et al., 2021) on enabling the kind of statistical analysis done after randomized trials on data collected via online bandit algorithms.

**Explore-then-Commit (ETC).** We start with a simple two-stage algorithm: Explore-then-Commit (ETC). In the first stage of ETC, the learner plays every arm for a fixed number of times (m) and obtain estimates of the expected rewards. In the second stage, the learner commits to the best arm according to the estimates in the first stage. For every arm a, let  $\hat{\mu}_a(t)$  denote the estimated expected reward up to time t:

$$\hat{\mu}_a(t) = \frac{1}{T_a(t)} \sum_{s=1}^t \mathbb{1}_{\{A_s=a\}} R_s,$$

where  $T_a(t) = \sum_{s=1}^t \mathbb{1}_{\{A_s=a\}}$  is the number of times action a has been performed up to round t.

With the above definitions, we are ready to present the ETC in Algorithm 2. The overall performance of ETC crucially depends on the parameter m. If m is too small, the algorithm cannot estimate the performance of every arm accurately, so it is likely to exploit a sub-optimal arm in the second stage, which leads to high regret. If m is too big, the first stage (explore step) plays with sub-optimal arms for too many times, so that the regret can be large again. The art is to choose an optimal value for m in order to minimize the total regret incurred in both stages. Specifically, ETC achieves  $O(T^{2/3})$  worst-case regret<sup>3</sup> by choosing  $m = O(T^{2/3})$  (Lattimore and Szepesvári, 2020). Sub-linear regret  $O(T^{2/3})$  performance is good as a starting point. Next, we will introduce another two classic algorithms which incur even less regret.

Upper Confidence Bound (UCB). There are several types of exploration strategies to select actions such as greedy, Boltzmann, optimism and pessimism. Suppose the agent has reward estimates  $\hat{\mu}_a$  for all actions. A greedy exploration strategy simply selects the action with the highest  $\hat{\mu}_a$ . A Boltzmann exploration strategy picks each action with probability proportional to  $\exp(\eta \hat{\mu}_a)$ , where  $\eta$  is a tuning parameter. Boltzmann becomes greedy as  $\eta$  goes to infinity. For an optimism strategy, one picks the action with the highest reward estimate plus some bonus term, i.e.  $\operatorname{argmax}_a \hat{\mu}_a + \operatorname{bonus}_a$ . In contrast, a pessimism strategy would pick the action:  $\operatorname{argmax}_a \hat{\mu}_a - \operatorname{bonus}_a$ .

<sup>&</sup>lt;sup>3</sup>Ignoring parameters other than T.

Algorithm 3 Upper Confidence Bound (UCB)

1: **Input:**  $\delta$ ,  $\mathcal{A}$  where  $K = |\mathcal{A}|$ .

2: for t = 1, ..., T do

- 3: play  $A_t = \operatorname{argmax}_{a \in \mathcal{A}} \operatorname{UCB}_a(t-1, \delta).$
- 4: receive reward  $R_t$  and update the upper confidence bound terms according to (1).

Out of these strategies, UCB algorithm follows the optimism strategy, in particular, a famous principle called *optimism in the face of uncertainty (OFU)*, which means that one should act as if the environment is the best possible one among those that are *plausible* given current experience. The reason OFU works is that misplaced optimism gets corrected when under-explored actions are tried and low rewards are observed. In contrast, pessimism does not work (at least in the online setting; for the offline setting things can be different (Jin et al., 2021)) since wrong beliefs about low performance of under-explored actions do not get a chance to get revised by collecting more data from those actions.

At every step t, the UCB algorithm updates a value called *upper confidence bound* defined for each action  $a \in \mathcal{A}$  and confidence level  $\delta \in [0, 1]$  as follows.

$$UCB_{a}(t-1,\delta) = \begin{cases} \infty, & \text{if } T_{a}(t-1) = 0\\ \hat{\mu}_{a}(t-1) + \sqrt{\frac{2\log(1/\delta)}{T_{a}(t-1)\vee 1}}, & \text{otherwise.} \end{cases}$$
(1)

The learner chooses the action with the highest UCB value at each step. Overall, UCB (Algorithm 3) guarantees a  $\tilde{O}(\sqrt{T})$  worst-case regret (where the informal  $\tilde{O}(\cdot)$  notation hides constants and logarithmic factors).

According to the construction of upper confidence bounds, an action will be selected under two circumstances: under-explored  $(T_a(t-1) \text{ small})$  or well-explored with good performance  $(\hat{\mu}_a(t-1) \text{ large})$ . The upper confidence bound for an action gets close to its true mean after being selected for enough times. A sub-optimal action will only be played if its upper confidence bound is larger than that of the optimal arm. However, this is unlikely to happen too often. The upper confidence bound for the sub-optimal action will eventually fall below that of the optimal action as we play the sub-optimal actions more times. We present the regret guarantee for UCB (Algorithm 3) in Theorem 1.

**Theorem 1** (Regret for UCB Algorithm). If  $\delta = 1/T$ , then the problem-dependent regret of UCB, as defined in Algorithm 3, on any 1-subgaussian bandit  $\nu$  is bounded by

$$\operatorname{Reg}_{T}(\operatorname{UCB},\nu) \leq 3\sum_{a\in\mathcal{A}}\Delta_{a} + \sum_{a:\Delta_{a}>0}\frac{16\log T}{\Delta_{a}},$$
(2)

where  $\Delta_a := \mu^*(\nu) - \mu_a(\nu)$  represents the corresponding gap term. The worst-case regret bound of UCB is:

$$\operatorname{Reg}_{T}(\operatorname{UCB}) = O(\sqrt{KT\log T}). \tag{3}$$

*Proof.* We only present the worst-case regret for simplicity. For the problem-dependent regret proof, we refer the reader to Chapter 7 in Lattimore and Szepesvári (2020).

We define a good event E as follows:

$$E := \left\{ \left| \hat{\mu}_a(t-1) - \mu_a \right| \le \sqrt{\frac{2\log(1/\delta)}{T_a(t-1)\vee 1}}, \forall t \in [T], a \in \mathcal{A} \right\}.$$

By Hoeffding inequality and union bound, one can show that  $\mathbb{P}(E^c) \leq 2TK\delta^4$ . Next, we decompose the

<sup>5:</sup> end for

regret.

$$\operatorname{Reg}_{T}(\operatorname{UCB}) = \mathbb{E}\left[\sum_{t=1}^{T} \mu^{*} - \operatorname{UCB}_{a^{*}}(t-1,\delta) + \operatorname{UCB}_{a^{*}}(t-1,\delta) - \operatorname{UCB}_{A_{t}}(t-1,\delta) + \operatorname{UCB}_{A_{t}}(t-1,\delta) - \mu_{A_{t}}\right]$$
$$\leq \mathbb{E}\left[\sum_{t=1}^{T} \mu^{*} - \operatorname{UCB}_{a^{*}}(t-1,\delta) + \operatorname{UCB}_{A_{t}}(t-1,\delta) - \mu_{A_{t}}\right].$$

The inequality in above expression is due to the action selection criterion in UCB algorithm. Condition on event E, we have  $\mu^* - \text{UCB}_{a^*}(t-1,\delta) \leq 0$  for all  $t \in [T]$ ; otherwise, the regret can be bounded by 2T. Combining these arguments, we have

$$\begin{aligned} \operatorname{Reg}_{T}(\operatorname{UCB}) &\leq \mathbb{P}(E^{c}) \cdot 2T + \mathbb{P}(E) \cdot \mathbb{E}\left[\sum_{t=1}^{T} \operatorname{UCB}_{A_{t}}(t-1,\delta) - \mu_{A_{t}} \mid E\right] \\ &\leq 4T^{2}K\delta^{4} + 2\mathbb{E}\left[\sum_{t=1}^{T}\sqrt{\frac{2\log\left(1/\delta\right)}{T_{A_{t}}(t-1)\vee 1}}\right] \text{ (By definition of event } E) \\ &\leq 4T^{2}K\delta^{4} + \sqrt{8\log\left(1/\delta\right)}\sum_{a\in\mathcal{A}}\sum_{t=1}^{T}\mathbb{E}\left[\sqrt{\frac{1}{T_{A_{t}}(t-1)\vee 1}}\mathbb{1}_{\{A_{t}=a\}}\right] \\ &\leq 4T^{2}K\delta^{4} + \sqrt{8\log\left(1/\delta\right)}\sum_{a\in\mathcal{A}}\int_{1}^{T_{a}(T)}\sqrt{1/s}\,\mathrm{d}s \\ &\leq 4T^{2}K\delta^{4} + \sqrt{8\log\left(1/\delta\right)}\sqrt{KT} = O(\sqrt{KT\log T}) \text{ (Set } \delta = 1/T). \end{aligned}$$

The last inequality in above is by Cauchy-Schwarz inequality.

The UCB family has many variants, one of which is to replace the upper confidence bound for every action by  $\hat{\mu}_a(t-1) + \sqrt{2\log(1+t\log^2(t))/T_a(t-1)}$ . Even though the regret dominant terms  $(\sqrt{KT\log T})$  and  $\sum_{a:\Delta_a>0} \frac{\log T}{\Delta_a}$  for this version has the same order as those of Algorithm 3, the leading constants for the two dominant terms become smaller.

Then one may ask the question: is it possible to further improve the regret bound of UCB and above variant? The answer is yes. Audibert et al. (2009) proposed an algorithm called *MOSS* (Minimax Optimal Strategy in the Stochastic case). *MOSS* replaces the upper confidence bounds in Algorithm 3 by

$$\hat{\mu}_a(t-1) + \sqrt{\frac{\max\left\{\log\left(\frac{T}{KT_a(t-1)}\right), 0\right\}}{T_a(t-1)}}$$

Under this construction, the worst-case regret of MOSS is guaranteed to be only  $O(\sqrt{KT})$ .

However, MOSS is not always good. One can easily construct regimes where the problem-dependent regret of MOSS is worse than UCB (Lattimore, 2015). On the other hand, the improved UCB algorithm proposed by Auer and Ortner (2010) satisfies a problem-dependent regret that is similar to (1), but the worstcase regret is  $O(\sqrt{KT \log K})$ . Later on, by carefully constructing the upper confidence bounds, Optimally Confidence UCB algorithm (Lattimore, 2015) and AdaUCB algorithm (Lattimore, 2018) are shown to achieve  $O(\sqrt{KT})$  worst-case regret and their problem-dependent regret bounds are also not worse than that of the UCB algorithm. There are many more UCB variants in the literature that we do not cover in this chapter. The reader may refer to Table 2 in Lattimore (2018) for a comprehensive summary.

**Successive Elimination (SE).** We now describe the SE algorithm that also relies on the upper confidence bound calculations. The idea is similar to UCB such that a sub-optimal arm is very unlikely to have large a

Algorithm 4 Successive Elimination (SE)

- 1: Input:  $\delta, \mathcal{A}$  where  $K = |\mathcal{A}|$ .
- 2: while  $\mathcal{A}$  contains more than one arm do
- 3: play every arm in  $\mathcal{A}$  once and update the UCBs and LCBs using (1) and (4).
- 4: Eliminate all arms a s.t.  $\exists a' \in \mathcal{A}$  with  $UCB_a(t-1,\delta) < LCB_{a'}(t-1,\delta)$ .
- 5: end while

upper confidence bound if it has been selected for enough times. At every round, SE maintains a confidence interval for the mean reward of every arm and removes all arms whose reward upper bound is smaller than the lower bound of the biggest estimated reward arm. The procedure ends when there is only one arm remained. We describe the SE algorithm in Algorithm 4 and define the UCB terms as (1) and LCB terms as

$$\operatorname{LCB}_{a}(t-1,\delta) = \begin{cases} \infty, & \text{if } T_{a}(t-1) = 0\\ \hat{\mu}_{a}(t-1) - \sqrt{\frac{2\log(1/\delta)}{T_{a}(t-1)\vee 1}}, & \text{otherwise.} \end{cases}$$
(4)

SE was first proposed in Even-Dar et al. (2006) along with a similar action elimination based algorithm: Median Elimination (ME). They studied the probably approximately correct (PAC) setting (Haussler and Warmuth, 2018). In particular, Even-Dar et al. (2006) shows that for given K arms, it suffices to pull the arms for  $O(\frac{K}{\varepsilon^2} \log(1/\delta))$  times to find an  $\varepsilon$ -optimal arm with probability at least  $1 - \delta$ . It is not hard to prove that SE also satisfies the following regret bound.

**Theorem 2** (Regret for SE Algorithm). If  $\delta = 1/T$ , the worst-case regret of SE over 1-subgaussian bandit environments is bounded by

$$\operatorname{Reg}_{T}(\operatorname{SE}) = O(\sqrt{KT\log T}).$$
(5)

Proof. Without loss of generality, we assume the optimal arm  $a^*$  is unique. Define the event E by  $\{|\hat{\mu}_a(t) - \mu_a| \leq c_a(t,\delta), \forall a,t\}$ , where  $c_a(t,\delta) = \sqrt{\frac{2\log(1/\delta)}{T_a(t)\vee 1}}$  denotes the confidence set width. By Hoeffding inequality and union bound, one can show that  $\mathbb{P}(E^c) \leq 2\delta^4 T K$ .

Define t as the last round when arm a is not eliminated yet. According to the elimination criterion in SE, the reward gap term can be bounded as:

$$\Delta_a := \mu^* - \mu_a \le 2(c_{a^*}(t, \delta) + c_a(t, \delta)) = O(c_a(t, \delta)).$$

The last equality holds as  $T_a(t)$  and  $T_{a^*}(t)$  differ at most by 1 by construction. Since t is the last round a being played, we have  $T_a(t) = T_a(T)$  and  $c_a(t) = c_a(T)$ , which implies below property for all non-optimal arms a:

$$\Delta_a \le O\left(\sqrt{\frac{\log(1/\delta)}{T_a(T)}}\right).$$

We thus obtain that under event E,

$$\sum_{t=1}^{T} \mathbb{E}[\mu^* - \mu_{a_t}|E] \le \sum_{a \in \mathcal{A} \setminus \{a^*\}} T_a(T) \Delta_a \le O(\sqrt{\log(1/\delta)}) \sum_{a \in \mathcal{A}} \sqrt{T_a(T)} \le O(\sqrt{KT \log(1/\delta)}),$$

where the last inequality is by Cauchy-Schwarz inequality. Take  $\delta = 1/T$ ,  $\text{Reg}_T(\text{SE}) = O(\sqrt{KT \log T})$  by conditional expectation calculations.

#### Algorithm 5 Thompson Sampling (TS)

1: Input: Bayesian bandit environment  $(\mathcal{E}, \mathcal{B}(\mathcal{E}), Q, P)$   $(\mathcal{B}(\cdot)$  is the Borel set), action set  $\mathcal{A}$  with  $|\mathcal{A}| = K$ .

- 2: for t = 1, ..., T do
- 3: Sample  $\nu_t \sim Q(\cdot | A_1, R_1, \dots, A_{t-1}, R_{t-1}).$
- 4: Play  $A_t = \operatorname{argmax}_{i \in [K]} \mu_i(\nu_t)$ .
- 5: end for

**Thompson Sampling (TS).** All of the methods we have mentioned so far select their actions based on a frequentist view. TS uses one of the oldest heuristic (Thompson, 1933) for choosing actions and addresses the exploration-exploitation dilemma based on a Bayesian philosophy of learning. The idea is simple. Before the game starts, the agent chooses a prior distribution over a set of possible bandit environments. At every round, the agent samples an environment from the posterior and acts according to the optimal action in that environment. The exploration in TS comes from the randomization over bandit environments. At the beginning, the posterior is usually poorly concentrated, then the policy will likely explore. As more data being collected, the posterior tends to concentrate towards the true environment and the rate of exploration decreases. We present the TS algorithm in Algorithm 5.

To formally describe how TS works, we start with several definition related to Bayesian bandits.

**Definition 2** (K-armed Bayesian bandit environment). A K-armed Bayesian bandit environment is a tuple  $(\mathcal{E}, \mathcal{G}, Q, P)$  where  $(\mathcal{E}, \mathcal{G})$  is a measurable space and Q is a probability measure on  $(\mathcal{E}, \mathcal{G})$  called the prior.  $P = (P_{\nu a} : \nu \in \mathcal{E}, a \in \mathcal{A})$  is the reward distribution for arms in bandit  $\nu$ , where  $|\mathcal{A}| = K$ .

Given a K-armed Bayesian bandit environment  $(\mathcal{E}, \mathcal{G}, Q, P)$  and a policy  $\pi$ , the Bayesian regret is defined as:

$$\operatorname{BReg}_T(\pi, Q) = \int_{\mathcal{E}} \operatorname{Reg}_T(\pi, \nu) dQ(\nu).$$

TS has been analyzed in both of the frequentist and the Bayesian settings and we will start with the Bayesian results.

**Theorem 3** (Bayesian Regret for TS Algorithm). For a K-armed Bayesian bandit environment  $(\mathcal{E}, \mathcal{G}, Q, P)$  such that  $P_{\nu a}$  is 1-subgaussian for all  $\nu \in \mathcal{E}$  and  $a \in [K]$  with mean in [0, 1]. Then the policy  $\pi$  of TS satisfies

$$BReg_T(\pi, Q) = O(\sqrt{KT\log T}).$$
(6)

*Proof.* The proof is quite similar to that of UCB. We abbreviate  $\mu_a = \mu_a(\nu)$  and let  $a^* = \operatorname{argmax}_{a \in [K]} \mu_a$  be the optimal arm. Note that  $a^*$  is a random variable depending on  $\nu$ . For every  $a \in [K]$ , we define a clipped upper bound term

UCB<sub>a</sub>(t - 1) = 
$$\hat{\mu}_a(t - 1) + \sqrt{\frac{2\log(1/\delta)}{1 \vee T_a(t - 1)}}$$

where  $\hat{\mu}_a(t-1)$  and  $T_a(t-1)$  are defined in the same way as those in UCB. We define event E such that for all  $t \in [T]$  and  $a \in \mathcal{A}$ ,

$$|\hat{\mu}_a(t-1) - \mu_a| < \sqrt{\frac{2\log(1/\delta)}{1 \vee T_a(t-1)}}$$

By Hoeffding inequality and union bound, one can show that  $\mathbb{P}(E^c) \leq 2TK\delta^4$ . This result will be used in later steps.

Let  $\mathcal{F}_t = \sigma(A_1, R_1, \dots, A_t, R_t)$  be the  $\sigma$ -algebra generated by the interaction sequence up to time t. The key insight for the whole proof is to observe below property from the definition of TS:

$$\mathbb{P}(a^* | \mathcal{F}_{t-1}) = \mathbb{P}(A_t | \mathcal{F}_{t-1}) \text{ a.s.}$$
(7)

Using above property and  $\operatorname{BReg}_T = \mathbb{E}\left[\sum_{t=1}^T (\mu_{a^*} - \mu_{A_t})\right] = \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}\left[(\mu_{a^*} - \mu_{A_t}) \mid \mathcal{F}_{t-1}\right]\right]$ , we have

$$\mathbb{E}\left[\mu_{a^*} - \mu_{A_t} \mid \mathcal{F}_{t-1}\right] = \mathbb{E}\left[\mu_{a^*} - \mathrm{UCB}_{a^*}(t-1) + \mathrm{UCB}_{A_t}(t-1) - \mu_{A_t} \mid \mathcal{F}_{t-1}\right],$$

and thus

$$BReg_T = \mathbb{E}\left[\sum_{t=1}^T (\mu_{a^*} - UCB_{a^*}(t-1)) + \sum_{t=1}^T (UCB_{A_t}(t-1) - \mu_{A_t})\right].$$

Conditioning on the high-probability event E, the first sum is negative and the second sum is of the order  $O(\sqrt{KT \log(1/\delta)})$ , while  $\operatorname{BReg}_T \leq 2T$  if conditioning on  $E^c$ . Take  $\delta = 1/T$ , one can verify that  $\operatorname{BReg}_T = O(\sqrt{KT \log T})$ .

Compared to the analysis for Bayesian regret, frequentist (worst-case) regret analysis for TS gets a lot more technical. The key reason behind this is that the worst-case regret does not have an expectation with respect to the prior and therefore the property in (7) cannot be used. Even though TS was well-known to be easy to implement and competitive with state of the art methods, it lacked worst-case regret analysis for a long time. Significant progress was made by Agrawal and Goyal (2012) and Kaufmann et al. (2012). In Agrawal and Goyal (2012), the first logarithmic bound on the frequentist regret of TS was proven. Kaufmann et al. (2012) provided a bound that matches the asymptotic lower bound of Lai and Robbins (1985). However, both of these bounds were problem-dependent. The first near-optimal worst-case regret  $O(\sqrt{KT \log T})$  was proved by Agrawal and Goyal (2013) for Bernoulli bandits with Beta prior, where the reward is either zero or one. For TS that uses Gaussian prior, the same work proved a  $O(\sqrt{KT \log K})$  worst-case regret. Jin et al. (2020) proposed a variant of TS called MOTS (Minimax Optimal TS) that achieves  $O(\sqrt{KT})$  regret.

#### 2.1.2 Adversarial Multi-armed Bandit

In stochastic bandit models, the rewards are assumed to be strictly i.i.d. given actions. This assumption can be violated easily in practice. For example, the health feedback for a patient after certain treatments may vary slightly across times and the way it varies is usually unknown. In such scenarios, a best action that maximizes the total reward still exists, but algorithms designed for stochastic bandit environments are no longer guaranteed to work. As a more robust counterpart to the stochastic model, we study the adversarial bandit model in this section, in which the assumption that a single action is good in hindsight is retained but the rewards are allowed to be chosen adversarially.

The adversarial bandit environment is often called as the adversary. In an adversarial bandit problem, the adversary secretly chooses reward vectors  $r := \{r_t\}_{t=1}^T$  where  $r_t \in [0, 1]^K$  corresponds to the rewards over all actions at time t. In every round, the agent chooses a distribution over the actions  $P_t$ . An action  $A_t \in [K]$  is sampled from  $P_t$  and the agent receives the reward  $r_{tA_t}$ . A policy  $\pi$  in this setting maps the history sequences to distributions over actions. We evaluate the performance of policy  $\pi$  by the expected regret, which is the cumulative reward difference between the best fixed action and the agent's selections:

$$\operatorname{Reg}_{T}(\pi, r) = \max_{a \in \mathcal{A}} \sum_{t=1}^{T} r_{ta} - \mathbb{E}\left[\sum_{t=1}^{T} r_{tA_{t}}\right]$$
(8)

The worst-case regret of policy  $\pi$  is defined by

$$\operatorname{Reg}_{T}(\pi) = \sup_{r \in [0,1]^{T \times K}} \operatorname{Reg}_{T}(\pi, r).$$
(9)

It may not be very clear at the beginning that why we define the regret by comparing to the fixed best action instead of the best action at every round. In the later case, the regret should be  $\operatorname{Reg}_T'(\pi, x) = \mathbb{E}\left[\sum_{t=1}^T \max_{a \in \mathcal{A}} r_{ta} - \sum_{t=1}^T r_{tA_t}\right]$ . However, this definition provides the adversary too much power so that for any policy, one can show  $\operatorname{Reg}_T'(\pi, r)$  can be  $\Omega(T)$  for certain reward vectors  $r \in [0, 1]^{K \times T}$ .

**Remark on randomized policy:** In stochastic bandit models, the optimal action is deterministic and the optimal policy is simply to select the optimal action at every round. However, in adversarial bandit setting, the adversary has great power in designing the reward. It may know the agent's policy and design the rewards accordingly, so that a deterministic policy can incur linear regret. For example, we consider there are two actions, whose reward is either 0 or 1 at any time. For a deterministic policy, the agent decides to choose an action at time t. Then the adversary knows it and can set the reward of that action at time t to be 0 and the reward of the unselected action to be 1. The cumulative regret will be T after T rounds. However, one can improve the performance by a randomized policy, e.g., choosing either action with probability 0.5, then the adversary cannot make you incur regret 1 at every round by manipulating the reward values for both actions.

**Exponential-weight algorithm for Exploration and Exploitation (EXP3).** We now study one of the most famous adversarial bandit algorithm called EXP3. Before describing the algorithm, we define some related terms below. In a randomized policy, the conditional probability of the action a being played is denoted by

$$P_{ta} = \mathbb{P}(A_t = a \mid A_1, R_1, \dots, A_{t-1}, R_{t-1}).$$

Assuming  $P_{ta} > 0$  almost surely for all policies, a natural way to define the importance-weighted estimator of  $r_{ta}$  is

$$\hat{R}_{ta} = \frac{\mathbb{1}_{\{A_t=a\}} R_t}{P_{ta}}.$$
(10)

Let  $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot | A_1, R_1, \dots, A_{t-1}, R_{t-1}]$ . A simple calculation shows that  $\hat{R}_{ta}$  is conditionally unbiased, i.e.  $\mathbb{E}_t[\hat{R}_{ta}] = r_{ta}$ . However, the variance of estimator  $\hat{R}_{ta}$  can be extremely large when  $P_{ta}$  is small and  $r_{ta}$  is non-zero. Let  $A_{ta} := \mathbb{1}_{\{A_t = a\}}$ , then the variance of the estimator is:

$$\mathbb{V}_t[\hat{R}_{ta}] = \mathbb{E}_t[\hat{R}_{ta}^2] - r_{ta}^2 = \mathbb{E}_t\left[\frac{A_{ta}r_{ta}^2}{P_{ta}^2}\right] - r_{ta}^2 = \frac{r_{ta}^2(1 - P_{ta})}{P_{ta}}.$$

An alternative estimator is:

$$\hat{R}_{ta} = 1 - \frac{\mathbb{1}_{\{A_t = a\}}}{P_{ta}} (1 - R_t).$$
(11)

This estimator is still unbiased and its variance is

$$\mathbb{V}_t[\hat{R}_{ta}] = y_{ta}^2 \frac{1 - P_{ta}}{P_{ta}},$$

where we define  $y_{ta} = 1 - r_{ta}$ .

The best choice of the estimator  $\hat{R}_{ta}$  depends on the actual rewards. One should use (10) for small rewards and (11) for large rewards. So far we have learned how to construct reward estimators for given sampling distributions  $P_{ta}$ . EXP3 algorithm provides a way to design the  $P_{ta}$  terms. Let  $\hat{S}_{t,a} = \sum_{s=1}^{t} \hat{R}_{sa}$  be the total estimated reward until the end of round t, where  $\hat{R}_{si}$  is defined in (11). We present EXP3 in Algorithm 6.

Surprisingly, even though adversarial bandit problems look more difficult than stochastic bandit problems due to the strong power of the adversary, one can show that the adversarial regret for EXP3 algorithm has the same order as before, i.e.  $O(\sqrt{KT \log T})$ .

Algorithm 6 Exponential-weight Algorithm for Exploration and Exploitation (Exp3)

1: Input:  $T, K, \eta$ 2: Set  $\hat{S}_{0,a} = 0$  for all a. 3: for t = 1, ..., T do 4: Calculate  $P_{ta} \leftarrow \frac{\exp(\eta \hat{S}_{t-1,a})}{\sum_{a' \in \mathcal{A}} \exp(\eta \hat{S}_{t-1,a'})}$  for all  $a \in [K]$ 5: Sample  $A_t$  from  $P_t$  and receive reward  $R_t$ 6: Calculate  $\hat{S}_{t,a} \leftarrow \hat{S}_{t-1,a} + 1 - \frac{\mathbb{I}_{\{A_t = a\}}(1-R_t)}{P_{ta}}$ 7: end for

**Theorem 4** (Regret for EXP3 Algorithm). Let  $r \in [0, 1]^{T \times K}$ . With learning rate  $\eta = \sqrt{\log K/(KT)}$ , we have

$$\operatorname{Reg}_{T}(\operatorname{EXP3}, r) \le 2\sqrt{KT\log K}.$$
 (12)

*Proof.* The proof for EXP3 is different than those for the stochastic bandit algorithms. We first define the expected regret relative to using action a in T rounds:

$$\operatorname{Reg}_{Ta} = \sum_{t=1}^{T} r_{ta} - \mathbb{E}\left[\sum_{t=1}^{T} R_t\right].$$

According to the definition of the adversarial bandit regret, the final result will follow if we can bound  $R_{Ta}$  for every  $a \in \mathcal{A}$ . It's not hard to show  $\mathbb{E}[\hat{S}_{T,a}] = \sum_{t=1}^{T} r_{ta}$  and  $\mathbb{E}_t[R_t] = \sum_{a \in \mathcal{A}} P_{ta} r_{ta} = \sum_{i=1}^{K} P_{ta} \mathbb{E}[\hat{R}_{ta}]$  hold using the definition of  $\hat{R}_{ta}$ . Then we can re-write  $R_{Ta}$  as  $\mathbb{E}[\hat{S}_{T,a} - \hat{S}_T]$ , where  $\hat{S}_T = \sum_{t=1}^{T} \sum_{a \in \mathcal{A}} P_{ta} \hat{R}_{ta}$ . Let  $W_t = \sum_{a \in \mathcal{A}} \exp(\eta \hat{S}_{t,a})$ ,  $\hat{S}_{0,a} = 0$  and  $W_0 = K$ , then one can show that

$$\exp\left(\eta \hat{S}_{T,a}\right) \leq \sum_{a' \in \mathcal{A}} \exp\left(\eta \hat{S}_{T,a'}\right) = W_T = K \prod_{t=1}^T \frac{W_t}{W_{t-1}} = K \prod_{t=1}^T \sum_{a' \in \mathcal{A}} P_{ta'} \exp\left(\eta \hat{R}_{ta'}\right).$$

We next bound the ratio term

$$\frac{W_t}{W_{t-1}} \le 1 + \eta \sum_{a' \in \mathcal{A}} P_{ta'} \hat{R}_{ta'} + \eta^2 \sum_{a \in \mathcal{A}} P_{ta'} \hat{R}^2_{ta'} \le \exp\left(\eta \sum_{a' \in \mathcal{A}} P_{ta'} \hat{R}_{ta'} + \eta^2 \sum_{a' \in \mathcal{A}} P_{ta'} \hat{R}^2_{ta'}\right),$$

using inequalities  $e^x \leq 1 + x + x^2$  for  $x \leq 1$  and  $1 + x \leq e^x$  for  $x \in \mathbb{R}$ .

Combining with previous results, we have

$$\exp\left(\eta \hat{S}_{T,a}\right) \le K \exp\left(\eta \hat{S}_T + \eta^2 \sum_{t=1}^T \sum_{a' \in \mathcal{A}} P_{ta'} \hat{R}_{ta'}^2\right).$$

Taking logarithm on both sides and rearranging give us

$$\hat{S}_{T,a} - \hat{S}_T \le \frac{\log K}{\eta} + \eta \sum_{t=1}^T \sum_{a' \in \mathcal{A}} P_{ta'} \hat{R}_{ta'}^2.$$
(13)

To bound  $R_{Ta}$ , we only need to bound the expectation of the second term in above. By standard (conditional) expectation calculations, one can get

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{a'\in\mathcal{A}}P_{ta'}\hat{R}_{ta'}^2\right] \leq TK.$$

By substituting above inequality into (13), we get

$$\operatorname{Reg}_{Ta} \le \frac{\log K}{\eta} + \eta T K = 2\sqrt{KT\log K},\tag{14}$$

where we choose  $\eta = \sqrt{\log K/(TK)}$ . By definition, the overall regret  $\operatorname{Reg}_T(\operatorname{EXP3}, r)$  has the same upper bound as above.

We just proved the expected regret of EXP3. However, if we consider the distribution of the random regret, EXP3 is not good enough. Define the random regret as  $\widehat{\operatorname{Reg}}_T = \max_{a \in \mathcal{A}} \sum_{t=1}^T r_{ta} - \sum_{t=1}^T R_t$ . One can show that for all large enough T and reasonable choices of  $\eta$ , there exists a bandit such that the random regret of EXP3 satisfies  $\mathbb{P}(\widehat{\operatorname{Reg}}_T \geq T/4) \geq c > 0$ , where c is a constant. That means EXP3 sometimes can incur linear regret with non-trivial probability, which makes EXP3 unsuitable for practical problems. This phenomenon is caused by the high variance of the regret distribution. In next section, we will discuss how to resolve this problem by slightly modifying EXP3.

**EXP3-IX (EXP3 with Implicit eXploration).** We have learned that small  $P_{ta}$  terms can cause enormous variance on the reward estimator, which then leads to high variance on the regret distribution. Thus, EXP3-IX (Neu, 2015) redefines the loss-estimator as

$$\hat{Y}_{ta} = \frac{\mathbb{1}_{\{A_t=a\}} Y_t}{P_{ta} + \gamma},$$
(15)

where  $Y_t = 1 - R_t$  denotes the loss at round t and  $\gamma > 0$ .  $\hat{Y}_{ta}$  is a biased estimator for  $y_{ta} = 1 - r_{ta}$  due to  $\gamma$ , but the variance can be reduced. An optimal choice for  $\gamma$  needs to balance the bias and variance. Other than this slight change on the loss estimator, the remaining procedures remain the same as EXP3. The name of 'IX' (Implicit eXploration) can be justified by the following argument:

$$\mathbb{E}_t[\hat{Y}_{ta}] = \frac{P_{ta}y_{ta}}{P_{ta} + \gamma} \le y_{ta}.$$

The effect of adding a  $\gamma$  to the denominator is that EXP3-IX tries to decrease the large losses for some actions, so that such actions can still be chosen occasionally. As a result, EXP3-IX explores more than EXP3. Neu (2015) has proved the following high probability regret bound for EXP3-IX.

**Theorem 5** (Regret for EXP3-IX Algorithm). With  $\eta = 2\gamma = \sqrt{\frac{2 \log K}{KT}}$ , EXP3-IX guarantees that

$$\widehat{\operatorname{Reg}}_{T}(\operatorname{EXP3-IX}) \leq 2\sqrt{2KT\log K} + \left(\sqrt{\frac{2KT}{\log K}} + 1\right)\log(2/\delta)$$
(16)

with probability at least  $1 - \delta$ .

### 2.1.3 Lower Bound for MAB Problems

We have discussed two types of bandit models and their corresponding algorithms in regret minimization. Then a natural question are: what is the minimal regret bound we can hope for? To answer this question, we will introduce two types of lower bound results: minimax lower bound and instance dependent lower bound. Both of them are useful for describing the hardness of a class of bandit problems and are often used to evaluate the optimality of an existing algorithm. For example, suppose the worst-case regret of a policy  $\pi$  matches the minimax lower bound up to a universal constant, we say that the policy  $\pi$  is minimax-optimal.

**Minimax Lower Bounds.** We consider a Gaussian bandit environment, in which the reward for every arm is Gaussian-distributed. We denote the class of Gaussian bandits with unit variance by  $\mathcal{E}_{\mathcal{N}}^{K}(1)$  and use  $\mu \in \mathbb{R}^{K}$  as the reward mean vector. In particular,  $\nu_{\mu} \in \mathcal{E}_{\mathcal{N}}^{K}(1)$  is a Gaussian bandit for which the *i*th arm has reward distribution  $\mathcal{N}(\mu_{i}, 1)$ . The following result provides a minimax lower bound for the Gaussian bandit class  $\mathcal{E}_{\mathcal{N}}^{K}(1)$ .

**Theorem 6** (Minimax Lower Bound for Gaussian Bandit Class). Let K > 1 and  $T \ge K - 1$ . For any policy  $\pi$ , there exists a mean vector  $\mu \in [0, 1]^K$  such that

$$\operatorname{Reg}_{T}(\pi,\nu_{\mu}) = \Omega(\sqrt{KT}). \tag{17}$$

*Proof.* To prove the lower bound, we start with constructing two bandit instances that are very similar to each other and hard to distinguish. Let  $\mu = (\Delta, 0, 0, ..., 0)$  denote the mean vector for the first unit variance Gaussian bandit. We use  $\mathbb{P}_{\mu}$  and  $\mathbb{E}_{\mu}$  to denote the probability and expectation induced by environment  $\nu_{\mu}$  and policy  $\pi$  up to time T. To choose the second environment, let

$$i = \operatorname{argmin}_{i>1} \mathbb{E}_{\mu}[T_j(T)].$$

Define the reward mean vector for the second bandit as  $\mu' = (\Delta, 0, \dots, 0, 2\Delta, 0, \dots, 0)$ , where  $\mu'_i = 2\Delta$ . Decomposing the regret leads to

$$\operatorname{Reg}_{T}(\pi,\nu_{\mu}) \geq \mathbb{P}_{\mu}(T_{1}(T) \leq T/2) \frac{T\Delta}{2},$$
  
$$\operatorname{Reg}_{T}(\pi,\nu_{\mu'}) > \mathbb{P}_{\mu'}(T_{1}(T) > T/2) \frac{T\Delta}{2}.$$

Then, applying the Bretagnolle-Huber inequality, we get

$$\operatorname{Reg}_{T}(\pi,\nu_{\mu}) + \operatorname{Reg}_{T}(\pi,\nu_{\mu'}) \geq \frac{T\Delta}{4} \exp(-\operatorname{KL}(\mathbb{P}_{\mu},\mathbb{P}_{\mu'})).$$

It remains to upper bound the KL-divergence term in above. By divergence decomposition, one can show that

$$\operatorname{KL}(\mathbb{P}_{\mu}, \mathbb{P}_{\mu'}) = \sum_{i=1}^{K} \mathbb{E}_{\mu}[T_i(T)] \operatorname{KL}(\mathbb{P}_i, \mathbb{P}'_i) = \mathbb{E}_{\mu}[T_i(T)] \operatorname{KL}(\mathcal{N}(0, 1), \mathcal{N}(2\Delta, 1))$$
$$= \mathbb{E}_{\mu}[T_i(T)] \frac{(2\Delta)^2}{2} \le \frac{2T\Delta^2}{K-1}.$$

In above, we use  $\mathbb{P}_i$  and  $\mathbb{P}'_i$  denote the reward distribution of the *i*th arm in  $\nu_{\mu}$  and  $\nu_{\mu'}$ , respectively. For the last inequality, since  $\sum_{j=1}^{K} \mathbb{E}_{\mu}[T_j(T)] = T$ , it holds that  $\mathbb{E}_{\mu}[T_i(T)] \leq \frac{T}{K-1}$ . Combining with previous results, we know that

$$\operatorname{Reg}_{T}(\pi,\nu_{\mu}) + \operatorname{Reg}_{T}(\pi,\nu_{\mu'}) \geq \frac{T\Delta}{4} \exp\left(-\frac{2T\Delta^{2}}{K-1}\right)$$

Choosing  $\Delta = \sqrt{(K-1)/4T} \le 1/2$ , the result follows.

An algorithm is called minimax-optimal if its worst-case regret matches with the minimax lower bound.

**Instance Dependent Lower Bounds.** An algorithm with nearly minimax-optimal regret is not always preferred, since it may fail to take advantage of environments that are not the worst case. In practice, what is more desirable is to have algorithms that are near minimax-optimal, while their performance gets better on "easier" instances Lattimore and Szepesvári (2020). This motivates the study of instance dependent regret. In this section, we present two types of lower bound for instance dependent regret: one is asymptotic, the other is finite-time.

We first define consistent policy and present the asymptotic instance-dependent lower bound result.

**Definition 3** (Consistent Policy). A policy  $\pi$  is consistent if over bandit environment  $\mathcal{E}$  if for all bandits  $\nu \in \mathcal{E}$  and for all p > 0 it holds that

$$\operatorname{Reg}_T(\pi,\nu) = O(T^p) \text{ as } n \to \infty.$$

**Theorem 7** (Asymptotic Instance Dependent Lower Bound for Gaussian Bandits (Lattimore and Szepesvári, 2020)). For any policy  $\pi$  consistent over K-armed unit-variance Gaussian environments  $\mathcal{E}_{N}^{K}(1)$  and any  $\nu \in \mathcal{E}_{N}^{K}(1)$ , it holds that

$$\liminf_{T \to \infty} \frac{\operatorname{Reg}_T(\pi, \nu)}{\log T} \ge \sum_{i:\Delta_i > 0} \frac{2}{\Delta_i}$$

A policy is called asymptotically optimal if the equality in above theorem holds. Interestingly, building on the similar idea of Theorem 7, one can also develop a finite-time instance dependent lower bound result.

**Theorem 8** (Instance Dependent Lower Bound for Gaussian Bandits (Lattimore and Szepesvári, 2020)). Let  $\nu \in \mathcal{E}_{\mathcal{N}}^{K}(1)$  be a K-armed Gaussian bandit with mean vector  $\mu \in \mathbb{R}^{K}$  and suboptimality gaps  $\Delta \in [0, \infty)^{K}$ . Define a bandit environment:

$$\mathcal{E}(\nu) = \{\nu' \in \mathcal{E}_{\mathcal{N}}^{K}(1) : \mu_{i}(\nu') \in [\mu_{i}, \mu_{i} + 2\Delta_{i}]\}.$$

Suppose C > 0 and  $p \in (0,1)$  are constants and  $\pi$  is a policy such that  $\operatorname{Reg}_T(\pi,\nu') \leq CT^p$  for all T and  $\nu' \in \mathcal{E}(\nu)$ . Then the regret for instance  $\nu$  is lower bounded by

$$\operatorname{Reg}_{T}(\pi,\nu) \geq \frac{2}{(1+\varepsilon)^{2}} \sum_{i:\Delta_{i}>0} \left(\frac{(1-p)\log T + \log\left(\frac{\varepsilon\Delta_{i}}{8C}\right)}{\Delta_{i}}\right)^{+},\tag{18}$$

where  $(x)^+ := \max\{x, 0\}.$ 

The proof for above two instance dependent lower bounds are too technical that we do not want to go beyond. Interested readers are referred to Chapter 16 in Lattimore and Szepesvári (2020) for details.

## 2.2 Contextual Bandit

In many real-world applications, some side information is available to facilitate decision making. For instance, in healthcare applications, it is important to make decisions based on contextual information such as the person's demographic information, genetic information, life history, biomarkers, and environmental exposures. Ignoring such contextual information may lead to suboptimal decision making. To this end. the contextual bandits setting was introduced to model the side information that determines the reward distribution. The term "contextual bandit" is due to Langford and Zhang (2007) but similar settings have been considered earlier under the names "bandit problems with covariates" (Clayton, 1989; Sarkar, 1991; Yang et al., 2002) and "bandit problems with side information" (Wang et al., 2005; Goldenshluger and Zeevi, 2011). Let  $\mathcal{C}$  be a context space for the side information. The contextual bandits setting considered in this section is the same as the MAB described in Section 2.1 except for an extra contextual information  $C_t \in \mathcal{C}$  at each step t, respectively. We extend our modeling of the environment from  $\nu = (P_a : a \in \mathcal{A})$ to  $\nu = (P_{a,c} : a \in \mathcal{A}, c \in \mathcal{C})$ , i.e. the distribution of rewards depends on both action and context. The trajectory is extended to a sequence of triples,  $(A_t, R_t, C_t)$ , which represents the action, rewards and context at step t. An instantiation of the contextual bandit protocol in the healthcare setting is given in Algorithm 7. Contextual bandits enable precision medicine by taking appropriate actions for particular patients rather than the same action for a variety of patients. Many papers, like Rindtorff et al. (2019) on precision oncology with in vitro data and Zhou et al. (2019) on cancer treatment with genetic data, have effectively applied contextual bandits to precision medicine. Other applications can be found in Shrestha and Jain (2021).

Throughout the section, we discuss the setting where  $C_t$  is presented adversarially, rather than assuming  $C_t$  is sampled stochastically from an underlying distribution. For the stochastic setting, Goldenshluger and Zeevi (2013) gives an algorithm with some parametric assumptions and Yang et al. (2002) with non-parametric algorithm. In fact, the stronger stochastic assumption does not improve the worst-case regret bound over the adversarial setting in linear contextual bandit that we will introduce later. The stochastic setting does help when the instance-dependent regret bound is considered.

Algorithm 7 Contextual Bandit Framework in Healthcare

**Input:** Available treatment options  $(\mathcal{A})$ , treatment period length T. **for**  $t = 1, \ldots, T$  **do** Agent receives the context information  $(C_t)$  for the current patient. A treatment decision  $(A_t)$  is made based on the history and the context information. Record the post treatment health outcome  $(R_t)$  for the patient. **end for** 

In this section, we review the online learning literature on bandit problems with adversarial context. We will first see a naive one bandit per context approach, which does not make any assumption on the data generating process. We will also introduce the linear contextual bandit setting with linear assumption on the reward distribution. In the end, we introduce some variants of the standard contextual bandits that address sparsity or different types of outputs.

#### 2.2.1 One bandit per context

The simplest solution for contextual bandits is possibly one bandit per context (OBPC), which solves the MAB for each context separately. Given any algorithm that solves MAB, OBPC maintains a trajectory for each context. On each step t, if  $C_t = c$ , the agent applies the algorithm on the trajectory corresponding to c and selects an action. Combined with the EXP3 algorithms, which we call EXP3 per Context, OBPC achieves a regret bound of  $\sqrt{K \log(K) |\mathcal{C}|T}$ . OBPC does not make any further assumption on the data generating process  $\nu$ , so it only works when  $\mathcal{C}$  is finite and small. As pointed out in Lattimore and Szepesvári (2020), the awareness of the context information improves the performance over the non-contextual algorithm only when the context actually leads to a non-stationary environment and causes a large gap between the best context-aware policy  $\sum_{c \in \mathcal{C}} \max_{a \in [K]} \sum_{t \in [n]: X_t = c} R_{t,a}$  and non-context policy  $\max_{a \in [K]} \sum_{t=1}^n R_{t,a}$ .

#### 2.2.2 Bandits with expert advice

OBPC fails when the context space is large and each context does not receive enough samples. In practice, context space has some internal structure that allows information sharing among different contexts. For instance, there can be thousands of categorical variables in precision health applications and many of them may be irrelevant to the reward distribution. Thus, grouping context based on these variables can reduce the context space while keeping the same performance.

Let  $\mathcal{P} \subset 2^{\mathcal{C}}$  be a partition over context space, i.e. any  $P, P' \in \mathcal{P}, P \cap P' = \emptyset$  and  $\bigcup_{P \in \mathcal{P}} P = \mathcal{C}$ . The above structural information can be seen as restricting the policy space within

$$\Pi(\mathcal{P}) = \{\pi : \mathcal{C} \to [K]; \forall c, c' \in \mathcal{C} \text{ s.t. } c, c' \in P \text{ for some } P \in \mathcal{P}, \pi(c) = \pi(c') \}$$

In other words, we represent the structural information as a restricted policy space that is expected to be essentially smaller than the whole policy space. To allow more flexibility in structural information, we study the general restricted policy space  $\Pi \subset \{\pi : \mathcal{C} \mapsto [K]\}$ . The aim of agent is now to compete against a fixed class of policies, instead of competing against a set of actions which is a special case of the former.

The recommendations of the policies in the restricted set can be thought of as "expert advice". We can therefore design an algorithm to use such advice by extending the EXP3 algorithm. The extended algorithm EXP4 (Exponential weighting for Exploration and Exploitation with Experts) is given in Algorithm 8.

Algorithm 8 Exponential-weight Algorithm for Exploration and Exploitation with Experts (Exp4)

1: Input: T, K, hyper-parameter  $\eta \in (0, 1]$  and policy space  $\Pi$ 

- 2: Set the distribution  $Q_{1,\pi} = 1/|\Pi|$  for all  $\pi \in \Pi$
- 3: for t = 1, ..., T do
- Receive context  $C_t$ 4:
- For all  $i \in [K]$ , compute  $P_{t,i} = \sum_{\pi \in \Pi} Q_{t,\pi} \mathbb{1}\{\pi(C_t) = i\}$ Sample  $A_t$  from  $P_t$  and receive reward  $R_t$ Set  $\hat{R}_{t,i} = 1 \frac{\mathbb{1}\{A_t=i\}}{P_{t,i}}(1-R_t)$ 5:
- 6:
- 7:
- Propagate the rewards to the experts:  $\tilde{R}_{t,\pi} = \sum_{i \in [K]} \mathbb{1}\{\pi(C_t) = i\} \hat{R}_{t,i}$  for all  $\pi \in \Pi$ 8:
- Update  $Q_{t+1,\pi} = \exp(\eta \tilde{R}_{t,\pi}) Q_{t,\pi} / \sum_{\pi' \in \Pi} \exp(\eta \tilde{R}_{t,\pi'}) Q_{t,\pi}$  for all  $\pi \in \Pi$ 9:

10: end for

Following the idea of EXP3, EXP4 first uses the same importance-weighted estimator for each action, i.e.  $\dot{R}_{t,i} = 1 - \mathbb{1}_{\{A_t=i\}}(1-R_t)/P_{ti}$ . Then EXP4 estimates the reward for each expert (policy)

$$\tilde{R}_{t,\pi} = \sum_{i \in [K]} \mathbb{1} \left\{ \pi \left( C_t \right) = i \right\} \hat{R}_{t,i},$$

and keeps a distribution  $Q_{t,\pi}$  over the policy space. The distribution is updated using exponentially weighting of the rewards:

$$Q_{t+1,\pi} = \exp(\eta R_{t,\pi}) Q_{t,\pi} / \sum_{\pi' \in \Pi} \exp(\eta R_{t,\pi'}) Q_{t,\pi},$$

where  $\eta \in (0, 1]$  is the learning rate.

We define regret not with respect to the best action but with respect to the best expert:

$$\operatorname{Reg}_{T} = \mathbb{E}\left[\max_{\pi \in \Pi} \sum_{t=1}^{T} \boldsymbol{r}_{t}^{T} \pi(C_{t}) - \sum_{t=1}^{n} R_{t}\right],$$

where  $\mathbf{r}_t = (r_{t,i})_i$  is the K-dimensional vector of rewards of the individual actions and  $\pi(C_t)$  is the Kdimensional vector of the probability of selecting each action using policy  $\pi$  under context  $C_t$ .

**Theorem 9.** Let  $\eta = \sqrt{2 \log(|\Pi|/(TK))}$ . The regret of EXP4 satisfies

$$\operatorname{Reg}_T(\operatorname{EXP4}) \le \sqrt{2\log(|\Pi|)KT}.$$

*Proof.* Similarly to the way we show (13) in the proof in Theorem 4, we can also show that the for any  $\pi^*\in\Pi:$ 

$$\sum_{t=1}^{T} \tilde{R}_{t,\pi^*} - \sum_{t=1}^{T} \sum_{\pi \in \Pi} Q_{t,\pi} \tilde{R}_{t,\pi} \le \frac{\log(|\Pi|)}{\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \sum_{\pi \in \Pi} Q_{t,\pi} \left(1 - \tilde{R}_{t,\pi}\right)^2.$$
(19)

Let  $\hat{\mathbf{R}}_t = (\hat{R}_{t,i})_{i=1}^K$  and  $\tilde{\mathbf{R}}_t = (\tilde{R}_{t,\pi})_{\pi \in \Pi}$  denote the vector of  $R_{t,i}$ 's and  $\tilde{R}_{t,\pi}$ 's. Since the estimator is unbiased,  $\mathbb{E}_t[\hat{R}_t] = \mathbf{r}_t$  and letting  $\mathbf{E}^{(t)}$  be the  $|\Pi| \times K$  matrix with each row corresponding to a  $\pi \in \Pi$  that gives the row vector  $(\mathbb{1}\{\pi(C_t) = i\})_{i=1}^K$ , we have

$$\mathbb{E}_t[\tilde{\boldsymbol{R}}_t] = \mathbb{E}_t[\boldsymbol{E}^{(t)}\hat{\boldsymbol{R}}_t] = \boldsymbol{E}^{(t)}\mathbb{E}_t[\hat{\boldsymbol{R}}_t] = \boldsymbol{E}^{(t)}\boldsymbol{r}_t.$$

Taking expectation over both side of Eq. (19) and using the fact that  $Q_{t,\pi}$  is  $\mathcal{F}_t$ -measurable leads to

$$\operatorname{Reg}_{T} \leq \frac{\log(|\Pi|)}{\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \sum_{\pi \in \Pi} \mathbb{E}Q_{t,\pi} \left(1 - \hat{R}_{t,\pi}\right)^{2}.$$
(20)

Let  $\hat{Y}_{t,i} = 1 - \hat{R}_{t,i}$ ,  $y_{t,i} = 1 - r_{t,i}$  and  $\tilde{Y}_{t,\pi} = 1 - \tilde{R}_{t,\pi}$  for all  $i \in [K]$  and  $\pi \in \Pi$ . Define vectors  $\tilde{Y}_t$  and  $\hat{Y}_t$  similarly. Note that  $\tilde{Y}_t = E^{(t)}\hat{Y}_t$  and let  $A_{t,i} = \mathbb{1}\{A_t = i\}$ , which gives us  $\hat{Y}_{t,i} = A_{t,i}y_{t,i}/P_{t,i}$  and

$$\mathbb{E}_{t}\left[\tilde{Y}_{t,\pi}^{2}\right] = \mathbb{E}_{t}\left[\left(\frac{E_{\pi,A_{t}}^{(t)}y_{t,A_{t}}}{P_{t,A_{t}}}\right)^{2}\right] = \sum_{i=1}^{K} \frac{\left(E_{\pi,i}^{(t)}y_{t,i}\right)^{2}}{P_{t,i}} \le \sum_{i=1}^{K} \frac{E_{\pi,i}^{(t)}}{P_{t,i}}$$

Thus, by the definition of  $P_{t,i}$ ,

$$\mathbb{E}[\sum_{\pi \in \Pi} Q_{t,\pi} (1 - \tilde{R}_{t,\pi})] \le \left[\sum_{\pi \in \Pi} Q_{t,\pi} \sum_{i=1}^{K} \frac{E_{\pi,i}^{(t)}}{P_{t,i}}\right] = K.$$

Substituting into Eq. (20) gives us

$$\operatorname{Reg}_T \leq \frac{\log(|\Pi|)}{\eta} + \frac{\eta TK}{2} \leq \sqrt{2TK\log(|\Pi|)},$$

when  $\eta = \sqrt{2\log(|\Pi|)/(TK)}$ .

Now we take a look at the context space partition example introduced before. Assume we have a *M*-partition of the context space C. The regret bound is reduced from  $\sqrt{2|C|\log(K)KT}$  to  $\sqrt{2M\log(K)KT}$  with the structural information.

#### 2.2.3 Linear bandits and contextual linear bandits

One bandit per context method learns a distinct environment for each context, which fails when the context space is large or even infinite. Many healthcare applications, especially in mobile health, have continuous context including heart rate, sleeping time, etc. To model such applications with efficient algorithms, we need to assume some nice relationship between the context and environment. The simplest and most frequently used model of this kind is the contextual linear bandit.

Before we introduce the setting for contextual linear bandit, we first define linear bandit. In linear bandit, we do not directly model contexts but assume that the action space  $\mathcal{A}_t \subset \mathbb{R}^d$  (we will soon reveal the reason behind choosing a time varying action space) for some positive integer d, the reward on step t is sampled from a linear model:

$$R_t = \langle \theta^*, A_t \rangle + \eta_t, \text{ for some } \theta^* \in \mathbb{R}^d, \tag{21}$$

and the noise  $\eta_t$  is usually assumed to be subgaussian. Note that, depending on the choice of  $\mathcal{A}_t$ , linear bandits can subsume other bandit problems. For example, to capture the standard stochastic bandit environment with d actions, we set  $\mathcal{A}_t = \{e_1, \ldots, e_d\}$  for unit vectors  $(e_i)_i$ . Note that this choice of action set is not time varying. The flexibility to choose time varying action sets is critical to capture contextual bandits. We simply let  $\mathcal{A}_t = \{\psi(C_t, i) \in \mathbb{R}^d : i \in [K]\}$ , for some known function  $\psi : \mathcal{C} \times [K] \mapsto \mathbb{R}^d$ . A common setup (Li et al., 2010) considers one context per arm, where  $\mathcal{C} \subset \mathbb{R}^{d \times K}$  and the function  $\psi(c, i)$  extracts the *i*-th column of the matrix *c*. Another formulation assumes the same context  $c_t \in \mathbb{R}^d$  for every arm but assumes each arm *i* has a parameter  $\theta_i^* \in \mathbb{R}^d$ :  $R_t = \langle \theta_i^*, c_t \rangle + \eta_t$ , when arm *i* is chosen. Note that this is a special case of the general formulation in (21). We can write  $\theta^* = (\theta_1^{*T}, \ldots, \theta_K^{*T})^T$  and  $\mathcal{A}_t = \{(\mathbf{0}_{d(i-1)}, C_t^T, \mathbf{0}_{d(K-i)})\}_{i=1}^K \subset \mathbb{R}^{dK}$ , where  $\mathbf{0}_n$  is the all-zero row vector of dimension *n*.

Linear Upper Confidence Bound (LinUCB). The regret for linear bandit is defined as

$$\operatorname{Reg}_{T} = \mathbb{E}\left[\sum_{t=1}^{T} \max_{a_{t} \in \mathcal{A}_{t}} \langle \theta^{*}, a_{t} \rangle - \sum_{t=1}^{T} R_{t}\right].$$

Algorithm 9 Linear Upper Confidence Bound (LinUCB)

- 1: Input:  $\delta, d, (\mathcal{A}_t)_t \subset \mathbb{R}^d, (\beta_t)_t, \lambda.$
- 2: for t = 1, ..., T do
- 3: Calculate  $\hat{\theta}_t$  using (22).
- 4: Compute the confidence set  $\Theta_t = \{\theta \in \mathbb{R}^d : \|\theta \hat{\theta}_{t-1}\|_{V_{t-1}}^2 \leq \beta_t \}.$
- 5: Play  $A_t = \operatorname{argmax}_{a \in \mathcal{A}_t} \operatorname{max}_{\theta \in \Theta_t} \langle \theta, a \rangle.$
- 6: Receive reward  $R_t$  and update  $V_t$  using (23).

```
7: end for
```

We have seen UCB for the regular bandit environment as a powerful algorithm to balance the exploration and exploitation. In this section, we introduce a new UCB algorithm called LinUCB for linear bandit setting (Auer, 2002; Dani et al., 2008; Rusmevichientong and Tsitsiklis, 2010).

As the case in UCB, we follow the principle *optimism in the face of uncertainty (OFU)*. Instead of directly building confidence set over the mean function  $\mu_a$ , we build confidence set  $\Theta_t \subset \mathbb{R}^d$  over the linear coefficient  $\theta^*$  that satisfies the following two properties: 1)  $\Theta_t$  contains the optimal  $\theta^*$  with a high probability; 2) the set  $\Theta_t$  should be as small as possible.

For that purpose, we consider the regularised least-square estimator up to step t,

$$\hat{\theta}_t = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \left( \sum_{s=1}^t \left( R_s - \langle \theta, A_s \rangle \right)^2 + \lambda \|\theta\|_2^2 \right),$$
(22)

where  $\lambda > 0$  is the penalty factor. Equ. (22) gives

$$\hat{\theta}_t = V_t^{-1} \sum_{s=1}^t A_s R_s, \quad \text{where } V_0 = \lambda I_d \text{ and } V_t = V_0 + \sum_{s=1}^t A_s A_s^{\top}.$$
 (23)

Then the confidence set centering at  $\hat{\theta}_t$  is defined by

$$\Theta_{t+1} = \{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_t\|_{V_t}^2 \le \beta_{t+1} \},$$
(24)

where  $(\beta_t)_t$  is an increasing sequence of constants that controls the level of confidence with  $\beta_1 \ge 1$  as  $(V_t)_t$ also have increasing eigenvalues. Lemma 10 shows the correctness of the confidence set defined in (24). LinUCB chooses actions that maximize the expected reward over all the possible  $\theta$ 's in the confidence set. Details are given in Algorithm 9.

**Lemma 10.** Let  $\delta \in (0,1)$ . Then, with a probability at least  $1 - \delta$ , it holds that for all  $t \in \mathbb{N}$ ,  $\theta^* \in \Theta_t$  defined in (24) with any sequence  $(\beta_t)_t$  satisfying

$$eta_t \leq \sqrt{\lambda} \left\| \theta^* \right\|_2 + \sqrt{2 \log\left(rac{1}{\delta}
ight) + d \log\left(rac{\det V_t}{\lambda^d}
ight)}.$$

Now using the confidence set in Lemma 10, we can show the following regret bound under some mild assumption on the boundedness of the action space. We refer reader to Part V of Lattimore and Szepesvári (2020) for the proof of Lemma 10.

**Theorem 11.** Assuming  $\max_{t \in [n]} \sup_{a,b \in A_t} \langle \theta^*, a - b \rangle \leq 1$  and  $||a||_2 \leq L$  for all  $a \in \bigcup_{t=1}^n A_t$ , using the confidence set defined in Lemma 10, we have with a probability at least  $1 - \delta$ , the expected regret of LinUCB is bounded by

$$\operatorname{Reg}_T(LinUCB, \nu) \leq Cd\sqrt{T}\log(L/\delta), \text{ for some } C > 0.$$

*Proof.* We first analyze the regret on each step. Let  $\theta_t \in \Theta_t$  be the parameter that maximize the expected reward. Let  $A_t^*$  be the optimal action given context  $C_t$ . Then, using the fact that  $\theta^* \in \Theta_t$ , we have  $\langle \theta^*, A_t^* \rangle \leq \langle \theta, A_t \rangle$ . Using Cauchy-Schwarz inequality, we have

$$r_t = \langle \theta_*, A_t^* - A_t \rangle \le \langle \theta_t - \theta_*, A_t \rangle \le \|A_t\|_{V_{t-1}^{-1}} \|\tilde{\theta}_t - \theta_*\|_{V_{t-1}} \le 2\|A_t\|_{V_{t-1}^{-1}} \sqrt{\beta_t}.$$

Then the total regret is given by

$$\operatorname{Reg}_{T} \leq \sum_{t=1}^{T} 2 \|A_{t}\|_{V_{t-1}^{-1}} \sqrt{\beta_{t}} \leq \sqrt{T \sum_{t=1}^{T} r_{t}^{2}} \leq \sqrt{T \beta_{T} \sum_{t=1}^{T} \|A_{t}\|_{V_{t-1}^{-1}}^{2}}.$$
(25)

Using the elliptical potential lemma (Carpentier et al., 2020),  $\sum_{t=1}^{T} \left( 1 \wedge ||A_t||_{V_{t-1}}^2 \right) \leq 2 \log \left( \det V_T / \det V_0 \right)$ . The results follow by plugging this into (25) with the chosen  $\beta_T$  and the fact that

$$\det (V_T) = \prod_{i=1}^d \lambda_i \le \left(\frac{1}{d} \operatorname{trace} V_T\right)^d \le \left(\frac{\operatorname{trace} V_0 + TL^2}{d}\right)^d$$

where  $\lambda_1, \ldots, \lambda_d$  are eigenvalues of  $V_t$ .

LinUCB gives a regret bound that scales with d. However, this is not the optimal rate when the action set is finite. Auer (2002) proposed SupLinUCB which maintains a master algorithm that achieves a regret bound of  $\sqrt{dT \log(TK)}$  when the action space is finite.

LinUCB for contextual bandits. Chu et al. (2011) applied the LinUCB to the contextual bandit problems. To be specific, they assume a context and a unknown parameter  $\theta$  for each arm a and the reward is generated by

$$R_t = C_{t,A_t}^{\top} \theta_{A_t}^* + \eta_t$$
, where  $C_{t,a}, \theta_a^* \in \mathbb{R}^d$ .

Applying Theorem 11 gives us the regret bound of  $Kd\sqrt{T}$ . Li et al. (2010) showed a good performance of the contextual bandit with LinUCB on News Article Recommendation task.

#### 2.2.4 Sparse LinUCB

So far we have achieved regret bound of  $d\sqrt{T\log(T)}$ , which scales linearly with d. However, the bound can be vacuous for high-dimensional context ( $d \gg T$ ). Such high dimensional problems can occur in healthcare applications. For example, Bastani and Bayati (2020) incorporates the high-dimensional genetic profile and medical records with contextual bandits to design patient's optimal medication dosage. Following the idea of UCB, the key is to construct a confidence set that accounts for the sparsity. To tackle the problem, we introduce a powerful tool that converts online prediction problem to confidence set estimation, with which a UCB-style algorithm can be designed using any online predicting method and we can analyze its regret in a unified way.

**Online to Confidence Set Conversion.** There is a conversion from the online prediction to confidence sets. We consider online linear regression problem with a squared loss, which has been well-explored in the past decade (Huang et al., 2008; Javanmard and Montanari, 2014). The agent interacts with the environment in the following manner where in each round t:

- 1. The environment chooses a context  $C_t \in \mathbb{R}$  and  $A_t \in \mathbb{R}^d$  in an arbitrary way.
- 2. The value of  $A_t$  is revealed to the agent.

Algorithm 10 Online linear predictor UCB

1: Input:  $\delta \in (0, 1)$ , an online predictor with a regret bound  $B_t$ 

- 2: for t = 1, ..., T do
- 3: Receive action set  $\mathcal{A}_t$
- 4: Compute the confidence set  $\Theta_t$  using (26)
- 5: Play  $A_t = \operatorname{argmax}_{a \in \mathcal{A}_t} \operatorname{max}_{\theta \in \Theta_t} \langle \theta, a \rangle$  and receive reward  $R_t$
- 6: Feed  $A_t$  to the online linear predictor and obtain prediction  $\hat{R}_t$
- 7: Feed  $R_t$  and loss to the linear predictor as feedback.
- 8: end for
  - 3. The agent generate a prediction  $\hat{R}_t$ .
  - 4. The environment reveals  $R_t$  to the agent as well as the loss  $(R_t \hat{R}_t)^2$ .

The regret of an agent with respective to the best linear predictor is given by

$$\rho_T = \sum_{t=1}^T (R_t - \hat{R}_t)^2 - \inf_{\theta \subset \mathbb{R}^d} \sum_{t=1}^T (R_t - \langle \theta, A_t \rangle)^2.$$

**Theorem 12** (Online-to-Confidence-Set Conversion (Abbasi-Yadkori et al., 2012)). Let  $\delta \in (0,1)$  and assume that  $\theta^* \in \Theta$ , noise  $\eta_t$ 's are R-sub-Gaussian and  $\rho_t \leq B_t$ . Let

$$\beta_t(\delta) = 1 + 2B_t + 32R^2 \log\left(\frac{R\sqrt{8} + \sqrt{1 + B_t}}{\delta}\right)$$

Then

$$\Theta_{t+1} = \left\{ \theta \in \mathbb{R}^d : \sum_{s=1}^t (\hat{R}_s - \langle \theta, A_s \rangle)^2 \le \beta_t(\delta) \right\},\tag{26}$$

satisfies  $\mathbb{P}(\theta^* \in \Theta_{t+1} \forall t \in \mathbb{N}) \ge 1 - \delta$ .

Theorem 12 provides a tool for constructing confidence set for the prediction of any online learning algorithm with a regret guarantee. One can develop the Online linear predictor UCB (OLR-UCB) as shown in Algorithm 10.

Let  $\|\theta^*\|_2 \leq m_2$ . Using Theorem 12 and the standard regret decomposition for UCB-style algorithm, one can show the following regret bound:

$$R_T \le \sqrt{8dT \left(m_2^2 + \beta_{T-1}(\delta)\right) \log\left(1 + \frac{T}{d}\right)}.$$

Application to sparse online linear prediction. Note that our goal is to derive a regret bound with a lower dependency on the ambient dimension d. In the standard regression problem, learnability of high-dimensional data relies on the sparsity assumption which assumes the number of non-zero dimensions is less or equal to a positive integer s.

Assumption 1. (sparse parameter) The true parameter  $\theta^*$  satisfies  $\|\theta^*\|_0 \leq s$  for some s > 0.

To this end, we adopt SeqSEW (Gerchinovitz, 2011) as the online linear prediction algorithm, which achieves the online learning regret  $\rho_T = \mathcal{O}(s \log(T))$ . Combined with Theorem 12, we have the following result:

**Theorem 13.** With a high probability, the we bound the regret of OLR-UCB with SeqSEW as the online predictor by

$$\operatorname{Reg}_{T} = \tilde{\mathcal{O}}(\sqrt{dsT}). \tag{27}$$

This improves over the  $d\sqrt{T}$  by LinUCB in the previous section when  $s \ll d$ . As pointed out by Lattimore and Szepesvári (2020), for any algorithm, there exists an infinite action space such that  $\operatorname{Reg}_T = \Omega(\sqrt{dsT})$ . Thus, the bound in (27) is minimax optimal. Despite its optimality, our result still have the unavoidable dependence of  $\sqrt{d}$ . Bastani and Bayati (2020) derived a regret bound of  $\mathcal{O}(K(s \log(d) \log(T))^2)$ , when the action set is finite. More variants of high-dimensional sparse linear bandits are discussed in Hao et al. (2020).

## 2.3 Offline Learning

So far we have been discussing bandit problems in the *online learning* setting, where we are free to choose actions while collecting the dataset. In many real applications, collecting data in the online manner can be expensive and can cost a lot of time. In such situations we might want to evaluate a target policy using a large but fixed dataset that has been collected for years by a different policy. This setting is called *offline learning*.

Given a fixed contextual bandit environment, the value of a policy is defined as  $v_{\pi} = \mathbb{E}_{a_t \sim \pi(c_t)} r_t$  Formally, offline bandits, also referred as off-policy evaluation aims at evaluating the value  $v_{\pi}$  of a new policy using a sequence of interactions  $S_T = (c_t, a_t, r_t)_{t=1}^T$ , where  $c_t$ ,  $a_t$  and  $r_t$  are the context, action and reward respectively at step t. The actions are chosen by the data-generating policy  $\pi_D$ , which is also referred as behavior policy or logging policy. In this section,

One solution is to learn the reward and context distribution using a parametric or non-parametric method and then evaluate the new policy with the simulator by sampling context and rewards from the learnt distribution. Although this approach is straightforward, the modeling step is often very expensive and difficult, and more importantly, it often introduces modeling bias to the simulator, making it hard to justify reliability of the obtained evaluation results.

Li et al. (2011) proposed a simple algorithm (Algorithm 11) that uses an unbiased estimator of the policy value when the behavior policy picks arms uniformly at random. The method is to simply use all the interactions in the dataset that coincide with the action that the policy being evaluated would have taken. They show that  $\hat{v}_{\pi}$  is an unbiased estimator of  $v_{\pi}$  and with a high probability  $(\hat{v}_{\pi} - v_{\pi})^2 = \tilde{\mathcal{O}} (Kv_{\pi}/n)$ , where  $v_{\pi}$  is the true expected value of the policy as defined above. They also successfully applied the off-policy evaluation to the News Article Recommendation system. Mary et al. (2014) improved upon Algorithm 11 via bootstrapping techniques. They give a bootstrapped estimator for any quantile of the distribution of the estimated value  $\hat{v}_{\pi}$  that allows some evaluation of the uncertainty.

#### Algorithm 11 Policy Evaluator

1: Input: a policy  $\pi$ ; stream of interactions  $S_T$  of length T2:  $h_0 \leftarrow \emptyset$ ,  $\hat{V}_{\pi} \leftarrow 0$  and  $N \leftarrow 0$ 3: for t = 1, ..., T do 4: Get the *i*-th interaction  $(c_i, a_i, r_t)$  from  $S_T$ 5: if  $\pi(h_{t-1}, c_t) = a_t$  then 6:  $h_t \leftarrow h_{t-1} \cup (c_t, a_t, r_t)$ 7:  $\hat{V}_{\pi} \leftarrow \hat{V}_{\pi} + r_t$  and  $N \leftarrow N + 1$ 8: end if 9: end for 10: Output:  $\hat{v}_{\pi} = \hat{V}_{\pi}/N$ 

Minimax rate lower bound. To understand how good our algorithms perform, we start with establishing a minimax lower bound that characterizes the inherent hardness of the off-policy value estimation problem. Let  $R_{\max}, \sigma \in \mathbb{R}^+$ . Define the class of reward distributions  $\mathcal{V}(\sigma, R_{\max})$  with bounded mean and variance as

$$\mathcal{V}(\sigma, R_{\max}) \coloneqq \{ \nu = (P_{a,c})_{a,c \in \mathcal{A} \times \mathcal{C}} : 0 \le \mathbb{E}_{R \sim P_{a,c}}[R] \le R_{\max} \text{ and } \operatorname{Var}[R \mid a, c] \le \sigma^2, \ \forall a, c \in \mathcal{A} \times \mathcal{C} \}.$$

Any estimator  $\hat{v}$  is a function that maps  $(\pi, \pi_D, S_T)$  to an estimate of  $v_{\pi}$ . Let the context  $C \sim \lambda$  with a density relative to Lebesgue measure. The minimax rate is defined as

$$M_n(\pi, \lambda, \pi_D, \sigma, R_{\max}) \coloneqq \inf_{\hat{v}} \sup_{\nu \in \mathcal{V}(\sigma, R_{\max})} \mathbb{E}[(\hat{v}(\pi, \pi_D, S_T) - v_{\pi})^2],$$

where the expectation is taken over the distribution of the dataset  $S_T$  that depends on  $(\pi, \lambda, \pi_D)$ .

A critical value that determines the hardness of the problem is

$$V_1 \coloneqq \mathbb{E}_{C \sim \lambda, A \sim \pi_D(\cdot, C)} \left[ \pi^2(A, C) / \pi_D^2(A, C) \right].$$

Wang et al. (2017) (Corollary 1) gives the following lower bound for  $R_n$ :

$$M_n(\pi, \lambda, \pi_D, \sigma, R_{\max}) \ge \frac{V_1(\sigma^2 + R_{\max}^2)}{700T}.$$
(28)

**Importance sampling estimator.** Wang et al. (2017) give some analysis on the minimax rate of the offpolicy evaluation given a known data-generating policy, denoted by  $\pi_D$ . We slightly abuse the notation and let  $\pi(a, c)$  denote the probability of choose *a* using policy  $\pi$  under the context *c*. They consider Importance Sampling (IS) estimator (Charles et al., 2013) and Weighted Importance Sampling (WIS) estimator given by

$$\hat{v}_{\text{IS},\pi} = \frac{1}{T} \sum_{t=1}^{T} \frac{\pi(a_t, c_t)}{\pi_D(a_t, c_t)} r_t \quad \text{and} \quad \hat{v}_{\text{WIS},\pi} = \sum_{t=1}^{T} \frac{\frac{\pi(a_t, c_t)}{\pi_D(a_t, c_t)}}{\sum_{t'=1}^{T} \frac{\pi(a_{t'}, c_t)}{\pi_D(a_{t'}, c_t)}} r_t, \text{ respectively.}$$

It is shown in Dudík et al. (2014) that the IS estimator is unbiased and it achieves the minimax rate in (28) when n is sufficiently large. The WIS estimator is biased with a lower variance. It also enjoys the minimax optimality up to a logarithmic constant.

**Regression estimator.** Another common method, so called regression estimator or plug-in estimator, simply learns the reward distribution and plugs it in. Let  $\hat{r} : \mathcal{A} \times \mathcal{C} \mapsto \mathbb{R}$  is an estimator for the mean reward. The regression estimator is given by

$$\hat{v}_{\text{Reg},\pi} := \frac{1}{T} \sum_{t=1}^{T} \sum_{a} \pi(a, c_t) \hat{r}(a, c_t).$$

When the context space and action space are finite. One can use the simple sample average estimator  $\hat{r}(a,c) = (\sum_{t=1}^{T} \mathbb{1}(a_t = a, c_t = c)r_i)/(\sum_{t=1}^{T} \mathbb{1}(a_t = a, c_t = c))$ . Interestingly, one can write the regression estimator using sample average as

$$\hat{v}_{\text{Reg}} = \frac{1}{T} \sum_{t=1}^{T} \frac{\pi(a_t, c_t)}{\hat{\pi}_D(a_t, c_t)} r_t, \text{ where } \hat{\pi}_D(a, c) = \frac{\sum_t \mathbb{1}(a_t = a, c_t = c)}{T}.$$

There is a strong connection between IS and regression estimator. That is regression estimator simply replaces the  $\pi_D$  with its empirical estimator. Regression estimator is biased while its variance is normally lower than IS (Dudík et al., 2011; Li et al., 2015). Regression estimator is also shown to be minimax optimal when n is sufficiently large. The simple plugin method using sample average is not applicable when context space is infinite. When more information on how rewards depend on contexts and actions, we can posit a parametric or non-parametric model of  $\mathbb{E}[R \mid C, A]$  and fit it to obtain an estimator. Due to the difficulties of having a good estimation  $\hat{x}$ , which may suffer model-misspecification or high variance, pure regression estimator only works well in problem with finite actions. For more details, we refer reader to a more recent paper (Kim et al., 2021). **Doubly robust estimator.** Based on the above discussions about IS and Regression Estimator, it is natural to combine the both. Doubly Robust (DR) (Bang and Robins, 2005; Dudík et al., 2011; Jiang and Li, 2016) which is a combination of regression and IS estimator, and can achieve the low variance of regression and no (or low) bias of IS. The method is named by Doubly Robust because it is accurate if at least one of the estimator is accurate. DR estimator is defined as

$$\hat{v}_{DR} = \frac{1}{T} \sum_{t} \left[ \frac{\pi(a_t, c_t)}{\pi_D(a_t, c_t)} (r_t - \hat{r}(a_t, c_t)) + \sum_{a} \pi(a, c_t) \hat{r}(a, c_t) \right].$$

Informally, the estimator uses  $\hat{r}(a, c)$  as a baseline and corrects the baseline when the action more likely to be sampled from the target policy. DR enjoys both the low variance of Regression estimator and the low bias of the IS estimator and also matches the minimax lower bound in 28. Wang et al. (2017) further introduces the SWITCH estimator that takes account into the scale of the importance weight. When the importance weight is large, it switches to the regression estimator, which significantly reduces the variance while introducing only little bias. It performs better in the numerical experiments.

**Non-asymptotic regime.** All of the above estimators achieve minimax optimality when n is sufficiently large. Ma et al. (2021) discussed regime when the sample size is not large. They showed that there is a fundamental statistical gap between algorithm with and without the knowledge of the behavior policy in this regime. They proposed a *competitive ratio* that measures the MSE of an algorithm with unknown behavior policy relative to the lower bound of all algorithms with known behavior policy. Competitive ratio is always greater than 1. A competitive ratio closer to 1 indicates that the algorithm can performance as good as the algorithm with known behavior policy. Ma et al. (2021) showed that regression estimator has minimax optimal competitive ratio of the rate is  $\mathcal{O}(K)$ , where K is the number of actions. That means an algorithm without the knowledge of the behavior policy has to pay multiplicative factor of K compared to that with the knowledge.

## 3 Advanced Topics

We now review a selection of advanced topics in bandit algorithms that are relevant to applications in healthcare.

## 3.1 Non-stationarity

We have discussed two basic and extreme cases of bandit theory: stochastic bandit models and adversarial bandit models. In the former, the reward distribution for every arm is static whereas in the latter, it can change arbitrarily over time. In this section, we describe a more practical setting that sits in the "middle" of above two extreme models: non-stationary bandits. Specifically, the reward distributions change over time, but the amount of changes cannot be as arbitrary as the adversarial setting.

For a non-stationary multi-armed bandit environment  $\nu$ , we define the reward mean for arm a at time t by  $\mu_{a,t}(\nu)$  or  $\mu_{a,t}$  for simplicity. Then the dynamic regret is

$$\operatorname{Reg}_{T} = \sum_{t=1}^{T} \max_{a} \mu_{a,t} - \mathbb{E}\left[\sum_{t=1}^{T} \mu_{A_{t},t}\right].$$

Unlike the regret definition for the adversarial setting that is with respect to the best single arm, dynamic regret is defined with respect to the sequence of best arms at each time step. There are two popular constraints that quantify the reward changes. One is the total count of changes in the mean reward that occur before time T:

$$S_T = 1 + \sum_{t=2}^T \mathbb{1}\{\mu_{a,t-1} \neq \mu_{a,t}, \text{ for some } a\}.$$

Settings where the above quantity is assumed to be small are also called piece-wise stationary. Another constraint quantifies the total variation for the reward mean:

$$B_T = \sum_{t=2}^{T} \max_{a} |\mu_{a,t} - \mu_{a,t-1}|.$$

**Piecewise Stationary Setting.** The first algorithm designed for the non-stationary MAB under finite number of reward mean changes is Exp3.S (Auer et al., 2002), which is a variant of the Exp3 algorithm. The expected regret of Exp3.S is  $O(S_T \sqrt{KT} \log(KT))$  in general, but can be improved to  $O(\sqrt{S_T KT} \log(KT))$ if  $S_T$  is used to tune the algorithm parameters. Later on, Garivier and Moulines (2011) proved that no policy can achieve problem-dependent regret smaller than  $O(\sqrt{T})$  in the non-stationary case, and Exp3.S is thus optimal up to logarithmic factors of T. Garivier and Moulines (2011) also studied two UCB-based algorithms: discounted UCB (D-UCB) which was first proposed by Kocsis and Szepesvári (2006) and sliding-window UCB (SW-UCB). Their main idea is to encourage using more recent data in UCB reward estimations. When S is given, both algorithms achieve  $O(K\sqrt{ST}\log T)$  regret. Another body of works explore the idea of monitoring the reward distributions by change-detection methods and reset the bandit algorithm accordingly. Using this idea, Monitored-UCB (M-UCB) (Cao et al., 2019) combines the UCB algorithm with a change-point detection component based on running sample means over a sliding window and achieves  $O(\sqrt{KST}\log T)$ regret under certain assumption on the reward mean change. In same year, Auer et al. (2019) proposed an action elimination algorithm ADSWITCH that again detects changes in the mean reward and restarts the learning algorithm accordingly. Without knowing S nor making other assumptions, ADSWITCH achieves  $O(\sqrt{KST \log T})$  regret.

Bounded Total Variation Setting. Comparing with piecewise stationary setting, bounded total variation is a softer constraint. Here nature has the power to change the reward at every round but only up to a total amount limit  $B_T$ . Besbes et al. (2014) proposed the Rexp3 policy that uses the famous Exp3 algorithm as a subroutine and restarts it at every batch. Rexp3 with a batch size tuned by  $B_T$  has  $O((B_T K \log K)^{1/3} T^{2/3})$  regret that nearly matches the lower bound of  $\Omega(KB_T)^{1/3} T^{2/3})$  in this setting (Besbes et al., 2014).

The non-stationary setup has also studied in linear bandits and contextual bandits. We briefly describe the latter one as an example. Define the distribution of the context-reward pairs  $\mathcal{C} \times [0,1]^K$  by  $\mathcal{D}_1, \ldots, \mathcal{D}_T$ . At each round, the environment samples  $(c_t, x_t)$  and reveals  $c_t$  to the agent, then the agent picks an arm  $A_t \in [K]$  and observes  $x_t(A_t)$ . For a fixed set of policies  $\Pi$  that contains mappings:  $\mathcal{C} \to [K]$ , the dynamic regret is defined as:

$$\operatorname{Reg}_T = \sum_{t=1}^T \max_{\pi \in \Pi} \mathbb{E}_{(c,x) \sim \mathcal{D}_t}[x(\pi(c))] - \sum_{t=1}^T x_t(A_t).$$

Similar to the MAB setting, there are two ways to measure the non-stationary of the environment: total number of changes and the total variation in below.

$$S = 1 + \sum_{t=2}^{T} \mathbb{1}\{\mathcal{D}_t \neq \mathcal{D}_{t-1}\}$$
$$B_T = \sum_{t=2}^{T} \|\mathcal{D}_t - \mathcal{D}_{t-1}\|_{TV}.$$

Chen et al. (2019) proposed ADA-ILTCB<sup>+</sup> algorithm that is parameter-free, efficient and achieves optimal regret  $\widetilde{O}(\min\{\sqrt{K(\log|\Pi|)ST}, \sqrt{K(\log|\Pi|)T} + (K\log|\Pi|B_T)^{1/3}T^{2/3}\})$  by randomly entering replay phases to detect non-stationarity. More recently, a generic reduction has been studied that allows one to convert certain algorithms for the stationary setting into algorithms for the non-stationary setting (Wei and Luo,

2021). What is nice about this work is that the resulting algorithms often *simultaneously* achieve optimal guarantees in terms of both S and  $B_T$  without prior knowledge of any of these parameters.

## 3.2 Robustness

So far, we have discussed stochastic and adversarial bandits. Stochastic bandit is an ideal setting, since it assumes a fixed reward distribution. Adversarial, on the other hand, considers a too extreme scenario. In real-world applications, some of the observed rewards may be corrupted resulting in deviation from purely stochastic behavior. For instance, there might be clerical errors while creating electronic health records or sensor errors in recording health status from wearable devices. If the reward is self-reported by patients, there can be corruptions due to mistakes, lack of attention, boredom, etc. In this section, we consider algorithms that are *robust* to possible corruptions. Note that this setting is between purely stochastic and adversarial but in a different way than the non-stationary bandit discussed in Section 3.1.

**Fraction corruption model.** The fraction corruption model limits the fraction of the total number of rounds that an adversarial corruption can happen. Assume corruptions happen with a probability  $\eta \in [0, 1]$ . The adversary closely follows the progress of arm pulling and reward generation. At each time step t, after the algorithm has decided to pull an arm  $A_t$ , the adversary first decides whether to corrupt this arm pull or not by performing a Bernoulli trail  $Z_t \in \{0, 1\}$  with a mean  $\eta \in [0, 1]$ . Then it generates a corruption  $\zeta$  arbitrarily. After this, the "clean reward"  $R_t^*$  is generated by the environment. The reward a player receives at step t is

$$R_t = \mathbb{1}\{Z_t = 0\}R_t^* + \mathbb{1}\{Z_t = 1\}\zeta_t.$$

Since the rewards may be contaminated, directly running the algorithms built for stochastic bandits can result in degraded performance. Several papers have proposed variants of UCB that use a more robust estimate. Kapoor et al. (2019) proposed RUCB-TUNE (Robust UCB) which makes two crucial changes to the classical UCB: RUCB-TUNE uses the median and a tuned variance estimate instead of the simple sample mean and variance to construct its upper bound. The algorithm requires an upper bound  $\eta_0$  on the corruption rate  $\eta$ . Their analyses are built on the environment with Gaussian reward distributions and they have a regret bound of  $\mathcal{O}((1 - \eta)\sqrt{KT \log(T)} + \eta_0(\mu^* + B)T)$ , where B is an upper bound on the scales of corruptions  $|\zeta_t|$  and  $\mu^*$  is the maximum expected reward of all the arms. They provide another algorithm RUCB with a similar result. RUCB does not require the knowledge of  $\eta_0$  but it requires an upper bound on the variance of rewards for all the arms.

Niss and Tewari (2020) considers a stronger requirement that all the arms have no more than  $\eta$ -fraction of corruptions. They proposed crUCB (contamination robust-UCB), which mimics RUCB in the general framework. The only difference is that they allow a variety of robust mean estimates including  $\alpha$ -trimmed mean, mean estimates that ignores the largest  $(1 - \alpha)$ -fraction of data,  $\alpha$ -shorth mean. Their algorithms are evaluated by uncontaminated regrets, the regret between true rewards of the best possible actions and the algorithm-chosen actions:

$$\operatorname{Reg}_{T} = \max_{a \in \mathcal{A}} \mathbb{E}[\sum_{t=1}^{T} R_{a}^{*}(t) - R_{A_{t}}^{*}(t)],$$

If the contamination fraction is small enough, both crUCB using  $\alpha$ -shorth mean and  $\alpha$ -trimmed mean achieve uncontaminated regret of  $\sqrt{KT \log(T)} + \sum_a \Delta_a$ , where  $\Delta_a$  is the gap between the expected reward of a and the optimal action.

**Budget-bounded corruption model.** In budget-bounded corruption model (Lykouris et al., 2018), all the rounds can be possibly corrupted but the total deviation between the "clean rewards" and the corrupted rewards is bounded. We define the total corruption budget as

$$B_T = \sum_{t=1}^T \|R_t^* - R_t\|.$$

Gupta et al. (2019) proposed BARBAR (Bandit Algorithm with Robustness: Bad Arms get Recourse) based on the arm elimination algorithm discussed in Section 2 with some crucial modifications that make it robust to corruptions. Gupta et al. (2019) show that their algorithm achieves a gap-dependent regret bound of

$$\operatorname{Reg}_{T} = \mathcal{O}\left(KB_{T} + \sum_{a \neq a^{\star}} \frac{\log T}{\Delta_{a}} \log\left(\frac{K}{\delta} \log T\right)\right).$$

without knowing the value of C.

### 3.3 Dealing with Constraints

In standard bandit settings, the only criterion to evaluate the performance of an algorithm is the regret (either cumulative or simple regret). In this section, we introduce constrained bandit problems, where the agent also needs to satisfy certain constraints while minimizing its regret. Depending on the actual problem's demands, different types of constraints have been considered. For healthcare applications, two types of constraints frequently arise in practice. First, delivering an intervention might consume resources such as doctor's time, patient's attention, phone battery and power, etc. We might therefore want to maximize rewards subject to a budget on resource consumption. Second, we might be worried that uncontrolled exploration might yield performance that is substantially worse than an existing standard of care. With these concerns in mind, in this section we consider the following two settings: bandits with knapsacks and conservative bandits.

**Bandits with Knapsacks (BwK).** The name of *bandits with knapsacks* comes from the well-known knapsack problem in combinatorial optimization that studies packing items into a fixed-size knapsack. In a knapsack problem, each item has its value and size. Its ultimate goal is to fill the knapsack with as large value as possible. The *bandits with knapsacks* then focus on a stochastic online version of the knapsack problem. Many practical scenarios can be captured by this framework. For example, in medical trials, physicians may be limited by the cost of treatment materials while optimizing the health condition of the patients. In online recommendations, the website designer may be constrained by the advertisers' budgets while maximizing its profit. We follow the terminologies and notations in Badanidiyuru et al. (2018) to formally describe this setting and their proposed algorithms.

An agent is given an action set  $\mathcal{A}$  and there are d resources being consumed. At every round, the agent selects an action  $A_t \in \mathcal{A}$  and observes a reward  $X_t$  along with a d-dimensional resource consumption vector  $\mathbf{c}_t$ . For each resource  $i \in [d]$ , the total consumption should not exceed the pre-specified budget  $B_i$  at any round. The agent stops immediately when the total consumption of certain resource exceeds its budget. Denote the stopping time by  $\tau$ , the goal is to maximize the total reward until time  $\tau$ . The worst-case regret for an algorithm is defined as the difference between the benchmark total reward OPT<sup>4</sup> and the algorithm's expected total reward until  $\tau$ .

To solve the problem, Badanidiyuru et al. (2018) proposed two algorithms. The first one is a primal-dual algorithm called PrimalDualBwK. At every round, it estimates the upper confidence bound for the expected reward and the lower confidence bound for the resource consumption for each arm. Then PrimalDualBwK plays the most "cost-effective" arm, i.e., the one with highest ratio of the expected reward's upper bound to the expected cost. The regret of PrimalDualBwK is proved to be  $\tilde{O}\left(\sqrt{|\mathcal{A}|\text{OPT}} + \text{OPT}\sqrt{|\mathcal{A}|/B}\right)$ , where  $B := \min_i B_i$ . Note that without the resource constraints, the regret bound becomes  $\tilde{O}\left(\sqrt{|\mathcal{A}|\text{OPT}}\right) = \tilde{O}\left(\sqrt{|\mathcal{A}|T}\right)$  that is optimal up to logarithmic factors since the only constraint is the time horizon T and we can simply set B = T in the general regret formula. The second algorithm is called BalancedExploration. Its design principle is to explore as much as possible while avoiding obviously sub-optimal strategies. They

 $<sup>^{4}</sup>$ An optimal dynamic policy that maximizes the expected total reward given prior knowledge on all latent distributions such as the reward distribution and the cost consumption for each action.

show that the regret of BalancedExploration is  $\tilde{O}\left(\sqrt{d|\mathcal{A}|\text{OPT}} + \text{OPT}\sqrt{d|\mathcal{A}|/B}\right)$ . Even though its regret has worse dependence on d comparing with PrimalDualBwK, BalancedExploration performs better in some special cases.

Comparing with the general constraint in above knapsack framework, researchers have also considered more specific settings. For example, other than the time horizon, there is a single resource with deterministic consumption and different arms consume the resource at different rates (Tran-Thanh et al., 2010, 2012; Ding et al., 2013; Xia et al., 2015; Zhou and Tomlin, 2018). Such problems are often called *budgeted bandits*. Agrawal and Devanur (2014) generalized the BwK model by allowing arbitrary concave rewards and convex constraints. Furthermore, similar constrained bandit problems are also studied in settings that includes contextual bandits (Agrawal and Devanur, 2014; Wu et al., 2015; Agrawal and Devanur, 2016) and even adversarial bandits (Sun et al., 2017; Immorlica et al., 2019).

**Conservative Bandits.** It is well-known that standard bandit algorithms often explore wildly in their early stages, so their regret over initial rounds can be very high. However, safety is one of the crucial concerns for designing adaptive experiments, so the early stage high regret cannot be tolerated in some high stakes scenarios. For instance, the UCB algorithm can assign treatments almost randomly to a patient during the first several rounds, which may cause severe problems for the patient's overall health. A company may also not be able to withstand the extremely low revenue initially if its operation is in great need of cash flow. In these examples, the real situation is that the physician or the company already have their favorite policies that operate well. They would like to explore new strategies to optimize their treatment performance or revenue while maintaining their performance to not be significantly worse than a baseline, uniformly over time. *Conservative bandits* (Wu et al., 2016) exactly model this problem. We follow their terminology and notation to describe the problem framework.

We use  $\{0, 1, \ldots, K\}$  as the indices of actions, in which the arm indexed by 0 correspond to the default action (the agent's typical strategy) and the other arms are the alternatives to be explored. The agent selects an arm  $A_t \in \{0, 1, \ldots, K\}$  in round t. Denote  $R_{t,i}$  as the random reward received at time t after playing arm i, then the regret and pseudo-regret are defined as  $\operatorname{Reg}_T = \max_{i \in \{0,1,\ldots,K\}} \sum_{t=1}^T R_{t,i} - R_{t,A_t}$ and  $\widetilde{\operatorname{Reg}}_T = \max_{i \in \{0,1,\ldots,K\}} \sum_{t=1}^T \mu_i - \mu_{A_t}$ . To ensure that the agent performs as good as his or her usual strategy, the algorithm needs to satisfy:

$$\sum_{s=1}^{t} R_{s,A_s} \ge (1-\alpha) \sum_{s=1}^{t} R_{s,0},$$
(29)

where  $0 < \alpha \leq 1$ . Above constraint guarantees that the reward collected by the agent is at least  $(1 - \alpha)$  fraction of the reward from simply playing arm 0. The objective of conservative bandits algorithms is to minimize regret while satisfying (29) for all t. To solve the problem, (Wu et al., 2016) proposed a novel algorithm called conservative UCB which is based on UCB with the novel twist of maintaining (29) being satisfied. Basically, the agent follows the UCB suggestion if (29) can be satisfied from their estimation, otherwise, the agent switches to the conservative arm 0. They show that conservative UCB achieves pseudo-regret  $\widetilde{\text{Reg}}_T = \widetilde{O}(\sqrt{KT} + \frac{K}{\alpha\mu_0})$ , where  $\mu_0$  denotes the expected reward for arm 0. For more details, including gap-dependent regret and regret analysis for the adversarial setting, we refer the reader to Wu et al. (2016).

The concept of conservative bandits, namely performing as good as a baseline uniformly over time, is quite general. Recently, it has been generalized to more settings, e.g., conservative contextual bandits (Kazerouni et al., 2016; Garcelon et al., 2020b) and conservative reinforcement learning (Garcelon et al., 2020a).

#### 3.4 Fairness

Algorithmic fairness has become an increasingly important topic in machine learning research (Barocas et al., 2021). Fairness is also a concern in healthcare applications due to the biases in data collection and algorithmic design (Paulus and Kent, 2020). Using healthcare cost as a proxy for healthcare needs leads to biased risk scores that hurt Black patients (Obermeyer et al., 2019). Using genetic datasets collected

mostly from patients of European ancestry leads to biased genetic risk scores on patients with non-European ancestries (Martin et al., 2019). Unfairness can also be rooted in algorithmic design. Algorithms that simply maximizing user responses can be unfair in how they allocate exposure to the items in recommendation systems (Singh and Joachims, 2018). In this section, we review several popular definitions of fairness in bandit literature and discuss how to calibrate unfairness.

Two types of fairness. Though there remains little agreement about what "fairness" should mean in different contexts, the literature can be divided into two broad families: those that target *group* fairness and those that target *individual* fairness. Group fairness requires the algorithm to maintain fairness across different demographic groups (say by race or gender) that are supposed to be treated equally. As we mentioned above, biased healthcare algorithms may hurt patients from certain demographic groups due to low representative. Individual fairness, on the other hand, asks for some constraints on the individual level, including two slightly different types: fairness through awareness, which requires that similar individuals be treated similarly and meritocratic fairness, which requires that less qualified individuals not be favored over more qualified individuals. Individual fairness in healthcare is also important: for two patients with similar demographic information or physical condition, we shall not develop an algorithm that performs much better for one patient while useless (or even harmful) for the other.

In the context of bandits, the literature can be further divided into those that model individuals as arms in the MAB setting and those that models individuals as contexts in the contextual bandits setting. We give some examples of fairness definitions from each category.

**Individual fairness.** We start from individual fairness in the MAB setting with each arm representing an individual. Joseph et al. (2016) proposed the use of meritocratic fairness. Formally, an algorithm is said to be  $\delta$ -fair if over T time steps, , for any pair of arms a, a', and any round t,

$$\pi_t(a) > \pi_t(a')$$
 only if  $\mu_a > \mu_{a'}$ 

with probability at least  $1 - \delta$ . This requires the algorithm to give equal probability to choose two arms unless one there is strong evidence that  $\mu_a > \mu_{a'}$ . Joseph et al. (2016) developed an algorithm that is  $\delta$ -fair based on the traditional UCB algorithm. In specific, they maintain a set of arms that chains to the arm with highest UCB. All the arms in the set are given the same probability to be selected. Their regret bound is  $\mathcal{O}(\sqrt{K^3T \ln \frac{TK}{\delta}})$ , with a higher cubic dependence on K, which is the price they pay to give equal probability to all the arms in the set. More generally, Liu et al. (2017) considers a fairness through awareness. An algorithm is said to be  $(\epsilon_1, \epsilon_2, \delta)$ -smooth fair <sup>5</sup>, if for any pair of arms a, a' and any t, with a probability at least  $1 - \delta$ ,

$$|\pi_t(a) - \pi_t(a')| \le \epsilon_1 |\mu_a - \mu_{a'}| + \epsilon_2.$$

They introduced a fairness regret to quantify the violation that occurs when the arm with the highest reward realization at a given time is not selected with the highest probability. Formally, they have

$$\operatorname{Reg}_{T} = \sum_{t=1}^{T} \mathbb{E} \left[ \sum_{i=1}^{K} \max\{P^{*}(i) - \pi_{t}(i), 0\} \right],$$

where  $P^*(i)$  is the probability that arm *i* has the highest regret. They developed a modified Thompson Sampling algorithm with an initial uniform exploration phase that achieve fairness regret of  $\mathcal{O}((KT)^{2/3})$ . Wang et al. (2021) also aims at the fairness through awareness by comparing to the optimal fair policy  $\pi^*$ ,

$$\frac{\pi^*(a)}{\pi^*(a')} = \frac{f(\mu_a)}{f(\mu_{a'})}$$
 for some metric function  $f$ .

 $<sup>^{5}</sup>$ Note that Liu et al. (2017) considers a general divergence function for the distribution of action selection and rewards, we introduce a special case here for a cleaner form.

They in turn target at the fairness regret, the cumulative  $L_0$  norm of  $\pi^* - \pi_t$ . Other individual fairness definition that simply requires a minimum number of pulls for each arms is considered by Patil et al. (2020) and Chen et al. (2020).

**Group fairness.** For group fairness, one may divide arms into n groups  $G_1, \ldots, G_n$ . Schumann et al. (2019) proposed an algorithm ensuring that the probability of pulling an arm does not change based on group membership:

P(pull  $a \mid a \in G_i$ ) = P(pull  $a \mid a \in G_j$ ),  $\forall i, j < n$  and  $a \in \mathcal{A}$ .

In other words, one cannot prefer one arm over the other one based on the group information.

Some literature treats each context  $c_t$  as an individual and aims at the fairness among different contexts. (Huang et al., 2020) divides the context space into two groups, a privileged group  $G^+$  and a protected group  $G^-$ . A group-level cumulative mean reward is defined as

$$\bar{R}^G = \frac{1}{|T_G|} \sum_{t \in T_G} R_t$$
, where  $T_G$  is the set of all rounds with  $g_t$  in group  $G$ .

Their group fairness requires that  $|\mathbb{E}[\bar{R}^{G^+} - \bar{R}^{G^-}]| \leq \tau$ , where  $\tau \in \mathbb{R}^+$  reflects the tolerance degree of unfairness. Assuming a linear bandit setting, they proposed a Fair-LinUCB algorithm, that penalize the arms that are unfair by decreasing the corresponding UCB values.

**Counterfactual fairness.** Another line of work (Kusner et al., 2017) models fairness through causal inference. Their fairness definition is normally referred to as Counterfactual Fairness. In the group fairness, we require the distributions of the value of interest are the same across different demographic groups. However, what we really want is that the demographic group information does not cause the unfairness, which can not be fully inferred from the group fairness definition. Counterfactual fairness, on the other hand, requires the distribution of the value of interest does not change when the demographic group were changed from one to the other while keeping all the other context variable unchanged.

General fairness metric. All the above methods requires certain fairness definition, it is natural to ask whether there is an universal algorithm that adapts to various fairness definition. In real-applications, we may not know the exact form of fairness metric. Instead, we may know whether certain policy violates fairness condition. Instead of studying a specific fairness metric, Gillen et al. (2018) developed an algorithm that achieves fairness through a fairness oracle, which tells the algorithm whether fairness is violated. This also allows great flexibility in the choice of fairness metric. At each round, each arm  $i \in [K]$  is given a context  $C_{t,i} \in \mathbb{R}^d$ . The probability of selecting action i at the step t is denoted by  $\pi_{t,i}$ , which is a vector denoting the probability of selecting. We denote the vector of context (policy) at round t by  $c_t$  ( $\pi_t$ ). The fairness oracle is define as following. Intuitively, it tells the algorithm, which pairs of actions violate the fairness metric.

**Definition 4.** Let  $\Delta([K])$  be the set of all distributions over [K]. Given a user-specified distance function d, a fairness oracle  $O_d$  is a function  $O_d : \mathbb{R}^{d \times K} \times \Delta([K]) \to 2^{[K] \times [K]}$ , defined such that

$$O_d(c_t, \pi_t) = \{(i, j) : |\pi_{t,i} - \pi_{t,j}| > d(C_{t,i}, C_{t,j})\}$$

They developed an algorithm that only has access to the fairness oracle. The algorithm has the number of fairness violations depending only logarithmically on T and a regret bound of the optimal rate  $\mathcal{O}(\sqrt{T})$  with respect to the best fair policy.

### 3.5 Benefiting from Causal Knowledge

Causal bandit is an example of a structured bandit problem where actions are composed of interventions on variables of a causal graph. According to the underlying causal dynamics, performing an intervention can help learn another intervention's reward distribution. Algorithms for causal bandit problems exploit the causal dependency among interventions to reduce the regret or sample complexity.

Bareinboim et al. (2015) was the first to connect bandit problems with causal models. One of the main issues that make causal inference hard is the existence of confounders defined below, which can confuse correlation and causation. Bareinboim et al. (2015) pointed out the importance of capturing the confounders before actually performing actions. In particular, they proposed a new criterion for choosing actions called regret decision criterion (RDC). Under RDC rule, the agent collects her *intention* (includes information on confounders) on selecting actions and re-thinks based on the estimated reward statistics for the current *intention* and then decides which action to play. Since the actual selected action can be different than the agent's *intention*, the agent decides in a counterfactual way.

**Definition 5** (Confounder). A confounder is a variable that influences both the dependent variable and independent variable.

Later on, Lattimore et al. (2016) formally proposed the causal bandit framework via causal graphs. A causal model consists of a directed acyclic graph G over a set of random variables  $\mathcal{V} = \{V_1, \ldots, V_n\}$  and a joint distribution P that factorizes over G. A size m hard intervention (action) is denoted by do( $\mathbf{V} = \mathbf{v}$ ), which assigns the values  $\mathbf{v} = \{v_1, \ldots, v_m\} \subset \mathcal{V}$  to the corresponding variables  $\mathbf{V} = \{V_1, \ldots, V_m\}$ . In causal bandits, the action set is defined as

$$\mathcal{A} := \{ \operatorname{do}(\mathbf{V} = \mathbf{v}) | \mathbf{V} \subset \mathcal{V}, \mathbf{v} \in \operatorname{Dom}(\mathbf{V}) \},\$$

or its subset and the optimal intervention is  $a^* := \operatorname{argmax}_{a \in \mathcal{A}} \mathbb{E}[R|a].$ 

Many practical problems can be modeled via above framework. For example, in healthcare applications, the physician adjusts several features such as dose levels on different medicines or life-style advices to achieve some desirable clinical outcomes (Liu et al., 2020). Genetic engineering also involves direct manipulation of genes using biotechnology to produce improved organisms. In these problems, the number of interventions can be exponentially large in the number of manipulable variables so that standard MAB algorithms cannot work efficiently. To address this issue, recent work on causal bandits has developed methods that exploit causal information to achieve simple or cumulative regret that does not scale with the number of interventions (Lattimore et al., 2016; Sen et al., 2017; Lee and Bareinboim, 2018; Lu et al., 2020, 2021a). We summarize existing methodologies below.

**Refining the Policy Space.** Assuming the causal graph is known (confounders may also exist), Lee and Bareinboim (2018) proposed an intervention set reduction algorithm. They showed that not all interventions are worthwhile to be played, because there are equivalences among interventions and some interventions can be proved to be no better than others. Their algorithm filters out redundant interventions using the causal graph structure before applying any standard MAB algorithm on the reduced intervention set, also called *possibly-optimal minimal intervention set* (POMIS). Since the size of POMIS is usually much smaller than  $|\mathcal{A}|$ , the corresponding regret can be reduced. Follow-up work in Lee and Bareinboim (2019) further showed that the expected reward for interventions in a POMIS can be estimated from each other in certain cases. Their method also works well when some variables are non-manipulable. We remark that the step of finding the POMIS solely relies on the given causal graph structure and does not involve performing any intervention.

Accelerating Using Causal Graph Side Information in Standard MAB Algorithms. Recall that a key step in bandit algorithm is estimating the expected reward for each arm  $\hat{\mu}_a$ . Intuitively speaking, one needs to play each arm for enough number of times in order to get accurate estimations if interventions are independent. But in causal bandits, interventions are correlated with each other so that performing one intervention can help with estimating others according to the shared causal model. Lattimore et al. (2016) studied the best arm identification problem via importance weighting. Their algorithm achieves simple regret  $O(\sqrt{m/T})$ , where m can be smaller than  $|\mathcal{A}|$  in many scenarios. Sen et al. (2017) further generalized the algorithm in Lattimore et al. (2016) to a soft intervention setting and proved gap-dependent regret guarantees. In soft interventions, the agent does not directly set values to specific variables, but instead change their conditional probabilities given their direct parents. Lu et al. (2020) proposed causal UCB and causal TS algorithms that achieve  $\tilde{O}(\sqrt{ZT})$  cumulative regret, where Z is a graph dependent number and can be exponentially smaller than  $|\mathcal{A}|$ . A causal linear bandit framework and its efficient algorithms were also well studied. Furthermore, gap-dependent worst-case regret bound and a budgeted causal bandit setting was also considered in follow-up works (Nair et al., 2021).

## 3.6 Multi-task Learning

Multitask learning is the learning paradigm in machine learning that aims at leveraging information in multiple tasks to achieve a better sample efficiency by solving them jointly rather than separately (Zhang and Yang, 2021). Multitask learning algorithms usually assume some similarity between tasks that allows information sharing. In the context of multitask bandit, each bandit is called a task and various similarity assumptions on the reward distribution can be made. This can model many interesting scenarios. For instance, a group of patient with the same disease can be seen as multiple tasks. We can not directly use the same policy, as a personalized treatment should be applied to each patient. However, some patients' responses may be similar to the other's such that one can transfer the information across patients. Recall that we use  $\nu = (P_a, a \in \mathcal{A})$  denote a bandit instance. A general multitask bandit framework can be summarized below:

Given a set of bandits  $\nu_1, \ldots, \nu_n$ , Discrepancy $(\nu_i, \nu_j)$  is small for some  $i, j \in [n]$ ,

for some discrepancy measure on bandit instances.

In this subsection, we will review some interesting works on multitask bandit and compare their assumptions on discrepancy and the corresponding regret bounds. We will also cover the potential applications of multitask bandits.

**Parametric Model.** Early work on multitask bandit was done by Cesa-Bianchi et al. (2013) who considered a gang of contextual bandits connected with a graph  $\{V, E\}$ , where  $V = \{1, \ldots, n\}$  is the node set and E is the edge set. Each node i is assumed to be a linear bandit with rewards generated by  $R_i(c) = \theta_i^T c + \epsilon_i(c)$  given a context c, where  $\epsilon_i(c)$  is a conditionally zero-mean noise. Any two nodes that are adjacent on the graph have low discrepancy compared to the total scale:  $\sum_{i \in V} ||\theta_i||^2 \gg \sum_{(i,j) \in E} ||\theta_i - \theta_j||^2$ . Their discrepancy measure is simply the  $L_2$  distance on the unknown parameters. However, their analysis only shows a benefit of a factor n in the logarithmic term. In the literature of supervised multitask learning, the benefit could be a factor of  $n^{1/2}$  outside the logarithmic term (Maurer et al., 2016). Another property makes it hard for their algorithm to generalize to other settings is that the discrepancy information on tasks similarities is partially known to the agent through the graph structure. Generally speaking, the agent does not know which tasks are similar and have to adaptively learn the information.

As an improvement, Gentile et al. (2014); Li et al. (2016) consider an unknown graph with clustered structure, and all the bandits within the same cluster share a common reward distribution, which is also from a linear model. The algorithm is to simply cluster the bandits based on the feedbacks and pool the data from all the bandits from the same cluster. Instead of have a dependency of n as for the baseline algorithm that learns all the bandit separately, their regret is  $\tilde{O}(dm\sqrt{T})$ , where  $m \ll n$  is the number of clusters.

**Non-parametric Model.** Despite the encouraging results by Gentile et al. (2014); Li et al. (2016), the clustered structure and linear payoffs are strong assumptions. Deshmukh et al. (2017) proposed a kernel regression model that allows general similarities. Formally, they assume the reward is generated by mean function  $f : \mathcal{Z} \times \mathcal{C} \mapsto \mathbb{R}$ , where  $\mathcal{Z}$  is called the task similarity space,  $\mathcal{C}$  is the context from the context space  $\mathcal{C}$ . Let  $\tilde{k}$  be a SPD (semi-positive definite) kernel on  $\mathcal{Z} \times \mathcal{C}$  with the form

$$k\left((z,c),(z',c')\right) = k_{\mathcal{Z}}\left(z,z'\right)k_{\mathcal{C}}\left(c,c'\right),$$

where  $k_{\mathcal{Z}}$  and  $k_{\mathcal{C}}$  are SPD kernels on  $\mathcal{Z} = \{1, \ldots, n\}$  representing *n* tasks and  $\mathcal{C}$ . They assume that *f* is from the RKHS,  $\mathcal{H}_{\tilde{k}}$  corresponding to  $\tilde{k}$ . Kernel  $k_{\mathcal{Z}}$  characterized the similarity between different bandits, for instance:

- 1. Independent:  $k_{\mathcal{Z}}(z, z') = \mathbb{1}_{z=z'}$ , i.e. observations from different tasks are not sharing any information.
- 2. Pooled:  $k_{\mathcal{Z}} \equiv 1$ , i.e all the observations from different tasks are treated the same.
- 3. Multi-task:  $k_{\mathcal{Z}}$  is a PDS matrix reflecting task similarity.

Their algorithm **KMTL-UCB** gives an regret bound of  $\tilde{\mathcal{O}}(\sqrt{Tr_zr_c})$ , where  $r_z$  is the rank of the  $n \times n$  matrix  $K_Z = [k_Z(z_i, z_i)]_{i=1}^n$  and  $r_c$  is the rank of  $K_{C_T} := [k_C(c_{a_t,t}, c_{a_{t'},t'})]_{t,t'=1}^T$ . To compare, the regular regret bound of learning one task independently is  $\tilde{\mathcal{O}}(\sqrt{Tr_c})$  and all tasks independently gives  $\tilde{\mathcal{O}}(n\sqrt{Tr_c})$ . The scaling of  $r_z$  determines the benefits of multitask learning:  $r_z = n$  corresponds to the first case above with  $k_Z(z, z') = \mathbb{1}_{z=z'}$  and  $r_z = 1$  corresponds to the second case with  $k_Z \equiv 1$ . Though Deshmukh et al. (2017) gives a general way to characterize the multitask bandit benefits, their algorithm still relies on the knowledge of the similarity kernel  $k_Z$ . In practice, one can learn a kernel, but little theory has been established to justify this practice.

Adaptive Learning. Tomkins et al. (2021) proposed Intelligent Pooling, a generalization of a Thompson sampling contextual bandit for learning personalized treatment policies. They assume the reward for user  $i \in [n]$  is generated by

$$R_i = \phi(C_i, A_i)^T \theta_i + \epsilon_i,$$

where  $\phi : \mathcal{C} \times \mathcal{A} \mapsto \mathbb{R}^d$  is the pre-specified mapping from contexts and actions and  $\theta_i \in \mathbb{R}^d$  is the personalized parameter to learn. IntelligentPooling imposes each  $w_i$  as a random-effects upon a population-level parameter  $w_{pop}$ , i.e.  $w_i = w_{pop} + u_i$  with  $u_i$  the Gaussian random effect. Prior distribution of  $w_{pop}$  is also assumed to be Gaussian. IntelligentPooling runs by sampling  $w_{i,t}$  from the posterior distribution of  $w_i$ . The algorithm uses random effects to adaptively pool users' data based on the degree to which users exhibit heterogeneous rewards that is away from the population behavior  $(w_{pop})$ . They showed a regret bound of  $\tilde{O}(dn\sqrt{T})$ .

## 4 Conclusion

We reviewed basic topics in bandit algorithms that are helpful in understanding their increasing role in mobile health, and more broadly, precision medicine. We also discussed a selection of advanced topics that might inform the design of the next generation of bandit driven mobile health applications. We hope that the reader has gained an appreciation of the beauty and power of bandit algorithms and of their relevance to mobile health and precision medicine. However, like any methodology, bandit algorithms do have their limitations. First, bandit algorithms assume that an appropriate reward is available upfront. In practice, especially in healthcare applications, the design of rewards to be optimized is a difficult problem with no easy solution. The issues involved in specified a reward function that is aligned with the goals of various human stakeholders have been brought into focus more recently under the term "the alignment problem" (Christian, 2020). Second, bandit algorithms do not take into account delayed impacts of the agent's actions. When actions have delayed impacts, it might be beneficial to take an action with low immediate reward if it increases the probability of higher rewards later on. The full treatment of this problem of delayed impact of actions requires techniques from reinforcement learning (Sutton and Barto, 2018). In fact, bandit problems can be thought of as the simplest of reinforcement learning problems.

## References

Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. Advances in neural information processing systems, 24:2312–2320.

- Abbasi-Yadkori, Y., Pal, D., and Szepesvari, C. (2012). Online-to-confidence-set conversions and application to sparse stochastic bandits. In Artificial Intelligence and Statistics, pages 1–9. PMLR.
- Agrawal, S. and Devanur, N. (2016). Linear contextual bandits with knapsacks. Advances in Neural Information Processing Systems, 29:3450–3458.
- Agrawal, S. and Devanur, N. R. (2014). Bandits with concave rewards and convex knapsacks. In *Proceedings* of the fifteenth ACM conference on Economics and computation, pages 989–1006.
- Agrawal, S. and Goyal, N. (2012). Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pages 39–1. JMLR Workshop and Conference Proceedings.
- Agrawal, S. and Goyal, N. (2013). Further optimal regret bounds for thompson sampling. In Artificial intelligence and statistics, pages 99–107. PMLR.
- Aguilera, A., Figueroa, C. A., Hernandez-Ramos, R., Sarkar, U., Cemballi, A., Gomez-Pathak, L., Miramontes, J., Yom-Tov, E., Chakraborty, B., Yan, X., et al. (2020). mhealth app using machine learning to increase physical activity in diabetes and depression: clinical trial protocol for the diamante study. *BMJ* open, 10(8):e034723.
- Ameko, M. K., Beltzer, M. L., Cai, L., Boukhechba, M., Teachman, B. A., and Barnes, L. E. (2020). Offline contextual multi-armed bandits for mobile health interventions: A case study on emotion regulation. In Fourteenth ACM Conference on Recommender Systems, pages 249–258.
- Audibert, J.-Y., Bubeck, S., et al. (2009). Minimax policies for adversarial and stochastic bandits. In COLT, volume 7, pages 1–122.
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. Journal of Machine Learning Research, 3(Nov):397–422.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. SIAM journal on computing, 32(1):48–77.
- Auer, P., Gajane, P., and Ortner, R. (2019). Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Conference on Learning Theory*, pages 138–158. PMLR.
- Auer, P. and Ortner, R. (2010). Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65.
- Badanidiyuru, A., Kleinberg, R., and Slivkins, A. (2018). Bandits with knapsacks. Journal of the ACM (JACM), 65(3):1–55.
- Bang, H. and Robins, J. M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973.
- Bareinboim, E., Forney, A., and Pearl, J. (2015). Bandits with unobserved confounders: A causal approach. Advances in Neural Information Processing Systems, 28:1342–1350.
- Barocas, S., Hardt, M., and Narayanan, A. (2021). *Fairness and Machine Learning*. fairmlbook.org. http://www.fairmlbook.org.
- Bastani, H. and Bayati, M. (2020). Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294.
- Besbes, O., Gur, Y., and Zeevi, A. (2014). Stochastic multi-armed-bandit problem with non-stationary rewards. Advances in neural information processing systems, 27:199–207.

- Bubeck, S., Munos, R., and Stoltz, G. (2009). Pure exploration in multi-armed bandits problems. In *International conference on Algorithmic learning theory*, pages 23–37. Springer.
- Cao, Y., Wen, Z., Kveton, B., and Xie, Y. (2019). Nearly optimal adaptive procedure with change detection for piecewise-stationary bandit. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 418–427. PMLR.
- Carpentier, A., Vernade, C., and Abbasi-Yadkori, Y. (2020). The elliptical potential lemma revisited. arXiv preprint arXiv:2010.10182.
- Cesa-Bianchi, N., Gentile, C., and Zappella, G. (2013). A gang of bandits. In Proceedings of the 26th International Conference on Neural Information Processing Systems-Volume 1, pages 737–745.
- Cesa-Bianchi, N. and Lugosi, G. (2012). Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422.
- Charles, D., Chickering, M., and Simard, P. (2013). Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14.
- Chen, Y., Cuellar, A., Luo, H., Modi, J., Nemlekar, H., and Nikolaidis, S. (2020). The fair contextual multiarmed bandit. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent* Systems, pages 1810–1812.
- Chen, Y., Lee, C.-W., Luo, H., and Wei, C.-Y. (2019). A new algorithm for non-stationary contextual bandits: Efficient, optimal and parameter-free. In *Conference on Learning Theory*, pages 696–726. PMLR.
- Christian, B. (2020). The Alignment Problem: Machine Learning and Human Values. W. W. Norton.
- Chu, W., Li, L., Reyzin, L., and Schapire, R. (2011). Contextual bandits with linear payoff functions. In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, pages 208–214. JMLR Workshop and Conference Proceedings.
- Clayton, M. K. (1989). Covariate models for bernoulli bandits. Sequential Analysis, 8(4):405–426.
- Dani, V., Hayes, T. P., and Kakade, S. M. (2008). Stochastic linear optimization under bandit feedback.
- Deshmukh, A. A., Dogan, U., and Scott, C. (2017). Multi-task learning for contextual bandits. arXiv preprint arXiv:1705.08618.
- Ding, W., Qin, T., Zhang, X.-D., and Liu, T.-Y. (2013). Multi-armed bandit with budget constraint and variable costs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27.
- Dudík, M., Erhan, D., Langford, J., and Li, L. (2014). Doubly robust policy evaluation and optimization. Statistical Science, 29(4):485–511.
- Dudík, M., Langford, J., and Li, L. (2011). Doubly robust policy evaluation and learning. arXiv preprint arXiv:1103.4601.
- Even-Dar, E., Mannor, S., Mansour, Y., and Mahadevan, S. (2006). Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning* research, 7(6).
- Forman, E. M., Kerrigan, S. G., Butryn, M. L., Juarascio, A. S., Manasse, S. M., Ontañón, S., Dallal, D. H., Crochiere, R. J., and Moskow, D. (2019). Can the artificial intelligence technique of reinforcement learning use continuously-monitored digital data to optimize treatment for weight loss? *Journal of behavioral medicine*, 42(2):276–290.

- Garcelon, E., Ghavamzadeh, M., Lazaric, A., and Pirotta, M. (2020a). Conservative exploration in reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1431–1441. PMLR.
- Garcelon, E., Ghavamzadeh, M., Lazaric, A., and Pirotta, M. (2020b). Improved algorithms for conservative exploration in bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3962–3969.
- Garivier, A. and Moulines, E. (2011). On upper-confidence bound policies for switching bandit problems. In International Conference on Algorithmic Learning Theory, pages 174–188. Springer.
- Gentile, C., Li, S., and Zappella, G. (2014). Online clustering of bandits. In International Conference on Machine Learning, pages 757–765. PMLR.
- Gerchinovitz, S. (2011). Sparsity regret bounds for individual sequences in online linear regression. In Proceedings of the 24th Annual Conference on Learning Theory, pages 377–396. JMLR Workshop and Conference Proceedings.
- Gillen, S., Jung, C., Kearns, M., and Roth, A. (2018). Online learning with an unknown fairness metric. arXiv preprint arXiv:1802.06936.
- Goldenshluger, A. and Zeevi, A. (2011). A note on performance limitations in bandit problems with side information. *IEEE transactions on information theory*, 57(3):1707–1713.
- Goldenshluger, A. and Zeevi, A. (2013). A linear response bandit problem. *Stochastic Systems*, 3(1):230–261.
- Gupta, A., Koren, T., and Talwar, K. (2019). Better algorithms for stochastic bandits with adversarial corruptions. In *Conference on Learning Theory*, pages 1562–1578. PMLR.
- Hao, B., Lattimore, T., and Wang, M. (2020). High-dimensional sparse linear bandits. arXiv preprint arXiv:2011.04020.
- Haussler, D. and Warmuth, M. (2018). The probably approximately correct (pac) and other learning models. The Mathematics of Generalization, pages 17–36.
- Huang, J., Ma, S., and Zhang, C.-H. (2008). Adaptive lasso for sparse high-dimensional regression models. *Statistica Sinica*, pages 1603–1618.
- Huang, W., Labille, K., Wu, X., Lee, D., and Heffernan, N. (2020). Achieving user-side fairness in contextual bandits. arXiv preprint arXiv:2010.12102.
- Immorlica, N., Sankararaman, K. A., Schapire, R., and Slivkins, A. (2019). Adversarial bandits with knapsacks. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), pages 202–219. IEEE.
- Javanmard, A. and Montanari, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. The Journal of Machine Learning Research, 15(1):2869–2909.
- Jiang, N. and Li, L. (2016). Doubly robust off-policy value evaluation for reinforcement learning. In International Conference on Machine Learning, pages 652–661. PMLR.
- Jin, T., Xu, P., Shi, J., Xiao, X., and Gu, Q. (2020). Mots: Minimax optimal thompson sampling. arXiv preprint arXiv:2003.01803.
- Jin, Y., Yang, Z., and Wang, Z. (2021). Is pessimism provably efficient for offline rl? In International Conference on Machine Learning, pages 5084–5096. PMLR.

- Joseph, M., Kearns, M., Morgenstern, J., and Roth, A. (2016). Fairness in learning: Classic and contextual bandits. arXiv preprint arXiv:1605.07139.
- Kapoor, S., Patel, K. K., and Kar, P. (2019). Corruption-tolerant bandit learning. Machine Learning, 108(4):687–715.
- Kaufmann, E., Korda, N., and Munos, R. (2012). Thompson sampling: An asymptotically optimal finite-time analysis. In *International conference on algorithmic learning theory*, pages 199–213. Springer.
- Kazerouni, A., Ghavamzadeh, M., Abbasi-Yadkori, Y., and Van Roy, B. (2016). Conservative contextual linear bandits. arXiv preprint arXiv:1611.06426.
- Kim, W., Kim, G.-s., and Paik, M. C. (2021). Doubly robust thompson sampling for linear payoffs. arXiv preprint arXiv:2102.01229.
- Klasnja, P., Hekler, E. B., Shiffman, S., Boruvka, A., Almirall, D., Tewari, A., and Murphy, S. A. (2015). Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology*, 34(S):1220.
- Kocsis, L. and Szepesvári, C. (2006). Discounted ucb. In 2nd PASCAL Challenges Workshop, volume 2.
- Konstantinidou, M. K., Karaglani, M., Panagopoulou, M., Fiska, A., and Chatzaki, E. (2017). Are the origins of precision medicine found in the corpus hippocraticum? *Molecular diagnosis & therapy*, 21(6):601–606.
- Kusner, M. J., Loftus, J. R., Russell, C., and Silva, R. (2017). Counterfactual fairness. arXiv preprint arXiv:1703.06856.
- Lai, T. L. and Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. Advances in applied mathematics, 6(1):4–22.
- Langford, J. and Zhang, T. (2007). Epoch-greedy algorithm for multi-armed bandits with side information. Advances in Neural Information Processing Systems (NIPS 2007), 20:1.
- Lattimore, F., Lattimore, T., and Reid, M. D. (2016). Causal bandits: Learning good interventions via causal inference. arXiv preprint arXiv:1606.03203.
- Lattimore, T. (2015). Optimally confident ucb: Improved regret for finite-armed bandits. arXiv preprint arXiv:1507.07880.
- Lattimore, T. (2018). Refining the confidence level for optimistic bandit strategies. The Journal of Machine Learning Research, 19(1):765–796.
- Lattimore, T. and Szepesvári, C. (2020). Bandit algorithms. Cambridge University Press.
- Lee, S. and Bareinboim, E. (2018). Structural causal bandits: where to intervene? Advances in Neural Information Processing Systems 31, 31.
- Lee, S. and Bareinboim, E. (2019). Structural causal bandits with non-manipulable variables. In *Proceedings* of the AAAI Conference on Artificial Intelligence, volume 33, pages 4164–4172.
- Lei, H., Nahum-Shani, I., Lynch, K., Oslin, D., and Murphy, S. A. (2012). A" smart" design for building individualized treatment sequences. Annual review of clinical psychology, 8:21–48.
- Li, L., Chu, W., Langford, J., and Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670.

- Li, L., Chu, W., Langford, J., and Wang, X. (2011). Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the fourth ACM international conference on* Web search and data mining, pages 297–306.
- Li, L., Munos, R., and Szepesvári, C. (2015). Toward minimax off-policy value estimation. In Artificial Intelligence and Statistics, pages 608–616. PMLR.
- Li, S., Karatzoglou, A., and Gentile, C. (2016). Collaborative filtering bandits. In Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, pages 539– 548.
- Liao, P., Greenewald, K., Klasnja, P., and Murphy, S. (2020). Personalized heartsteps: A reinforcement learning algorithm for optimizing physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable* and Ubiquitous Technologies, 4(1):1–22.
- Liu, S., See, K. C., Ngiam, K. Y., Celi, L. A., Sun, X., and Feng, M. (2020). Reinforcement learning for clinical decision support in critical care: comprehensive review. *Journal of medical Internet research*, 22(7):e18477.
- Liu, Y., Radanovic, G., Dimitrakakis, C., Mandal, D., and Parkes, D. C. (2017). Calibrated fairness in bandits. arXiv preprint arXiv:1707.01875.
- Lu, Y., Meisami, A., and Tewari, A. (2021a). Causal bandits with unknown graph structure. arXiv preprint arXiv:2106.02988.
- Lu, Y., Meisami, A., and Tewari, A. (2021b). Low-rank generalized linear bandit problems. In International Conference on Artificial Intelligence and Statistics, pages 460–468. PMLR.
- Lu, Y., Meisami, A., Tewari, A., and Yan, W. (2020). Regret analysis of bandit problems with causal background knowledge. In Conference on Uncertainty in Artificial Intelligence, pages 141–150. PMLR.
- Lykouris, T., Mirrokni, V., and Paes Leme, R. (2018). Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 114–122.
- Ma, C., Zhu, B., Jiao, J., and Wainwright, M. J. (2021). Minimax off-policy evaluation for multi-armed bandits. arXiv preprint arXiv:2101.07781.
- Martin, A. R., Kanai, M., Kamatani, Y., Okada, Y., Neale, B. M., and Daly, M. J. (2019). Clinical use of current polygenic risk scores may exacerbate health disparities. *Nature genetics*, 51(4):584–591.
- Mary, J., Preux, P., and Nicol, O. (2014). Improving offline evaluation of contextual bandit algorithms via bootstrapping techniques. In *International Conference on Machine Learning*, pages 172–180. PMLR.
- Maurer, A., Pontil, M., and Romera-Paredes, B. (2016). The benefit of multitask representation learning. Journal of Machine Learning Research, 17(81):1–32.
- Nair, V., Patil, V., and Sinha, G. (2021). Budgeted and non-budgeted causal bandits. In International Conference on Artificial Intelligence and Statistics, pages 2017–2025. PMLR.
- Neu, G. (2015). Explore no more: Improved high-probability regret bounds for non-stochastic bandits. arXiv preprint arXiv:1506.03271.
- Niss, L. and Tewari, A. (2020). What you see may not be what you get: Ucb bandit algorithms robust to varepsilon-contamination. In Conference on Uncertainty in Artificial Intelligence, pages 450–459. PMLR.
- Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453.

- Paredes, P., Gilad-Bachrach, R., Czerwinski, M., Roseway, A., Rowan, K., and Hernandez, J. (2014). Poptherapy: Coping with stress through pop-culture. In *Proceedings of the 8th International Confer*ence on Pervasive Computing Technologies for Healthcare, pages 109–117.
- Patil, V., Ghalme, G., Nair, V., and Narahari, Y. (2020). Achieving fairness in the stochastic multi-armed bandit problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5379–5386.
- Paulus, J. K. and Kent, D. M. (2020). Predictably unequal: understanding and addressing concerns that algorithmic clinical prediction may increase health disparities. NPJ digital medicine, 3(1):1–8.
- Piette, J. D., Farris, K. B., Newman, S., An, L., Sussman, J., and Singh, S. (2015). The potential impact of intelligent systems for mobile health self-management support: Monte carlo simulations of text message support for medication adherence. Annals of Behavioral Medicine, 49(1):84–94.
- Piette, J. D., Krein, S. L., Striplin, D., Marinec, N., Kerns, R. D., Farris, K. B., Singh, S., An, L., and Heapy, A. A. (2016). Patient-centered pain care using artificial intelligence and mobile health tools: protocol for a randomized study funded by the us department of veterans affairs health services research and development program. JMIR research protocols, 5(2):e53.
- Rabbi, M., Aung, M. H., Zhang, M., and Choudhury, T. (2015). Mybehavior: automatic personalized health feedback from user behaviors and preferences using smartphones. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 707–718.
- Rabbi, M., Klasnja, P., Choudhury, T., Tewari, A., and Murphy, S. (2019). Optimizing mhealth interventions with a bandit. In Baumeister, H. and Montag, C., editors, *Mobile Sensing and Digital Phenotyping: New Developments in Psychoinformatics*. Springer.
- Rindtorff, N. T., Lu, M., Patel, N. A., Zheng, H., and D'Amour, A. (2019). A biologically plausible benchmark for contextual bandit algorithms in precision oncology using in vitro data. arXiv preprint arXiv:1911.04389.
- Rusmevichientong, P. and Tsitsiklis, J. N. (2010). Linearly parameterized bandits. Mathematics of Operations Research, 35(2):395–411.
- Sarkar, J. (1991). One-armed bandit problems with covariates. The Annals of Statistics, pages 1978–2002.
- Schumann, C., Lang, Z., Mattei, N., and Dickerson, J. P. (2019). Group fairness in bandit arm selection. arXiv preprint arXiv:1912.03802.
- Sen, R., Shanmugam, K., Dimakis, A. G., and Shakkottai, S. (2017). Identifying best interventions through online importance sampling. In *International Conference on Machine Learning*, pages 3057–3066. PMLR.
- Shrestha, S. and Jain, S. (2021). A bayesian-bandit adaptive design for n-of-1 clinical trials. Statistics in Medicine, 40(7):1825–1844.
- Singh, A. and Joachims, T. (2018). Fairness of exposure in rankings. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pages 2219–2228.
- Sun, W., Dey, D., and Kapoor, A. (2017). Safety-aware algorithms for adversarial contextual bandit. In International Conference on Machine Learning, pages 3280–3288. PMLR.
- Sutton, R. S. and Barto, A. G. (2018). Reinforcement learning: An introduction. MIT press.
- Tewari, A. and Murphy, S. A. (2017). From ads to interventions: Contextual bandits in mobile health. In Rehg, J., Murphy, S. A., and Kumar, S., editors, *Mobile Health: Sensors, Analytic Methods, and Applications*. Springer.

- 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.
- Tomkins, S., Liao, P., Klasnja, P., and Murphy, S. (2021). Intelligent pooling: practical thompson sampling for mhealth. *Machine Learning*, pages 1–43.
- Tran-Thanh, L., Chapman, A., De Cote, E. M., Rogers, A., and Jennings, N. R. (2010). Epsilon-first policies for budget-limited multi-armed bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24.
- Tran-Thanh, L., Chapman, A., Rogers, A., and Jennings, N. (2012). Knapsack based optimal policies for budget–limited multi–armed bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26.
- Wang, C.-C., Kulkarni, S. R., and Poor, H. V. (2005). Arbitrary side observations in bandit problems. Advances in Applied Mathematics, 34(4):903–938.
- Wang, L., Bai, Y., Sun, W., and Joachims, T. (2021). Fairness of exposure in stochastic bandits. arXiv preprint arXiv:2103.02735.
- Wang, Y.-X., Agarwal, A., and Dudık, M. (2017). Optimal and adaptive off-policy evaluation in contextual bandits. In *International Conference on Machine Learning*, pages 3589–3597. PMLR.
- Wei, C.-Y. and Luo, H. (2021). Non-stationary reinforcement learning without prior knowledge: An optimal black-box approach. arXiv preprint arXiv:2102.05406.
- Wu, H., Srikant, R., Liu, X., and Jiang, C. (2015). Algorithms with logarithmic or sublinear regret for constrained contextual bandits. arXiv preprint arXiv:1504.06937.
- Wu, Y., Shariff, R., Lattimore, T., and Szepesvári, C. (2016). Conservative bandits. In International Conference on Machine Learning, pages 1254–1262. PMLR.
- Xia, Y., Li, H., Qin, T., Yu, N., and Liu, T.-Y. (2015). Thompson sampling for budgeted multi-armed bandits. In Proceedings of the 24th International Conference on Artificial Intelligence, pages 3960–3966.
- Yang, Y., Zhu, D., et al. (2002). Randomized allocation with nonparametric estimation for a multi-armed bandit problem with covariates. *The Annals of Statistics*, 30(1):100–121.
- Yao, J., Brunskill, E., Pan, W., Murphy, S., and Doshi-Velez, F. (2020). Power-constrained bandits. arXiv preprint arXiv:2004.06230.
- Yom-Tov, E., Feraru, G., Kozdoba, M., Mannor, S., Tennenholtz, M., and Hochberg, I. (2017). Encouraging physical activity in patients with diabetes: intervention using a reinforcement learning system. *Journal of* medical Internet research, 19(10):e338.
- Zhang, K. W., Janson, L., and Murphy, S. A. (2021). Statistical inference with m-estimators on bandit data. arXiv preprint arXiv:2104.14074.
- Zhang, Y. and Yang, Q. (2021). A survey on multi-task learning. IEEE Transactions on Knowledge and Data Engineering.
- Zhou, D. and Tomlin, C. (2018). Budget-constrained multi-armed bandits with multiple plays. In *Proceedings* of the AAAI Conference on Artificial Intelligence, volume 32.
- Zhou, Z., Wang, Y., Mamani, H., and Coffey, D. G. (2019). How do tumor cytogenetics inform cancer treatments? dynamic risk stratification and precision medicine using multi-armed bandits. Dynamic Risk Stratification and Precision Medicine Using Multi-armed Bandits (June 17, 2019).