

Best-of-three-worlds Analysis for Linear Bandits with Follow-the-regularized-leader Algorithm

Fang Kong
Canzhe Zhao
Shuai Li*

FANGKONG@SJTU.EDU.CN
CANZHEZHAO@SJTU.EDU.CN
SHUAILI8@SJTU.EDU.CN

John Hopcroft Center for Computer Science, Shanghai Jiao Tong University, Shanghai, China

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

The linear bandit problem has been studied for many years in both stochastic and adversarial settings. Designing an algorithm that can optimize the environment without knowing the loss type attracts lots of interest. Lee et al. (2021) propose an algorithm that actively detects the loss type and then switches between different algorithms specially designed for specific settings. However, such an approach requires meticulous designs to perform well in all environments. Follow-the-regularized-leader (FTRL) is another type of popular algorithm that can adapt to different environments. This algorithm is of simple design and the regret bounds are shown to be optimal in traditional multi-armed bandit problems compared with the detect-switch type. Designing an FTRL-type algorithm for linear bandits is an important question that has been open for a long time. In this paper, we prove that the FTRL algorithm with a negative entropy regularizer can achieve the best-of-three-world results for the linear bandit problem. Our regret bounds achieve the same or nearly the same order as the previous detect-switch type algorithm but with a much simpler algorithmic design.

Keywords: Best-of-three-worlds, Linear bandits, Follow-the-regularized-leader, Negative entropy

1. Introduction

The linear bandit problem is a fundamental sequential decision-making problem, in which the learner needs to choose an arm to pull and will receive a loss of the pulled arm as a response in each round t (Auer, 2002; Awerbuch and Kleinberg, 2004; Dani et al., 2007, 2008; Abbasi-Yadkori et al., 2011; Chu et al., 2011; Bubeck et al., 2012). Particularly, the expected value of the loss of each arm is the inner product between the feature of this arm and an *unknown* loss vector. The learning objective is to minimize the *regret*, defined as the difference between the learner’s cumulative loss and the loss of some best arm in hindsight. There are two common environments characterizing the schemes of the loss received by the learner. In the stochastic environment, the loss vector is fixed across all rounds and the learner only observes a stochastic loss sampled from a distribution with the expectation being the inner product between the feature vector of the pulled arm and the loss vector. In this line of research, it has been shown that the instance-optimal regret bound is of order $O(\log T)$ (omitting all the dependence on other parameters) (Lattimore and Szepesvari, 2017). In the adversarial environment, the loss vector can arbitrarily change in each round. In this case, the minimax-optimal regret has been proved to be $O(\sqrt{T})$ (Bubeck et al., 2012).

While several algorithms are able to achieve the above (near) optimal regret guarantees in either stochastic environment (Lattimore and Szepesvari, 2017) or adversarial environment (Bubeck

* Corresponding author.

et al., 2012), these algorithms require prior knowledge about the regime of the environments. If such knowledge is not known a priori, the algorithms tailored to the adversarial environment will have too conservative regret guarantees in a stochastic environment. Also, without knowing such knowledge, the algorithms tailored to the stochastic environment may have linear regret guarantees in the adversarial environment and thus completely fail to learn. To cope with these issues, recent advances have emerged in simultaneously obtaining the (near) optimal regret guarantees in both the stochastic and adversarial environments, which is called the best-of-both-worlds (BoBW) result. For the special case of multi-armed bandit (MAB) problems, algorithms with (near) optimal regret in both adversarial and stochastic environments have been proposed (Bubeck and Slivkins, 2012; Seldin and Slivkins, 2014; Auer and Chiang, 2016; Seldin and Lugosi, 2017; Wei and Luo, 2018; Zimmert and Seldin, 2019; Ito, 2021b). In general, these algorithms fall into two lines. The first line is based on conducting statistical tests to distinguish between the stochastic and adversarial environments and use the specially designed algorithm for the detected setting (Bubeck and Slivkins, 2012; Auer and Chiang, 2016). The other approach is finding an appropriate regularizer for the follow-the-regularized-leader (FTRL) type algorithm which is originally designed for adversarial bandits to achieve improved regret guarantee in the stochastic environment, without compromising the adversarial regret guarantee (Seldin and Slivkins, 2014; Seldin and Lugosi, 2017; Wei and Luo, 2018; Zimmert and Seldin, 2019; Ito, 2021b). Compared with the first line, the second approach has better adaptivity since it does not need to detect the nature of the environment and is also shown to achieve the optimal result in both environments as well as the intermediate corruption setting (Zimmert and Seldin, 2019), which are called the best-of-three-worlds (BoTW) results. This approach also attracts great interest in more general problems including combinatorial MAB (Zimmert et al., 2019; Ito, 2021a), graph feedback (Erez and Koren, 2021; Ito et al., 2022) and Markov decision processes (MDPs) (Jin and Luo, 2020; Jin et al., 2021). But for the canonical linear bandit problem, only the detect-switch-based method (Lee et al., 2021) is available. A natural question is still open:

Does there exist an FTRL-type algorithm for linear bandits, with the same adaptivity to both the adversarial and (corrupted) stochastic environments?

In this paper, we give an affirmative answer to the above question. We investigate the FTRL algorithm with a negative entropy regularizer and show that such an algorithm can simultaneously achieve $O(\log T)$ regret in the stochastic setting and $O(\sqrt{T})$ regret in the adversarial setting. Our work provides a more concise algorithm for the BoTW problem in linear bandits given that only a detect-switch-based method was available previously. Table 1 provides the comparison between our work and previous works. As can be seen, our results achieve the same or nearly the same results compared with Lee et al. (2021), but with a much simpler algorithmic design. On the other hand, our results also unleash the potential of the FTRL-type algorithms to solve problems with linear structures. As a preliminary step, we believe the analysis for linear bandit also contributes some new ideas for the BoTW problem in other online learning models with linear structures such as linear MDPs.

2. Related Works

Linear bandits. Linear bandit is a fundamental model in online sequential decision-making problems. Its stochastic setting is originally proposed by Abe and Long (1999), and first solved by Auer (2002) using the optimism principle. Subsequently, the regret guarantees of this problem have been further improved by several algorithms (Dani et al., 2008; Rusmevichientong and Tsitsiklis, 2010;

Table 1: Comparisons of regret bounds with state-of-the-art results in different environments for linear bandits. T is the horizon, D is the finite arm set, C is the corruption level, Δ_{\min} is the minimum loss gap between the optimal and sub-optimal arms, ℓ is the loss vector in the stochastic setting, and $c(D, \ell)$ is the solution to the optimization problem in the lower bound of stochastic linear bandits with the natural upper bound d/Δ_{\min} .

	Stochastic with corruptions	Adversarial
Bubeck et al. (2012)		$O\left(\sqrt{dT \log D }\right)$
Lattimore and Szepesvari (2017)	$O(c(D, \ell) \log T)$ ($C = 0$)	
Lee et al. (2021)	$O(c(D, \ell) \log^2 T + C)$	$O\left(\sqrt{dT \log^2 T}\right)$
Ours	$O\left(d \log^2 T / \Delta_{\min} + \sqrt{Cd \log^2 T / \Delta_{\min}}\right)$	$O\left(\sqrt{dT \log^2 T}\right)$

¹ Particular care is required when comparing the results in the corrupted setting due to differences in the detailed problem definitions.

[Abbasi-Yadkori et al., 2011](#); [Chu et al., 2011](#)). The adversarial setting dates back to [Awerbuch and Kleinberg \(2004\)](#). The first $O(\sqrt{T})$ expected regret bound is obtained by the Geometric Hedge algorithm proposed by [Dani et al. \(2007\)](#). Latter, [Abernethy et al. \(2008\)](#) establish the first computationally efficient algorithm achieving $\tilde{O}(\sqrt{T})$ regret guarantee based on the FTRL framework. [Bubeck et al. \(2012\)](#) further obtain the regret guarantee that is minimax optimal up to a logarithmic factor. The recent work of [Lee et al. \(2021\)](#) achieves the (near) optimal instance-dependent regret and minimax optimal regret in stochastic and adversarial environments respectively, via establishing a detect-switch type algorithm. For linear bandits in a corrupted stochastic setting, [Li et al. \(2019\)](#) first obtain an instance-dependent regret guarantee, with an additional corruption term depending on the corruption level C linearly and on T logarithmically. Subsequently, [Bogunovic et al. \(2021\)](#) attain the instance-independent bound in the same setting. For the more general setting of linear contextual bandits where the arm set may vary in each round and can even be infinite, [Zhao et al. \(2021\)](#); [Ding et al. \(2022\)](#); [He et al. \(2022\)](#) establish the sublinear instance-independent and/or instance-dependent regret guarantees. Remarkably, [He et al. \(2022\)](#) obtain the near-optimal instance-independent regret guarantees for both the corrupted and uncorrupted cases with the leverage of a weighted ridge regression scheme. Due to the differences in the problem settings, the comparisons among these works require particular care.

Best-of-both-worlds algorithms. The best-of-both-worlds (BoBW) algorithms aim to obtain the (near) optimal theoretical guarantees in both the stochastic and adversarial environments simultaneously, and have been extensively investigated for a variety of problems in the online learning literature, including MAB problem ([Bubeck and Slivkins, 2012](#); [Seldin and Slivkins, 2014](#); [Auer and Chiang, 2016](#); [Seldin and Lugosi, 2017](#); [Wei and Luo, 2018](#); [Zimmert and Seldin, 2019](#); [Ito, 2021b](#)), linear bandit problem ([Lee et al., 2021](#)), combinatorial MAB problem ([Zimmert et al., 2019](#); [Ito, 2021a](#)), online learning with graph feedback ([Erez and Koren, 2021](#); [Kong et al., 2022](#); [Ito et al., 2022](#); [Rouyer et al., 2022](#)), online learning to rank ([Chen et al., 2022](#)), the problem of prediction with expert advice ([Gaillard et al., 2014](#); [De Rooij et al., 2014](#); [Luo and Schapire, 2015](#); [Mourtada and Gaïffas, 2019](#)), finite-horizon tabular Markov decision processes (MDPs) ([Jin and](#)

Luo, 2020; Jin et al., 2021), online submodular minimization (Ito, 2022), and partial monitoring problem (Tsuchiya et al., 2022).

In addition to investigating the BoBW results for stochastic and adversarial environments, devising algorithms that can simultaneously learn in the stochastic environment with corruptions also gains lots of research attention recently (Zimmert and Seldin, 2019; Erez and Koren, 2021; Ito, 2021c; Lee et al., 2021; Ito, 2021a). Particularly, algorithms with (near) optimal regret guarantees in adversarial, stochastic, and corrupted stochastic environments are typically called the best-of-three-worlds (BoTW) algorithms. In this work, we are interested in achieving the BoTW results in the linear bandit problem, but without using a detect-switch type algorithm similar to that of Lee et al. (2021). Instead, we aim to achieve the BoTW results for linear bandits by devising an FTRL-type algorithm, whose key idea is to leverage the self-bounding property of regret. Such types of algorithms have been proposed to obtain BoBW/BoTW results in MABs (Wei and Luo, 2018; Zimmert and Seldin, 2019; Ito, 2021b,c), online learning with graph feedback (Erez and Koren, 2021), combinatorial MAB problem (Zimmert et al., 2019; Ito, 2021a), online learning to rank (Chen et al., 2022), finite-horizon tabular MDPs (Jin and Luo, 2020; Jin et al., 2021), online submodular minimization (Ito, 2022), and partial monitoring problem (Tsuchiya et al., 2022).

3. Preliminaries

In this section, we present the basic preliminaries regarding the linear bandit problem, in both the adversarial environments and stochastic environments (with possibly adversarial corruptions).

We denote by $D \subseteq \mathbb{R}^d$ the *finite* arm set with d being the ambient dimension of the arm feature vectors (Lattimore and Szepesvari, 2017; Lee et al., 2021). Without loss of generality, we further assume that D spans \mathbb{R}^d and $\|x\|_2 \leq 1$ for all $x \in D$ as previous works (Lee et al., 2021; Bubeck et al., 2012). In each round t , the learner needs to select an arm $x_t \in D$. Meanwhile, there will be an *unknown* loss vector $\theta_t \in \mathbb{R}^d$ determined by the environment. Without loss of generality, we assume $\|\theta_t\|_2 \leq 1$ for any t . Subsequently, the loss $\ell_t(x_t)$ of the chosen arm x_t will be revealed to the learner. Particularly, $\ell_t(x_t) = \langle x_t, \theta_t \rangle + \epsilon_t(x_t)$ in linear bandits, with $\epsilon_t(x_t)$ being some independent zero-mean noise. The performance of one algorithm is measured by the *regret* (a.k.a., *expected pseudo-regret*), defined as

$$R(T) = \mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \sum_{t=1}^T \ell_t(x^*) \right],$$

where the expectation is taken over the randomness of both the loss sequence $\{\ell_t\}_{t=1}^T$ and the internal randomization of the algorithm and $x^* \in \operatorname{argmin}_{x \in D} \mathbb{E} \left[\sum_{t=1}^T \ell_t(x) \right]$ is the best arm in expectation in hindsight.

For linear bandits in *adversarial environment*, θ_t is chosen by an adversary arbitrarily in each round t . In this work, we consider the non-oblivious adversary, which means that the choice of θ_t may potentially depend on the history of the learner's actions $\{x_\tau\}_{\tau=1}^{t-1}$ up to round $t-1$, in addition to the adversary's own internal randomization. Also, for ease of exposition, we assume $\epsilon_t(x_t) = 0$ for all t in the adversarial environment.

The *stochastic* and *corrupted stochastic* linear bandits can be contained in the following adversarial regime with a self-bounding constraint which is first formalized by Zimmert and Seldin (2021).

Definition 1 (Adversarial regime with a self-bounding constraint) Let $C \geq 0$ and $\Delta \in [0, 2]^{|D|}$. An environment is defined as the adversarial regime with a (Δ, C, T) self-bounding constraint, if for any algorithm, the regret satisfies

$$R(T) \geq \sum_{t=1}^T \sum_{x \in D} \Delta(x) \mathbb{P}(x_t = x) - C. \quad (1)$$

The stochastic environment satisfies $\theta_t = \theta$ for some fixed but *unknown* θ . Hence it has regret $R(T) = \mathbb{E} \left[\sum_{t=1}^T \Delta(x_t) \right]$ with $\Delta(x) = \langle x - x^*, \theta \rangle$ for all $x \in D$. It follows that a stochastic environment is also an adversarial regime with a $(\Delta, 0, T)$ self-bounding constraint. For this case, we assume the optimal arm x^* to be unique as previous works studying stochastic linear bandits (Lattimore and Szepesvari, 2017; Lee et al., 2021) and BoTW problems (Bubeck and Slivkins, 2012; Zimmert and Seldin, 2019; Lee et al., 2021; Kong et al., 2022; Ito et al., 2022). For the corrupted stochastic environment, it is also satisfied that $\theta_t = \theta$ for some *unknown* θ , but the learner now receives loss $\ell_t(x_t) = \langle x_t, \theta \rangle + \epsilon_t(x_t) + c_t$, where c_t is the corrupted value chosen by the adversary in each round t and may depend on the historical information $\{(\ell_\tau, x_\tau)\}_{\tau=1}^{t-1}$. We note that the corrupted stochastic environment considered in this work is slightly more general than that considered by Lee et al. (2021), since we do not impose any structural assumptions over c_t , while they further assume c_t is linear in x_t . As aforementioned, the corrupted stochastic environment is also an instance of the adversarial regime with a self-bounding constraint. To see this, let $C = \sum_{t=1}^T |c_t|$ be the corruption level, indicating the total amount of the corruptions. Then a corrupted stochastic environment has regret satisfying Eq. (1), and hence it is an instance of the adversarial regime with a (Δ, C, T) self-bounding constraint.

4. Algorithm

In this section, we present the proposed algorithm for linear bandits in both adversarial and (corrupted) stochastic environments, and detail its pseudo-code in Algorithm 1.

Algorithm 1 FTRL for Linear Bandits

- 1: Input: regularizer $\{\psi_t\}_t$, exploration rate $\{\gamma_t\}_t$, arm set D and its G-optimal design π
 - 2: **for** each round $t = 1, 2, \dots$ **do**
 - 3: $q_t \in \operatorname{argmin}_{p \in \Delta(D)} \left\{ \sum_{s=1}^{t-1} \langle \hat{\ell}_s, p \rangle + \psi_t(p) \right\}$
 - 4: $p_t = \gamma_t \cdot \pi + (1 - \gamma_t) \cdot q_t$
 - 5: Sample action $x_t \sim p_t$ and observe the loss $\ell_t(x_t)$
 - 6: Update $\Sigma_t = \sum_{x \in D} p_t(x) x x^\top$; $\hat{\ell}_t(x) = x^\top \Sigma_t^{-1} x_t \ell_t(x_t)$ for all $x \in D$
 - 7: **end for**
-

In general, our algorithm follows the basic idea of FTRL that chooses the arm which minimizes the cumulative loss of past rounds together with the regularizer. Let $\Delta(D)$ be the probability simplex over D , $\hat{\ell}_s$ be the loss estimate in round s to be defined later and $\psi_t(p)$ be the regularizer. At each round t , our algorithm computes the regularized leader q_t by solving the optimization problem (Line 3). In particular, different from the complex (hybrid) regularizers used in previous works studying FTRL-type BoTW algorithms (Zimmert et al., 2019; Ito, 2021a; Erez and Koren, 2021), we simply

choose the negative Shannon entropy as the regularizer:

$$\psi_t(p) = -\beta_t H(p) = \sum_{x \in D} p(x) \ln p(x),$$

where β_t is the time-varying learning rate, whose concrete value is postponed to Theorem 1.

Then, to permit enough exploration and control the variance of the loss estimates, we mix the regularized leader q_t with a distribution π of the G-optimal design over the arm set D on Line 4, where γ_t is a time-varying hyper-parameter to be chosen later. In particular, the G-optimal design distribution π is chosen such that $\pi \in \operatorname{argmin}_{\rho \in \Delta(D)} g(\rho)$, where $g(\rho) = \max_{x \in D} \|x\|_{V(\rho)^{-1}}^2$ and $V(\rho) = \sum_{x \in D} \rho(x) x x^\top$ for any $\rho \in \Delta(D)$. Subsequently, the algorithm samples action x_t from the mixed distribution p_t , and observes the corresponding loss $\ell_t(x_t)$ (Line 5).

Denote by $\Sigma_t = \sum_{x \in D} p_t(x) x x^\top$ the covariance matrix in terms of p_t . At the end of each round t , we construct a least squares loss estimate $\hat{\ell}_t(x)$ for each $x \in D$ (Line 6).

4.1. Theoretical Results

In this section, we provide the theoretical guarantees of Algorithm 1 and the corresponding discussions. The following theorem presents the regret bounds for Algorithm 1 in adversarial and (corrupted) stochastic environments.

Theorem 1 *Choose $\psi_t(p) = -\beta_t H(p)$ and $\gamma_t = \min \{g(\pi)/\beta_t, 1/2\}$ where*

$$\beta_t := 2g(\pi) + \sqrt{d \ln T / \ln |D|} + \sum_{\tau=1}^{t-1} \frac{\sqrt{d \ln T / \ln |D|}}{\sqrt{1 + (\ln |D|)^{-1} \sum_{s=1}^{\tau} H(q_s)}}$$

and $H(p) := -\sum_{x \in D} p(x) \ln p(x)$ is the negative Shannon entropy of distribution p . Our Algorithm 1 satisfies that:

- In the adversarial environment, regret can be upper bounded by

$$R(T) \leq O\left(\sqrt{dT \ln T \ln(|D|T)}\right).$$

- In the stochastic environment with corruption level C , the regret can be bounded by

$$R(T) \leq O\left(\frac{d \ln T \ln(|D|T)}{\Delta_{\min}} + \sqrt{\frac{Cd \ln T \ln(|D|T)}{\Delta_{\min}}}\right), \quad (2)$$

where $\Delta_{\min} = \min_{x \in D, x \neq x^*} \Delta(x)$ is the minimum suboptimality gap over all sub-optimal arms.

Importance of the results. The FTRL-type algorithm has been demonstrated to achieve BoTW results for many problems including MAB (Zimmert and Seldin, 2019; Ito, 2021b), combinatorial MAB (Zimmert et al., 2019; Ito, 2021a), graph feedback (Erez and Koren, 2021; Ito et al., 2022), and MDP (Jin and Luo, 2020; Jin et al., 2021). However, the question of whether this algorithm can achieve BoTW results for the canonical linear bandit problem remains open. To the best of our knowledge, Theorem 1 provides the first BoTW results for linear bandits with FTRL algorithm. Although our results have not reached the optimal order in separate settings (Lattimore and Szepesvari, 2017; Bubeck et al., 2012), the convergence result itself is significant because it not only provides a more concise algorithm for this problem with only detect-switch-based methods available previously but also discovers the potential of FTRL-type algorithms.

Technical challenge. While the FTRL algorithm with a negative entropy regularizer is commonly used in the adversarial linear bandit literature, discovering the optimal exploration rate in the stochastic setting and achieving BoTW results is challenging. Similar to previous FTRL algorithms for the multi-armed bandit problem (Zimmert and Seldin, 2019), the instance-dependent regret in the stochastic setting mainly relies on the self-bounding inequality that lower bounds the regret by the selecting probability of arms. However, using the product of the sub-optimality gap and the selection probabilities of each sub-optimal arm as the self-bounding inequality and following the FTRL technique for MAB (Zimmert and Seldin, 2019) introduce a dependence on the number of arms, which is not ideal in the linear bandit case. By leveraging the linear structure of the loss estimators and using standard deduction of FTRL, we discover that the regret can be upper bounded using the dimension d , the learning rate, and the cumulative entropy of the output decision vector. The former reflects the problem hardness of linear structured problems, and the latter can be further exploited to obtain the cumulative selecting times of all sub-optimal arms. This observation leads us to define the self-bounding inequality (Lemma 2) on the selecting times of all sub-optimal arms and use the minimum gap Δ_{\min} to measure the problem hardness.

Lemma 2 (*Self-bounding inequality*) *Algorithm 1 with Definition 1 satisfies that*

$$R(T) \geq \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T (1 - q_t(x^*)) \right] \Delta_{\min} - C.$$

The proof is deferred to Section 5.1. Such an observation connects the cumulative entropy with the regret definition and guarantees convergence in the stochastic setting.

To simultaneously achieve an $O(\sqrt{T})$ result in the adversarial setting and $O(\log T)$ result in the stochastic setting, we need to balance the choice of the learning rate. Based on the standard analysis of the FTRL algorithm, the regret can be decomposed as

$$R(T) \leq O \left(\mathbb{E} \left[\sum_{t=1}^T \left(\frac{1}{\beta_t} + (\beta_{t+1} - \beta_t) H(q_{t+1}) \right) \right] \right).$$

In the adversarial setting, we hope to get an $O(\sqrt{T})$ regret, the natural choice is to select $\beta_t = O(\sqrt{t})$ as both $\sum_t 1/\sqrt{t} = O(\sqrt{t})$ and $\sum_t (\sqrt{t+1} - \sqrt{t}) = O(\sqrt{t})$. While in the stochastic setting, we aim to get an $O(\log T)$ regret, so it would be natural to choose $\beta_t = O(t)$ as $H(q_t)$ tends to be 0 when the policy always selects the unique optimal arm. Inspired by the motivation that $H(q_t)$ can be regarded as a constant term in the adversarial setting and a decreasing sequence that eventually tends to 0 in the stochastic setting, we set the learning rate β_t as $O(\sum_{s=1}^{t-1} 1/\sqrt{1 + \sum_{s=1}^s H(q_s)})$.

The cooperation of the specially designed self-bounding inequality and the choice of $\{\beta_t\}_t$ contributes to the final results. And here we mainly follow the technique of Ito et al. (2022) to achieve these results. The detailed proof is deferred to the next section.

Comparisons with concurrent works. There is also a concurrent work studying the FTRL-type algorithm for the BoTW problem in linear bandits (Ito and Takemura, 2023). Ito and Takemura (2023) study the data-dependent regret for this problem. In the worst case, their algorithm achieves $O(d^{3/2} \sqrt{T \log T})$ regret in the adversarial setting and $O(d^3 \log T / \Delta_{\min} + \sqrt{C d^3 \log T / \Delta_{\min}})$ regret in the corrupted stochastic setting. Compared with these results, our regret bound is $O(d^2 / \log T)$

better in the stochastic setting and $O(d/\sqrt{\log T})$ better in the adversarial setting. We also note that their results are data-dependent bounds, which may be better in environments with certain characteristics.

5. Theoretical Analysis

This section provides the proof of Theorem 1. We first present some useful lemmas and the main proof would come later.

We start from the following lemma which shows that the regret can be decomposed in the form of the cumulative entropy of exploitation probability q_t .

Lemma 3 (*Regret Decomposition*) *By choosing $\beta_t = 2g(\pi) + c + \sum_{\tau=1}^{t-1} \frac{c}{\sqrt{1+\ln|D|^{-1} \sum_{s=1}^{\tau} H(q_s)}}$ and $\gamma_t = \min\{g(\pi)/\beta_t, 1/2\}$, the regret can be bounded as*

$$R(T) = O\left(\left(\frac{d \ln T}{c\sqrt{\ln|D|}} + c\sqrt{\ln|D|}\right) \sqrt{\mathbb{E}\left[\sum_{t=1}^T H(q_t)\right]} + (2g(\pi) + c) \ln|D|\right). \quad (3)$$

The next lemma further links the cumulative entropy with the expected selecting times of all sub-optimal arms.

Lemma 4 (*Relationship between $H(q_t)$ and expected selecting times of all sub-optimal arms*)

$$\sum_{t=1}^T H(q_t) \leq \left(\sum_{t=1}^T (1 - q_t(x^*))\right) \ln \frac{e|D|T}{\sum_{t=1}^T (1 - q_t(x^*))}.$$

The detailed proof for both lemmas are deferred to Section 5.1. Based on Lemma 3 and 4, we are now able to upper bound the regret as below.

$$\begin{aligned} R(T) &\leq O\left(\left(\frac{d \ln T}{c\sqrt{\ln|D|}} + c\sqrt{\ln|D|}\right) \sqrt{\mathbb{E}\left[\sum_{t=1}^T H(q_t)\right]} + (2g(\pi) + c) \ln|D|\right) \\ &\leq O\left(\left(\frac{d \ln T}{c\sqrt{\ln|D|}} + c\sqrt{\ln|D|}\right) \sqrt{\mathbb{E}\left[\underbrace{\left(\sum_{t=1}^T (1 - q_t(x^*))\right) \ln \frac{e|D|T}{\sum_{t=1}^T (1 - q_t(x^*))}}_{\text{term A}}\right]} + (2g(\pi) + c) \ln|D|\right) \\ &\leq O\left(\left(\frac{d \ln T}{c\sqrt{\ln|D|}} + c\sqrt{\ln|D|}\right) \sqrt{\mathbb{E}\left[\ln(e|D|T) \max\left\{e, \sum_{t=1}^T (1 - q_t(x^*))\right\}\right]} + (2g(\pi) + c) \ln|D|\right) \\ &= O\left(\left(\frac{d \ln T}{c\sqrt{\ln|D|}} + c\sqrt{\ln|D|}\right) \sqrt{\ln(e|D|T) \max\left\{e, \mathbb{E}\left[\sum_{t=1}^T (1 - q_t(x^*))\right]\right\}} + (2g(\pi) + c) \ln|D|\right) \end{aligned} \quad (4)$$

$$=O\left(\sqrt{d\ln T} \cdot \sqrt{\ln(e|D|T) \max\left\{e, \mathbb{E}\left[\sum_{t=1}^T(1-q_t(x^*))\right]\right\}}\right), \quad (5)$$

where Eq. (4) holds since term $A \leq \sum_{t=1}^T(1-q_t(x^*)) \ln|D|T$ if $\sum_{t=1}^T(1-q_t(x^*)) > e$ and term $A \leq e \ln e|D|^2T/(|D|-1) = O(e \ln(e|D|T))$ otherwise as $\sum_{t=1}^T(1-q_t(x^*)) \geq 1 - q_1(x^*) = 1 - 1/|D|$; Eq. (5) is due to $c = \sqrt{d\ln T/\ln|D|}$.

Regret analysis in the adversarial setting. According to Eq. (5), since $\sum_{t=1}^T(1-q_t(x^*)) \leq T$, it holds that

$$R(T) \leq O\left(\left(\sqrt{d\ln T} \sqrt{\ln(e|D|T) \cdot T}\right)\right) = O\left(\sqrt{dT \ln T \ln(|D|T)}\right).$$

Regret analysis in the stochastic setting with corruptions. According to Lemma 2,

$$\begin{aligned} R(T) &= (1+\lambda)R(T) - \lambda R(T) \\ &\leq O\left(\left(1+\lambda\right) \cdot \sqrt{d\ln T} \cdot \sqrt{\ln(e|D|T) \max\left\{e, \mathbb{E}\left[\sum_{t=1}^T(1-q_t(x^*))\right]\right\}}\right. \\ &\quad \left.- \frac{1}{2}\lambda \mathbb{E}\left[\sum_{t=1}^T(1-q_t(x^*))\right] \Delta_{\min} + \lambda C\right) \\ &\leq O\left(\frac{(1+\lambda)^2 d \ln T \ln(|D|T)}{2\lambda \Delta_{\min}} + \lambda C\right) \end{aligned} \quad (6)$$

$$\begin{aligned} &= O\left(\frac{d \ln T \ln(|D|T)}{2\lambda \Delta_{\min}} + \frac{d \ln T \ln(|D|T)}{\Delta_{\min}} + \frac{\lambda d^2 \ln T \ln(|D|T)}{2\Delta_{\min}} + \lambda C\right) \\ &\leq O\left(\frac{d \ln T \ln(|D|T)}{\Delta_{\min}} + \sqrt{\frac{Cd \ln T \ln(|D|T)}{\Delta_{\min}}}\right), \end{aligned} \quad (7)$$

where Eq. (6) holds since $a\sqrt{x} - bx/2 = a^2/2b - 1/2\left(a/\sqrt{b} - \sqrt{bx}\right) \leq a^2/2b$ for any $a, b, x \geq 0$;

Eq. (7) holds by choosing $\lambda = \sqrt{\frac{d \ln T \ln(|D|T)/(2\Delta_{\min})}{d \ln T \ln(|D|T)/(2\Delta_{\min}) + C}}$.

5.1. Proof of Useful Lemmas

This section provides detailed proof for the lemmas used in the analysis. The first is the proof of Lemma 2 which is about self-bounding inequality.

Proof [Proof of Lemma 2]

$$\begin{aligned} R(T) &\geq \mathbb{E}\left[\sum_{t=1}^T(\ell_t(x_t) - \ell_t(x^*))\right] - C \\ &= \mathbb{E}\left[\sum_{t=1}^T\langle \ell_t, p_t - p^* \rangle\right] - C \end{aligned}$$

$$\begin{aligned}
 &\geq \mathbb{E} \left[\sum_{t=1}^T (1 - \gamma_t) \langle \ell_t, q_t - p^* \rangle \right] - C \\
 &= \mathbb{E} \left[\sum_{t=1}^T (1 - \gamma_t) \sum_x q_t(x) \Delta(x) \right] - C \\
 &\geq \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T (1 - q_t(x^*)) \right] \Delta_{\min} - C,
 \end{aligned}$$

where the last inequality holds by choosing $\gamma_t \leq 1/2$ for all t . ■

In the following, we provide the proof of Lemma 3 which is used in the analysis of Theorem 1.

Proof [Proof of Lemma 3]

$$\begin{aligned}
 R(T) &= \mathbb{E} \left[\sum_{t=1}^T \ell_t(x_t) - \ell_t(x^*) \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \langle \ell_t, p_t - p^* \rangle \right] \tag{8}
 \end{aligned}$$

$$= \mathbb{E} \left[\sum_{t=1}^T (1 - \gamma_t) \langle \hat{\ell}_t, q_t - p^* \rangle \right] + \mathbb{E} \left[\sum_{t=1}^T \gamma_t \langle \ell_t, \pi - p^* \rangle \right] \tag{9}$$

$$\leq \mathbb{E} \left[\sum_{t=1}^T (1 - \gamma_t) \langle \hat{\ell}_t, q_t - p^* \rangle \right] + 2 \sum_{t=1}^T \gamma_t \tag{10}$$

$$\begin{aligned}
 &\leq \mathbb{E} \left[\sum_{t=1}^T (1 - \gamma_t) \left(\langle \hat{\ell}_t, q_t - q_{t+1} \rangle - D_t(q_{t+1}, q_t) + \psi_t(q_{t+1}) - \psi_{t+1}(q_{t+1}) \right) \right] \\
 &\quad + \psi_{T+1}(p^*) - \psi_1(q_1) + 2 \sum_{t=1}^T \gamma_t \tag{11}
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E} \left[\sum_{t=1}^T (1 - \gamma_t) \left(\langle \hat{\ell}_t, q_t - q_{t+1} \rangle - D_t(q_{t+1}, q_t) + (\beta_{t+1} - \beta_t) H(q_{t+1}) \right) \right] \\
 &\quad + \psi_{T+1}(p^*) - \psi_1(q_1) + 2 \sum_{t=1}^T \gamma_t \\
 &\leq \mathbb{E} \left[\sum_{t=1}^T (1 - \gamma_t) \left(\underbrace{\langle \hat{\ell}_t, q_t - q_{t+1} \rangle - D_t(q_{t+1}, q_t)}_{\text{part 1}} + (\beta_{t+1} - \beta_t) H(q_{t+1}) \right) \right] + \beta_1 \ln |D| + 2 \sum_{t=1}^T \gamma_t \\
 &\leq \mathbb{E} \left[\sum_{t=1}^T \left(\frac{3d}{\beta_t} + \underbrace{(\beta_{t+1} - \beta_t) H(q_{t+1})}_{\text{part 2}} \right) \right] + \beta_1 \ln |D| \tag{12}
 \end{aligned}$$

$$\leq \mathbb{E} \left[\sum_{t=1}^T \frac{3d}{\beta_t} + 2c \sqrt{\ln |D| \sum_{t=1}^T H(q_t)} \right] + (2g(\pi) + c) \ln |D| \quad (13)$$

$$\leq \mathbb{E} \left[3d \sum_{t=1}^T \frac{\sqrt{1 + (\ln |D|)^{-1} \sum_{s=1}^t H(q_s)}}{c \cdot t} + 2c \sqrt{\ln |D| \sum_{t=1}^T H(q_t)} \right] + (2g(\pi) + c) \ln |D| \quad (14)$$

$$\leq O \left(\left(\frac{d \ln T}{c \sqrt{\ln |D|}} + c \sqrt{\ln |D|} \right) \sqrt{\mathbb{E} \left[\sum_{t=1}^T H(q_t) \right]} + (2g(\pi) + c) \ln |D| \right).$$

where Eq. (8) holds due to $x_t \sim p_t$, Eq. (9) holds since $\hat{\ell}_t$ is an unbiased estimator, Eq. (10) is because $\langle \ell_t, \pi - p^* \rangle \leq 2$, Eq. (11) is based on Lemma 5, Eq. (12) holds according to Lemma 6 and the choice of $\gamma_t = \min \{g(\pi)/\beta_t, 1/2\} \leq d/\beta_t$ (Kiefer and Wolfowitz, 1960), Eq. (13) and (14) are due to Lemma 8 and the choice of β_t , the last inequality is based on Jensen's inequality. ■

We next provide Lemma 5, 6 and 8 that is used in the proof of Lemma 3.

Lemma 5 (*Decomposition of FTRL, Exercise 28.12 in Lattimore and Szepesvári (2020)*)

$$\begin{aligned} \sum_{t=1}^T \langle \hat{\ell}_t, q_t - p^* \rangle &\leq \sum_{t=1}^T \left(\langle \hat{\ell}_t, q_t - q_{t+1} \rangle - D_t(q_{t+1}, q_t) + \psi_t(q_{t+1}) - \psi_{t+1}(q_{t+1}) \right) \\ &\quad + \psi_{T+1}(p^*) - \psi_1(q_1). \end{aligned}$$

Lemma 6 (*Bound Part 1*)

$$\mathbb{E} \left[(1 - \gamma_t) \left(\langle \hat{\ell}_t, q_t - q_{t+1} \rangle - D_t(q_{t+1}, q_t) \right) \right] \leq \frac{d}{\beta_t}.$$

Proof

$$\begin{aligned} &\mathbb{E} \left[(1 - \gamma_t) \left(\langle \hat{\ell}_t, q_t - q_{t+1} \rangle - D_t(q_{t+1}, q_t) \right) \right] \\ &\leq \mathbb{E} \left[(1 - \gamma_t) \beta_t \left(\sum_{x \in D} q_t(x) \left(\exp(-\hat{\ell}_t(x)/\beta_t) + \hat{\ell}_t(x)/\beta_t - 1 \right) \right) \right] \quad (15) \end{aligned}$$

$$\leq \mathbb{E} \left[(1 - \gamma_t) \beta_t \left(\sum_{x \in D} q_t(x) \left(\hat{\ell}_t^2(x)/\beta_t^2 \right) \right) \right] \quad (16)$$

$$\begin{aligned} &= \mathbb{E} \left[\frac{(1 - \gamma_t)}{\beta_t} \left(\sum_{x \in D} q_t(x) \hat{\ell}_t^2(x) \right) \right] \\ &\leq \mathbb{E} \left[\frac{1}{\beta_t} \left(\sum_{x \in D} p_t(x) \hat{\ell}_t^2(x) \right) \right], \quad (17) \end{aligned}$$

where Eq. (15) is based on Ito et al. (2022, Lemma 8), Eq. (16) holds due to $\exp(-a) + a - 1 \leq a^2$ when $a \geq -1$ and Lemma 7, Eq. (17) is from the definition of p_t .

Recall that $\hat{\ell}_t(x) = x^\top \Sigma_t^{-1} x_t \ell_t(x_t)$. We have

$$\hat{\ell}_t(x)^2 = \ell_t(x_t)^2 x_t^\top \Sigma_t^{-1} x x^\top \Sigma_t^{-1} x_t \leq x_t^\top \Sigma_t^{-1} x x^\top \Sigma_t^{-1} x_t$$

and further,

$$\sum_x p_t(x) \hat{\ell}_t(x)^2 \leq \sum_x p_t(x) x_t^\top \Sigma_t^{-1} x x^\top \Sigma_t^{-1} x_t = x_t^\top \Sigma_t^{-1} x_t = \text{trace}(x_t x_t^\top \Sigma_t^{-1}).$$

Above all,

$$\mathbb{E} \left[\sum_x p_t(x) \hat{\ell}_t(x)^2 \right] \leq \text{trace} \left(\sum_x p_t(x) x x^\top \Sigma_t^{-1} \right) = d.$$

Lemma 6 is then proved. ■

Lemma 7 Denote $\Sigma(\pi) = \sum_{x \in D} \pi(x) x x^\top$ and $g(\pi) = \max_{x \in D} x^\top \Sigma(\pi)^{-1} x$. By choosing $\gamma_t = \min \{g(\pi)/\beta_t, 1/2\}$, we have $\hat{\ell}_t(x)/\beta_t \geq -1$ for all $x \in D$.

Proof Since $\Sigma_t = \sum_x p_t(x) x x^\top = \sum_x ((1 - \gamma_t) q_t(x) + \gamma_t \pi(x)) x x^\top \succeq \sum_x \gamma_t \pi(x) x x^\top = \gamma_t \Sigma(\pi)$, we have

$$\Sigma_t^{-1} \preceq \frac{1}{\gamma_t} \Sigma(\pi)^{-1}.$$

Further,

$$\left| x^\top \Sigma_t^{-1} x_t \right| \leq \|x\|_{\Sigma_t^{-1}} \|x_t\|_{\Sigma_t^{-1}} \leq \max_{v \in D} v^\top \Sigma_t^{-1} v \leq \frac{1}{\gamma_t} \max_{v \in D} v^\top \Sigma(\pi)^{-1} v = \frac{g(\pi)}{\gamma_t}.$$

With the choice of γ_t , it holds that

$$\left| \frac{\hat{\ell}_t(x)}{\beta_t} \right| = \left| \frac{1}{\beta_t} x^\top \Sigma_t^{-1} x_t \ell_t(x_t) \right| \leq \left| \frac{1}{\beta_t} x^\top \Sigma_t^{-1} x_t \right| \leq \frac{g(\pi)}{\beta_t \gamma_t} \leq 1,$$

The last inequality is due to the choice of β_t . ■

Lemma 8 (Bound Part 2) By choosing $\beta_t = 2g(\pi) + c + \sum_{\tau=1}^{t-1} \frac{c}{\sqrt{1 + \ln |D|^{-1} \sum_{s=1}^{\tau} H(q_s)}}$,

$$\sum_{t=1}^T (\beta_{t+1} - \beta_t) H(q_{t+1}) \leq 2c \sqrt{\ln |D| \sum_{t=1}^T H(q_t)}.$$

Proof Based on the definition of β_t , it holds that

$$\begin{aligned}
 \sum_{t=1}^T (\beta_{t+1} - \beta_t) H(q_{t+1}) &= \sum_{t=1}^T \frac{c}{\sqrt{1 + \ln |D|^{-1} \sum_{s=1}^t H(q_s)}} H(q_{t+1}) \\
 &\leq \sqrt{\ln |D|} \sum_{t=1}^T \frac{c}{\sqrt{\ln |D| + \sum_{s=1}^t H(q_s)}} H(q_{t+1}) \\
 &= 2\sqrt{\ln |D|} \sum_{t=1}^T \frac{c}{\sqrt{\ln |D| + \sum_{s=1}^t H(q_s)} + \sqrt{\ln |D| + \sum_{s=1}^t H(q_s)}} H(q_{t+1}) \\
 &\leq 2\sqrt{\ln |D|} \sum_{t=1}^T \frac{c}{\sqrt{\sum_{s=1}^{t+1} H(q_s)} + \sqrt{\sum_{s=1}^t H(q_s)}} H(q_{t+1}) \\
 &= 2c\sqrt{\ln |D|} \sum_{t=1}^T \left(\sqrt{\sum_{s=1}^{t+1} H(q_s)} - \sqrt{\sum_{s=1}^t H(q_s)} \right) \\
 &= 2c\sqrt{\ln |D|} \left(\sqrt{\sum_{s=1}^{T+1} H(q_s)} - \sqrt{H(q_1)} \right) \\
 &\leq 2c\sqrt{\ln |D|} \sqrt{\sum_{t=1}^T H(q_t)},
 \end{aligned}$$

where the last two inequalities holds since $H(q_t) \leq H(q_1) = \ln |D|$ for any t . ■

And the last is the proof of Lemma 4 that is used in the analysis of Theorem 1.

Proof [Proof of Lemma 4] According to the definition of $H(p)$ for a distribution p over D and the fact that the optimal arm $x^* \in D$ is unique, we can bound $H(p)$ as

$$\begin{aligned}
 H(p) &= \sum_x p(x) \ln \frac{1}{p(x)} = p(x^*) \ln \frac{p(x^*) + 1 - p(x^*)}{p(x^*)} + \sum_{x \neq x^*} p(x) \ln \frac{1}{p(x)} \\
 &\leq p(x^*) \frac{1 - p(x^*)}{p(x^*)} + \sum_{x \neq x^*} p(x) \ln \frac{1}{p(x)} \\
 &\leq 1 - p(x^*) + (|D| - 1) \frac{\sum_{x \neq x^*} p(x)}{|D| - 1} \ln \frac{|D| - 1}{\sum_{x \neq x^*} p(x)} \\
 &= (1 - p(x^*)) \left(1 + \ln \frac{|D| - 1}{(1 - p(x^*))} \right),
 \end{aligned}$$

where the last inequality is due to Jensen's inequality. Thus,

$$\sum_{t=1}^T H(q_t) \leq \sum_{t=1}^T (1 - q_t(x^*)) \left(1 + \ln \frac{|D| - 1}{(1 - q_t(x^*))} \right)$$

$$\begin{aligned}
&\leq \left(\sum_{t=1}^T (1 - q_t(x^*)) \right) \left(1 + \ln \frac{T(|D| - 1)}{\sum_{t=1}^T (1 - q_t(x^*))} \right) \\
&\leq \left(\sum_{t=1}^T (1 - q_t(x^*)) \right) \ln \frac{e|D|T}{\sum_{t=1}^T (1 - q_t(x^*))}.
\end{aligned}$$

■

6. Conclusion

In this work, we aim to provide the theoretical analysis for an FTRL-type algorithm when solving the BoTW problem in linear bandits. We investigate the performances of the FTRL algorithm with the common negative entropy regularizer. By discovering its difference in different environments and defining the self-bounding inequality to cooperate with it, we show that the algorithm can achieve $O(\log T)$ regret in the stochastic setting and $O(\sqrt{T})$ regret in the adversarial setting. To the best of our knowledge, it is the first BoTW analysis for linear bandits with an FTRL-type algorithm. Our work not only provides a simpler algorithm for this problem compared with the previous detect-switch-based method (Lee et al., 2021), but also taps the potential of FTRL when dealing with sequential decision-making problems with linear structures. As a special case of linear MDP, we believe our analysis for linear bandits also provides new insights into the BoTW problem for this general model.

Stochastic linear bandit problem has a special property that the optimal result depends on $c(D, \ell)$ which is the solution to an optimization problem and has no explicit form (Lattimore and Szepesvari, 2017). A regret bound that adapts to $c(D, \ell)$ would better illustrate the problem’s structure. Lee et al. (2021) achieve this since they adopt a detect-switch-based method and the specially designed algorithm for the stochastic setting helps. One interesting future direction would be to investigate such optimal exploration rates for FTRL-type algorithms and derive regret bounds in the stochastic setting that enjoys the form of $c(D, \ell)$.

Acknowledgments

We thank Shinji Ito for valuable discussions and suggestions on the improvement of the learning rate. The corresponding author Shuai Li is supported by National Key Research and Development Program of China (2022ZD0114804) and National Natural Science Foundation of China (62006151).

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in neural information processing systems*, pages 2312–2320, 2011.
- Naoki Abe and Philip M Long. Associative reinforcement learning using linear probabilistic concepts. In *Proceedings of the Sixteenth International Conference on Machine Learning*, pages 3–11, 1999.

- Jacob Abernethy, Elad Hazan, and Alexander Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. In *Annual Conference on Learning Theory*, pages 263–274, 2008.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Peter Auer and Chao-Kai Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 116–120. PMLR, 2016.
- Baruch Awerbuch and Robert D Kleinberg. Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 45–53, 2004.
- Ilija Bogunovic, Arpan Losalka, Andreas Krause, and Jonathan Scarlett. Stochastic linear bandits robust to adversarial attacks. In *International Conference on Artificial Intelligence and Statistics*, pages 991–999. PMLR, 2021.
- Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 42–1. JMLR Workshop and Conference Proceedings, 2012.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, and Sham M Kakade. Towards minimax policies for online linear optimization with bandit feedback. In *Conference on Learning Theory*, pages 41–1, 2012.
- Cheng Chen, Canzhe Zhao, and Shuai Li. Simultaneously learning stochastic and adversarial bandits under the position-based model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6202–6210, 2022.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.
- Varsha Dani, Sham M Kakade, and Thomas Hayes. The price of bandit information for online optimization. *Advances in Neural Information Processing Systems*, 20, 2007.
- Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, pages 355–366. Omnipress, 2008.
- Steven De Rooij, Tim Van Erven, Peter D Grünwald, and Wouter M Koolen. Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*, 15(1):1281–1316, 2014.
- Qin Ding, Cho-Jui Hsieh, and James Sharpnack. Robust stochastic linear contextual bandits under adversarial attacks. In *International Conference on Artificial Intelligence and Statistics*, pages 7111–7123. PMLR, 2022.
- Liad Erez and Tomer Koren. Towards best-of-all-worlds online learning with feedback graphs. *Advances in Neural Information Processing Systems*, 34:28511–28521, 2021.
- Pierre Gaillard, Gilles Stoltz, and Tim Van Erven. A second-order bound with excess losses. In *Conference on Learning Theory*, pages 176–196. PMLR, 2014.

- Jiafan He, Dongruo Zhou, Tong Zhang, and Quanquan Gu. Nearly optimal algorithms for linear contextual bandits with adversarial corruptions. In *Advances in Neural Information Processing Systems*, 2022.
- Shinji Ito. Hybrid regret bounds for combinatorial semi-bandits and adversarial linear bandits. *Advances in Neural Information Processing Systems*, 34:2654–2667, 2021a.
- Shinji Ito. Parameter-free multi-armed bandit algorithms with hybrid data-dependent regret bounds. In *Conference on Learning Theory*, pages 2552–2583. PMLR, 2021b.
- Shinji Ito. On optimal robustness to adversarial corruption in online decision problems. *Advances in Neural Information Processing Systems*, 34:7409–7420, 2021c.
- Shinji Ito. Revisiting online submodular minimization: Gap-dependent regret bounds, best of both worlds and adversarial robustness. In *International Conference on Machine Learning*, pages 9678–9694. PMLR, 2022.
- Shinji Ito and Kei Takemura. Best-of-three-worlds linear bandit algorithm with variance-adaptive regret bounds. *arXiv preprint arXiv:2302.12370*, 2023.
- Shinji Ito, Taira Tsuchiya, and Junya Honda. Nearly optimal best-of-both-worlds algorithms for online learning with feedback graphs. In *Advances in Neural Information Processing Systems*, 2022.
- Tiancheng Jin and Haipeng Luo. Simultaneously learning stochastic and adversarial episodic mdps with known transition. *Advances in neural information processing systems*, 33:16557–16566, 2020.
- Tiancheng Jin, Longbo Huang, and Haipeng Luo. The best of both worlds: stochastic and adversarial episodic mdps with unknown transition. *Advances in Neural Information Processing Systems*, 34:20491–20502, 2021.
- Jack Kiefer and Jacob Wolfowitz. The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366, 1960.
- Fang Kong, Yichi Zhou, and Shuai Li. Simultaneously learning stochastic and adversarial bandits with general graph feedback. In *International Conference on Machine Learning*, pages 11473–11482. PMLR, 2022.
- Tor Lattimore and Csaba Szepesvári. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pages 728–737. PMLR, 2017.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, Mengxiao Zhang, and Xiaojin Zhang. Achieving near instance-optimality and minimax-optimality in stochastic and adversarial linear bandits simultaneously. In *International Conference on Machine Learning*, pages 6142–6151. PMLR, 2021.
- Yingkai Li, Edmund Y. Lou, and Liren Shan. Stochastic linear optimization with adversarial corruption. *CoRR*, abs/1909.02109, 2019.

- Haipeng Luo and Robert E Schapire. Achieving all with no parameters: Adanormalhedge. In *Conference on Learning Theory*, pages 1286–1304. PMLR, 2015.
- Jaouad Mourtada and Stéphane Gaïffas. On the optimality of the hedge algorithm in the stochastic regime. *Journal of Machine Learning Research*, 20:1–28, 2019.
- Chloé Rouyer, Dirk van der Hoeven, Nicolò Cesa-Bianchi, and Yevgeny Seldin. A near-optimal best-of-both-worlds algorithm for online learning with feedback graphs. *CoRR*, abs/2206.00557, 2022.
- Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.
- Yevgeny Seldin and Gábor Lugosi. An improved parametrization and analysis of the exp3++ algorithm for stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 1743–1759. PMLR, 2017.
- Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *International Conference on Machine Learning*, pages 1287–1295. PMLR, 2014.
- Taira Tsuchiya, Shinji Ito, and Junya Honda. Best-of-both-worlds algorithms for partial monitoring. *CoRR*, abs/2207.14550, 2022.
- Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Conference On Learning Theory*, pages 1263–1291. PMLR, 2018.
- Heyang Zhao, Dongruo Zhou, and Quanquan Gu. Linear contextual bandits with adversarial corruptions. *CoRR*, abs/2110.12615, 2021.
- Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 467–475. PMLR, 2019.
- Julian Zimmert and Yevgeny Seldin. Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits. *The Journal of Machine Learning Research*, 22(1):1310–1358, 2021.
- Julian Zimmert, Haipeng Luo, and Chen-Yu Wei. Beating stochastic and adversarial semi-bandits optimally and simultaneously. In *International Conference on Machine Learning*, pages 7683–7692. PMLR, 2019.