
Robust Stochastic Linear Contextual Bandits Under Adversarial Attacks

Qin Ding

Department of Statistics
University of California, Davis
qding@ucdavis.edu

Cho-Jui Hsieh

Department of Computer Science
University of California, Los Angeles
chohsieh@cs.ucla.edu

James Sharpnack

Amazon¹
Berkeley, CA
jsharpna@gmail.com

Abstract

Stochastic linear contextual bandit algorithms have substantial applications in practice, such as recommender systems, online advertising, clinical trials, etc. Recent works show that optimal bandit algorithms are vulnerable to adversarial attacks and can fail completely in the presence of attacks. Existing robust bandit algorithms only work for the non-contextual setting under the attack of rewards and cannot improve the robustness in the general and popular contextual bandit environment. In addition, none of the existing methods can defend against attacked context. In this work, we provide the first robust bandit algorithm for stochastic linear contextual bandit setting under a fully adaptive and omniscient attack with sub-linear regret. Our algorithm not only works under the attack of rewards, but also under attacked context. Moreover, it does not need any information about the attack budget or the particular form of the attack. We provide theoretical guarantees for our proposed algorithm and show by experiments that our proposed algorithm improves the robustness against various kinds of popular attacks.

1 INTRODUCTION

A stochastic linear contextual bandit problem models a repeated game between a player and an environment. At each round of the game, the player is given the information of K arms, usually represented by d -dimensional feature vectors, and the player needs to

make a decision by pulling an arm based on past observations. Only the reward associated with the pulled arm is revealed to the player and the relationship between the rewards and features follows a linear model. The goal of the player is to maximize the cumulative reward or minimize the cumulative regret over T rounds. Over the past few decades, bandit algorithms have enjoyed popularity in substantial real-world applications, such as recommender system (Li et al., 2010), online advertising (Schwartz et al., 2017) and clinical trials (Woodrooffe, 1979).

Classic stochastic linear contextual bandit algorithms such as Linear Upper Confidence Bound (LinUCB) (Li et al., 2010) and Linear Thompson Sampling (LinTS) (Agrawal and Goyal, 2013) algorithms can achieve an optimal regret upper bound $\tilde{O}(d\sqrt{T})$, where \tilde{O} hides any poly-logarithm factors. Both LinUCB and LinTS have regret upper bounds that match the lower bound $O(d\sqrt{T})$ up to logarithm factors in infinite-arm problems. In recent years, adversarial attacks have been extensively studied in order to understand the robustness of machine learning algorithms, including bandit algorithms. Attacks were designed for both the contextual bandit problem and the multi-armed bandit (MAB) problem. The types of attacks include both the attacks on rewards and context. For example, in MAB setting, Jun et al. (2018) proposes an oracle attack to modify the rewards of the pulled arms to be ϵ worse than the target arm. In contextual bandit setting, Garcelon et al. (2020) proposes to attack rewards by converting the reward of pulled arms into random noise. For attack of context, Garcelon et al. (2020) proposes to dilate the features of the pulled arms if they are not among the target arms of the attacker. Due to the optimality of LinUCB and LinTS algorithms, they will stop pulling the sub-optimal arms eventually. However, under adversarial attacks, the real optimal arm appears sub-optimal to the algorithm, which makes classic algorithms fail completely under attacks (Garcelon et al.,

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

¹Work done prior to joining Amazon

2020; Jun et al., 2018; Liu and Shroff, 2019; Ma et al., 2018).

Demonstrated by the aforementioned works, adversarial attacks have become major concerns before applying bandit algorithms in practice. For example, in recommender systems, adversarial attacks will trick the algorithm into imperfect recommendations and lead to disappointing user experiences; In online advertising, attacks can be made by triggering some ads and not clicking on these ads, which fools the system to think these ads have low click-through-rate and therefore benefit its competitors (Lykouris et al., 2018). Thus, it is urgent to design a robust bandit algorithm that works well under attacks.

Lots of efforts have been made to design robust bandit algorithms under the attack of rewards with a fixed attack budget C . However, all of them are proposed only for the non-contextual setting. Gupta et al. (2019); Lykouris et al. (2018) established nearly optimal algorithms for the corrupted multi-armed bandit (MAB) problem. Subsequently, Bogunovic et al. (2020) generalizes the idea of Lykouris et al. (2018) to the corruption-tolerant Gaussian process bandit problems. Recently, a robust phase elimination (PE) algorithm (Bogunovic et al., 2021) is derived for a fixed finite-arm stochastic linear bandit problem under attacks. However, PE algorithm needs to pull a fixed feature many times in each phase, implying that it does not work for (or at least non-trivial to be extended to) the contextual bandit setting. As suggested by Bogunovic et al. (2021), the introduction of context significantly complicates the problem and only under a “diversity context” assumption (Bogunovic et al., 2021), a Greedy algorithm can improve the robustness. Adversarial corruptions in linear contextual bandits was considered before by Kapoor et al. (2019), however, Kapoor et al. (2019) can only achieve $O(CT)$ regret which is trivial in bandit environment. To the best of our knowledge, there is no existing work that can achieve sub-linear regret and improve the robustness of contextual bandit algorithms under attacks of rewards in the general environment. Finally, we find that all the previous works only consider the case of corrupted rewards and cannot deal with fully adaptive attacked context, which is also a common type of attack in contextual bandit (Garcelon et al., 2020).

In this work, we consider the stochastic linear contextual bandit algorithm under unknown adversarial attacks with a fixed attack budget, where the feature vectors of arms are contextual and can change over time. In addition, the attackers can be omniscient and fully adaptive, in the sense that the attacker can adjust the attack strategy after observing the player’s decisions at each round. We first show that when the total attack budget C is known, a simple modification of the classic

LinUCB algorithm with an enlarged exploration parameter can obtain a regret upper bound $\tilde{O}(d(C+1)\sqrt{T})$. In contrast, the traditional LinUCB algorithm suffers from linear regret when the attack budget $C \sim O(\log T)$ (Garcelon et al. (2020)). When the total attack budget is unknown, we should learn C online. While learning hyper-parameters online is difficult in bandit setting, the bandit-over-bandit (BOB) algorithm (Cheung et al., 2019) in the non-stationary bandit environment successfully adjusts its sliding-window size online to adapt to changing models. Motivated by the BOB algorithm, we design a similar two-layer bandit algorithm called RobustBandit to deal with unknown attack budget. The high-level idea is to use the top layer to select an appropriate candidate attack budget in the adversarial environment and follow the classic LinUCB policy with increased exploration parameter based on the selected budget in the bottom layer. We show that with the two-layer design, our proposed RobustBandit algorithm can achieve $\tilde{O}(d^{\frac{1}{2}}(C+1)T^{\frac{3}{4}}) + \tilde{O}(d^{\frac{3}{2}}CT^{\frac{1}{4}})$ regret upper bound under unknown adversarial attacks. We also show in experiments that a simple modification of our RobustBandit algorithm achieves substantial improvements in robustness. Our algorithm works no matter the attack is for rewards or context and it is simple to implement. We summarize our contributions below:

1. We propose the RobustBandit algorithm, which is the first algorithm that can achieve sub-linear regret and improve the robustness of linear contextual bandit problems with no additional assumptions, such as the “diversity context” assumption, weak adversary assumption, etc.
2. Our proposed algorithm works for infinite-arm problems under the attack of rewards. It is also the first work that improves the robustness under the fully adaptive attack of context.
3. Our algorithm does not need to know the algorithm used for attack or the attack budget. Furthermore, the method does not need to know whether the attack is for rewards or context under the finite-arm problem.

Notations: We use θ to denote the true model parameter. For a vector $x \in \mathbb{R}^d$, we use $\|x\|$ to denote its l_2 norm and $\|x\|_A = \sqrt{x^T A x}$ to denote its weighted l_2 norm associated with a positive-definite matrix $A \in \mathbb{R}^{d \times d}$. Denote $[n] := \{1, 2, \dots, n\}$ and $\lceil b \rceil$ as the minimum integer such that $\lceil b \rceil \geq b$.

2 RELATED WORK

In this section, we will focus on discussing previous works on improving the robustness of bandit algorithms.

Note that very few works consider the attacks of context. Yang and Ren (2021) considers imperfect context and the attack of context cannot be adjusted by the attacker adaptively at each round based on previous observations, which makes the problem much easier. In the following, all the previous works discussed only consider the attacks of rewards.

In case of multi-armed bandit (MAB) problems under adversarial corruptions, Lykouris et al. (2018) designs a multi-layer active arm elimination race algorithm (MLAER) when the total attack budget is unknown. It obtains multiplicative regret upper bound $O(\frac{KC+\log T}{\Delta} \log(KT))$, where Δ is the mean reward gap between the optimal and sub-optimal arms. Gupta et al. (2019) improves the robustness of MAB problems by a novel restarting algorithm, BARBAR, and derives a near-optimal regret upper bound $O(KC + \frac{(\log T)^2}{\Delta})$, where the regret depends on C linearly. The key is to ensure the algorithm never permanently eliminates an arm and so the seemingly-not-so-good arms always get a small amount of resource. Both MLAER and BARBAR only work for a weaker adversary that cannot adjust the attack strategy based on the player’s decisions. This limits their applications when the attacks are fully adaptive, which is a popular assumption when designing bandit attacks (Garcelon et al., 2020; Jun et al., 2018; Liu and Shroff, 2019; Ma et al., 2018).

In the linear bandit setting with non-contextual features, Bogunovic et al. (2020) considers the Gaussian process bandit problems with a function of bounded RKHS norm, and extends the idea in MLAER to Fast-Slow GP-UCB algorithm with a $\tilde{O}(d(C+1)\sqrt{T})$ regret upper bound. Note that Bogunovic et al. (2020) assumes that the attacks cannot be fully adaptive, which makes it easier to defend against the corruptions. Built upon Gupta et al. (2019), Li et al. (2019) proposes support bias exploration algorithm (SBE) for stochastic linear bandits and derives a $O(\frac{d^{\frac{5}{2}}C \log T}{\Delta} + \frac{d^6(\log T)^2}{\Delta^2})$ regret upper bound. Most recently, Bogunovic et al. (2021) proposes a phase elimination algorithm (PE) to improve the robustness of linear bandit and obtains $\tilde{O}(\sqrt{dT} + C^2 + Cd^{\frac{3}{2}} \log T)$ regret upper bound. In each phase, the PE algorithm computes a design over the set of potential optimal arms and plays each arm in this set with number of times proportional to the computed design. In the meantime, it requires that every arm in its support is played at least some minimal number of times. Since PE needs to play an arm with a fixed feature many times, it does not work in the contextual bandit environment, where the features of arms are contextual and can change over time, which is a more general and popular problem in practice.

For the contextual bandit problem under attacks of

rewards, Bogunovic et al. (2021) shows that under a “diversity context” assumption (Bogunovic et al., 2021; Ding et al., 2021; Wu et al., 2020), a simple Greedy algorithm can obtain $\tilde{O}(\frac{1}{\lambda_0}(\sqrt{dT} + C \log T))$ regret upper bound in the presence of attacks. Here, λ_0 is the degree of the diversity, which is defined as the minimum eigenvalue of the covariance matrix of the feature vectors lying in any half space. However, due to the partial feedback setting, contextual feature vectors of the pulled arms are highly biased and so the “diversity context” assumption rarely holds in bandit environment. When “diversity context” assumption does not hold, $\lambda_0 = 0$ and this makes the regret bound of the Greedy algorithm meaningless. Moreover, under similar diversity assumptions (Wu et al., 2020), LinUCB was shown to obtain a constant regret, which makes the optimality of the Greedy algorithm under attacks questionable. Finally, we emphasize that there is no existing work that considers contextual bandit problems with fully adaptive attacked context, even under the “diversity context” assumption. We summarize the properties of different algorithms in Table 1.

3 PROBLEM SETTING

We study a K -armed stochastic linear contextual bandit setting under corrupted observations, where K can be infinite. Two types of adversarial attacks listed below are considered in this paper.

Attack rewards. At each round $t \in [T]$, the player is given a set of context including K feature vectors $\mathcal{A}_t = \{x_{t,a} | a \in [K]\} \subset \mathbb{R}^d$, where $x_{t,a}$ represents the true information of arm a at round t . Based on previous observations and \mathcal{A}_t , the player pulls an arm a_t at round t . For ease of notation, we denote $X_t := x_{t,a_t}$. The attacker observes a_t and assigns the attack c_t , where c_t may depend on a_t and other possible information. The player receives an attacked reward $\tilde{Y}_t = Y_t + c_t$, where $Y_t = X_t^T \theta + \epsilon_t$ is the true stochastic reward. Here, $\theta \in \mathbb{R}^d$ is the unknown true model parameter and ϵ_t is a random noise.

Attack context. At each round t before the player observes \mathcal{A}_t , the attacker calculates an attack based on past information and the true context \mathcal{A}_t at the current round. The attacker then modifies the true context into the attacked context $\tilde{x}_{t,a}$. The attack cost paid by the attacker at this round is defined as $c_t = \sum_{a=1}^K \|\tilde{x}_{t,a} - x_{t,a}\|$. The player can only observe the contextual features after the attack, i.e., $\tilde{\mathcal{A}}_t := \{\tilde{x}_{t,a} | a \in [K]\}$ and pull arm a_t according to the attacked context. The observed reward for this pulled arm is the true reward $Y_t = X_t^T \theta + \epsilon_t$.

In the following, we will use $\tilde{\mathcal{A}}_t$, \tilde{X}_t and \tilde{Y}_t to de-

Table 1: Comparison of different algorithms with unknown total attack budget. Column “Fully adaptive” means whether the algorithm works when the attacker is fully adaptive and can adapt his attack strategy based on the player’s decisions at the current round.

Algorithm	Contextual	Infinite-arm	Fully adaptive	Attacked Context
MLAER (Lykouris et al., 2018) (MAB only)	✗	✗	✗	✗
BARBAR (Gupta et al., 2019) (MAB only)	✗	✗	✗	✗
SBE (Li et al., 2019)	✗	✓	✗	✗
Fast-Slow GP-UCB (Bogunovic et al., 2020)	✗	✓	✗	✗
PE (Bogunovic et al., 2021)	✗	✗	✓	✗
Greedy ² (Bogunovic et al., 2021)	✓	✓	✓	✗
RobustBandit (This work)	✓	✓	✓	✓

note the player’s observed features set, observed features and rewards of the pulled arm at round t in the attacked environment respectively. Under the attack of rewards, $\tilde{X}_t = X_t$ and $\tilde{A}_t = A_t$. While under the attack of context, $\tilde{Y}_t = Y_t$. We use $\mathcal{F}_t = \sigma(a_1, \dots, a_t, \tilde{A}_1, \dots, \tilde{A}_t, \tilde{Y}_1, \dots, \tilde{Y}_t, c_1, \dots, c_t)$ to denote the σ -algebra generated by all the information up to round t . In both the above two attacks, the random noises ϵ_t are independent zero-mean σ -sub-Gaussian random variables with $\mathbb{E}[e^{b\epsilon_t} | \mathcal{F}_{t-1}] \leq e^{\frac{b^2 \sigma^2}{2}}$ for all t and $b \in \mathbb{R}$. Without loss of generality, we assume that $\|x_{t,a}\| \leq 1$ for all $t \in [T]$ and $a \in [K]$, $\|\theta\| \leq 1$ and the true observed reward $Y_t \in [0, 1]$. We do not make any assumption on the attacker. The attacker can be fully adaptive in our model and even have complete knowledge of the bandit algorithm the player is using and the bandit environment, i.e., θ, σ, A_t , random noises ϵ_t , the pulled arm a_t and so on. However, to make the attacker invisible to the player, the attack strategy should obey the rules to make sure $\tilde{Y}_t \in [0, 1]$ and $\|\tilde{x}_{t,a}\| \leq 1$. In addition, the attacker has an attack budget C which limits the total amount of perturbations by $\sum_{t=1}^T |c_t| \leq C$. Denote $a_t^* = \operatorname{argmax}_{a \in [K]} x_{t,a}^T \theta$ as the optimal arm at round t and $x_{t,*} = x_{t,a_t^*}$ the corresponding true feature vector. The goal of the player in both cases is to minimize the cumulative regret defined below

$$R(T) = \sum_{t=1}^T (x_{t,*} - X_t)^T \theta.$$

4 STOCHASTIC LINEAR CONTEXTUAL BANDIT UNDER KNOWN ATTACK BUDGET

In this section, we first show that if the player knows the attack budget C , then a simple modification of Lin-

²The results only hold for Greedy algorithm when “diversity context” assumption (Bogunovic et al., 2021) is satisfied.

UCB algorithm with an enlarged exploration parameter performs well under corruptions. In the non-corrupted setting, LinUCB (see Algorithm 2 in Appendix for more details) calculates the ridge regression (with regularization parameter λ) estimator $\hat{\theta}_t^0$ using all the past observations $\{(X_s, Y_s)\}_{s=1}^{t-1}$. By Abbasi-Yadkori et al. (2011), for all $x \in \mathbb{R}^d$, it satisfies $|x^T (\hat{\theta}_t^0 - \theta)| \leq \beta_t \|x\|_{V_t^{-1}}$ with probability at least $1 - \delta$, where

$$\beta_t = \sigma \sqrt{d \log \left(\frac{1+t/\lambda}{\delta} \right)} + \sqrt{\lambda} \quad (1)$$

$$V_t = \lambda I_d + \sum_{s=1}^{t-1} X_s X_s^T. \quad (2)$$

LinUCB policy then follows by pulling the arm with the biggest upper confidence bound with the exploration parameter β_t as

$$a_t = \operatorname{argmax}_a x_{t,a}^T \hat{\theta}_t^0 + \beta_t \|x_{t,a}\|_{V_t^{-1}}.$$

Under adversarial attacks, we cannot observe the true context X_t or true rewards Y_t . With the attacked context \tilde{X}_t or attacked rewards \tilde{Y}_t , we can only calculate the corrupted ridge regression estimator from the observation history $\{(\tilde{X}_s, \tilde{Y}_s)\}_{s=1}^{t-1}$ and denote it as $\hat{\theta}_t$. Define $\tilde{V}_t = \lambda I_d + \sum_{s=1}^{t-1} \tilde{X}_s \tilde{X}_s^T$. Take the attack of rewards as an example, and note that $\tilde{X}_t = X_t$, $\tilde{Y}_t = Y_t + c_t$ under attack of rewards, the difference $\hat{\theta}_t - \theta$ can be decomposed as in the following equation:

$$\begin{aligned} \hat{\theta}_t - \theta &= \tilde{V}_t^{-1} \sum_{s=1}^{t-1} \tilde{X}_s \tilde{Y}_s - \theta \\ &= \tilde{V}_t^{-1} \sum_{s=1}^{t-1} \tilde{X}_s (Y_s + c_s) - \theta \\ &= \hat{\theta}_t^0 - \theta + \tilde{V}_t^{-1} \sum_{s=1}^{t-1} c_s \tilde{X}_s. \end{aligned} \quad (3)$$

Under attacked context, the analysis of $\hat{\theta}_t - \theta$ is more tricky since the player can only observe attacked con-

text. However, by utilizing the concentration results in Abbasi-Yadkori et al. (2011), we can derive the following lemma that holds no matter the attack is on rewards or context.

Lemma 1. Define $\hat{\theta}_t$ as the solution to the ridge regression with regularization parameter λ with observations $\{(\tilde{X}_t, \tilde{Y}_t)\}_{s=1}^{t-1}$. It satisfies the following inequality for all vector $x \in \mathbb{R}^d$ with probability at least $1 - \delta$ no matter the attack is for rewards or context.

$$|x^T(\hat{\theta}_t - \theta)| \leq (\beta_t + \gamma_t C) \|x\|_{\tilde{V}_t^{-1}}, \quad (4)$$

where $\beta_t = \sigma \sqrt{d \log \left(\frac{1+t/\lambda}{\delta} \right)} + \sqrt{\lambda}$ and $\gamma_t = \sqrt{\sum_{s=1}^{t-1} \|\tilde{X}_s\|_{\tilde{V}_s^{-1}}^2}$.

Based on Lemma 1, if the player knows an upper bound C' of the attack budget such that $C \leq C'$, then we can simply enlarge the exploration parameter of LinUCB from β_t to $\beta_t + \gamma_t C'$, i.e., pull the arm

$$a_t = \operatorname{argmax}_{a=1, \dots, K} \tilde{x}_{t,a}^T \hat{\theta}_t + (\beta_t + \gamma_t C') \|\tilde{x}_{t,a}\|_{\tilde{V}_t^{-1}}. \quad (5)$$

By doing so, we can obtain a regret upper bound in Theorem 1 below.

Theorem 1. Let C' be a constant that satisfies $C' \geq C$, where C is the total attack budget. Denote $R_r(T)$ and $R_c(T)$ as the cumulative regret under attack of rewards and context respectively. Under the attack of rewards or context, if we pull arm a_t as defined in Equation 5 at round t , then with probability at least $1 - \delta$ the cumulative regret satisfies

$$R_r(T) \leq (2\beta_T + \gamma_T(C + C')) \gamma_{T+1} \sqrt{T} \quad (6)$$

$$R_c(T) \leq (2\beta_T + \gamma_T(C + C')) \gamma_{T+1} \sqrt{T} + C \quad (7)$$

Under both cases,

$$R(T) \leq \tilde{O}(d\sqrt{T}) + \tilde{O}(d(C + C')\sqrt{T}). \quad (8)$$

5 PROPOSED ALGORITHM - RobustBandit

5.1 Stochastic contextual bandit under unknown attacks

In the previous section, we have shown that when the total attack budget has a known and tight upper bound, then a simple LinUCB algorithm with an enlarged exploration parameter can improve the robustness under adversarial attacks. However, if the attack budget C is agnostic, then we have to adaptively enlarge the exploration parameter to make it robust to all possible amounts of attacks.

Algorithm 1 RobustBandit Algorithm

Input: T , β_t as in Equation 1, epoch length H .

- 1: Initialize candidate attack budget set $J = \{0\} \cup \{2^j\}_{j=0}^{\lceil \log_2 2KT \rceil}$.
 - 2: Initialize exponential weights $w_j(1) = 1$ for $j = 1, \dots, |J|$.
 - 3: Initialize the exploration parameter for EXP3 as $\alpha = \min \left\{ 1, \sqrt{\frac{|J| \log |J|}{(e-1) \lceil \frac{T}{H} \rceil}} \right\}$.
 - 4: **for** $i = 1$ **to** $\lceil \frac{T}{H} \rceil$ **do**
 - 5: Initialize $\hat{\theta}_t = 0$, $\tilde{V}_t = \lambda I_d$.
 - 6: Update probability distribution for pulling candidates in J as

$$p_j(i) = \frac{\alpha}{|J|} + (1 - \alpha) \frac{w_j(t)}{\sum_{i=1}^{|J|} w_i(t)}, \forall j = 1, \dots, |J|.$$
 - 7: Draw $j_i \leftarrow j \in [|J|]$ with probability $p_j(i)$.
 - 8: **for** $t = (i - 1)H + 1$ **to** $\min\{T, iH\}$ **do**
 - 9: Observe (attacked) contextual feature vectors $\tilde{x}_{t,a}$.
 - 10: Pull arm according to the following equation

$$a_t = \operatorname{argmax}_{a=1, \dots, K} \tilde{x}_{t,a}^T \hat{\theta}_t + (\beta_t + \gamma_t J_{j_i}) \|\tilde{x}_{t,a}\|_{\tilde{V}_t^{-1}}$$
 - 11: Observe reward \tilde{Y}_t .
 - 12: Update LinUCB components, i.e., $\tilde{V}_{t+1} = \tilde{V}_t + \tilde{X}_t \tilde{X}_t^T$ and $\hat{\theta}_{t+1} = \tilde{V}_{t+1}^{-1} \sum_{s=1}^t \tilde{X}_s \tilde{Y}_s$.
 - 13: **end for**
 - 14: Update EXP3 components: $\hat{y}_t(j) \leftarrow 0$ for all $j \neq j_i$, $\hat{y}_t(j) = \sum_{t=(i-1)H+1}^t \tilde{Y}_t / p_j(i)$ if $j = j_i$ and

$$w_j(t+1) = w_j(t) \times \exp \left(\frac{\alpha}{|J|} \hat{y}_t(j) \right).$$
 - 15: **end for**
-

From Theorem 1, we can see that the key to design a robust algorithm is to find an appropriate upper bound C' on the total attack budget C . Since the regret upper bound is $\tilde{O}(d(C + C')\sqrt{T})$ in Theorem 1, a good choice of C' should satisfy $C \leq C' \leq mC$, where $m \geq 1$ is a constant that should be as small as possible.

The intuition behind our method is that we accumulate experience about what the best choice of C' is. In the meantime, we cut the whole time horizon into epochs and restart the algorithm when we need to try a different choice of C' . The effectiveness of this restarting technique not only helps adaptively learn the best C' , but also voids the attacker's efforts in the previous epoch. While the attack budget of the attacker is being exhausted after restarting the epochs, our

algorithm will be able to learn the best arms gradually through the epochs.

Assume we restart the algorithm every H rounds, then there will be $L = \lceil \frac{T}{H} \rceil$ epochs and the last epoch does not necessarily have H rounds. Since both Y_t and $\|x_{t,a}\|$ are upper bounded by 1 and the attacker should make sure the corruption is undetectable, which means that $\tilde{Y}_t \in [0, 1]$ and $\|\tilde{x}_{t,a}\| \leq 1$, a natural upper bound on the attack budget spent by the attacker is $C \leq T$ under attack of rewards and $C \leq 2KT$ under attack of context. Assume C_i is the attack spent in the i -th epoch by the attacker, then $C_i \leq \min(C, H)$ under attack of rewards and $C_i \leq \min(C, 2KH)$ under attack of context for all epochs i .

If we denote $J = \{0\} \cup \{2^j\}_{j=0}^{\lceil \log_2 2KT \rceil}$ as a set of possible choices of C' . Then since $C_i \leq 2KH$, there must exist a $C^* \in J$ such that $C_i \leq C^* \leq 2C$ for all i , which makes C^* a proper choice for C' at the every epoch. Due to the existence of adversarial corruptions, it is natural to treat the choices of C' as another adversarial multi-armed bandit (MAB) problem Auer et al. (2002). This motivates our proposed algorithm. This idea is also very similar to the Bandit-Over-Bandit idea Cheung et al. (2019), where the authors use an EXP3 algorithm over the LinUCB algorithm in non-stationary bandit problems.

To be more specific, we propose a two-layer robust bandit algorithm. In the top layer, the algorithm uses an adversarial MAB policy, namely EXP3 (Auer et al., 2002), to pull candidate J_{j_i} from the set J in the beginning of every epoch. Here j_i is the index of the pulled candidate at the i -th epoch and J_j is the j -th element in the set J . In the i -th epoch, where $t = (i-1)H + 1, \dots, iH$, in the bottom layer, the algorithm enlarges the exploration parameter of LinUCB based on $C' = J_{j_i}$ and pulls arm a_t at round t according to

$$a_t = \operatorname{argmax}_{a=1, \dots, K} \tilde{x}_{t,a}^T \hat{\theta}_t + (\beta_t + \gamma_t J_{j_i}) \|\tilde{x}_{t,a}\|_{\tilde{V}_t^{-1}}. \quad (9)$$

after observing the (attacked) context. After receiving the (attacked) reward of the pulled arm, this reward is fed to the bottom layers to update the components of LinUCB. When an epoch ends, we restart the LinUCB algorithm and use the rewards accumulated during the previous epoch to update the EXP3 algorithms. The EXP3 algorithm will then update a decision about which C' to choose in the next epoch. Details can be found in Algorithm 1.

We emphasize here that our algorithm does not need to know whether the attack is on rewards or context in order to make decisions if there are finite arms. Our algorithm also works for infinite-arm problem under attacks of rewards by simply setting $J = \{0\} \cup \{2^j\}_{j=0}^{\lceil \log_2 T \rceil}$ since $C \leq T$ under attack of rewards.

5.2 Regret analysis

In this section, we formally analyze the regret of our proposed Algorithm 1. Proofs are similar to Cheung et al. (2019). For a round t in the i -th epoch, let $X_t(J_j)$ denote the true feature vector of the arm pulled at round t if the chosen C' in the beginning of the i -th epoch is J_j and the enlarged exploration parameter is based on J_j for all rounds in this epoch. Denote $\tilde{X}_t(J_j)$ as the corresponding (attacked) feature vector. Let C^* be the element in J such that $C_i \leq C^* \leq 2C$. We decompose the cumulative regret into two quantities below:

$$\begin{aligned} \mathbb{E}[R(T)] &= \mathbb{E} \left[\sum_{t=1}^T (x_{t,*}^T \theta - X_t^T \theta) \right] \\ &= \underbrace{\mathbb{E} \left[\sum_{t=1}^T (x_{t,*}^T \theta - X_t(C^*)^T \theta) \right]}_{\text{Quantity (A)}} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{t=1}^T (X_t(C^*)^T \theta - X_t(J_{j_i})^T \theta) \right]}_{\text{Quantity (B)}}. \quad (10) \end{aligned}$$

We bound the regret from Quantity (A) in the following lemma.

Lemma 2. *Quantity (A) is upper bounded by $\tilde{O}(\frac{dT}{\sqrt{H}} + dC\sqrt{H} + \frac{dCT}{\sqrt{H}})$.*

Proof. Quantity (A) can be further decomposed into epochs

$$\sum_{i=1}^{\lceil \frac{T}{H} \rceil} \mathbb{E} \left[\sum_{t=(i-1)H+1}^{\min(T, iH)} (x_{t,*}^T \theta - X_t(C^*)^T \theta) \right]$$

Under attack of rewards, according to Theorem 1 and the choice of C^* , the regret of each epoch from Quantity (A) is upper bounded by $\tilde{O}(d\sqrt{H}) + \tilde{O}(d(C_i + C^*)\sqrt{H})$, where C_i is the real attack budget used by the attacker in the i -th epoch. According to Equation 6, 7, and summing over $\lceil \frac{T}{H} \rceil$ epochs, we get

$$\begin{aligned} \text{Quantity (A)} &\leq \sum_{i=1}^{\lceil \frac{T}{H} \rceil} (2\beta_T + \gamma_T(C_i + C^*))\gamma_{T+1}\sqrt{H} \\ &= \frac{2\beta_T\gamma_{T+1}T}{\sqrt{H}} + \gamma_T\gamma_{T+1}(C\sqrt{H} + \frac{2CT}{\sqrt{H}}) \end{aligned}$$

Similarly, under attack of context, we have

$$\begin{aligned} \text{Quantity (A)} &\leq \frac{2\beta_T\gamma_{T+1}T}{\sqrt{H}} + \gamma_T\gamma_{T+1}(C\sqrt{H} + \frac{2CT}{\sqrt{H}}) + C \end{aligned}$$

We bound γ_{T+1} in Equation 14 in Appendix, and according to Equation 14, we have

$$\text{Quantity (A)} = \tilde{O}\left(\frac{dT}{\sqrt{H}} + dC\sqrt{H} + \frac{dCT}{\sqrt{H}}\right).$$

□

The regret in Quantity (B) comes from the learning of C^* using EXP3 algorithm. From the result in Auer et al. (2002), we have the following lemma and proof is deferred to Appendix.

Lemma 3. *For a random sequence of candidates $\{J_{j_1}, \dots, J_{j_L}\}$ pulled by the EXP3 layer in each epoch in Algorithm 1, where $L = \lceil T/H \rceil$, no matter the attack is for rewards or context, we have*

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T X_t(C^*)^T \theta \right] - E \left[\sum_{t=1}^T X_t(J_{j_t})^T \theta \right] \\ & \leq 2\sqrt{e-1} \sqrt{|TH| |J| \log |J|} = \tilde{O}(\sqrt{TH}). \end{aligned}$$

The regret upper bound of our proposed algorithm follows directly from Lemma 2 and Lemma 3. We show the conclusion in Theorem 2 below.

Theorem 2. *Denote $\nu = \sqrt{2d \log \frac{d\lambda+T}{d\lambda}}$. For any error probability $\delta \in (0, 1)$, Algorithm 1 with $\beta_t = \sigma \sqrt{d \log \left(\frac{1+t/\lambda}{\delta} \right)} + \sqrt{\lambda}$, $\alpha = \min \left\{ 1, \sqrt{\frac{|J| \log |J|}{(e-1)T}} \right\}$ and $\gamma_t = \sqrt{\sum_{s=1}^{t-1} \|\tilde{X}_s\|_{V_s^{-1}}^2}$ has the following expected total regret:*

$$\mathbb{E}[R(T)] \leq \tilde{O}\left(\frac{dT}{\sqrt{H}} + dC\sqrt{H} + \frac{dCT}{\sqrt{H}}\right) + \tilde{O}(\sqrt{TH})$$

Remark 1. *Since the choice of H should not depend on C and we want to minimize the cumulative regret with respect to H , we can choose $H = \tilde{O}(d\sqrt{T})$. The cumulative regret of our algorithm then satisfies*

$$\mathbb{E}[R(T)] \leq \tilde{O}(d^{\frac{1}{2}}(C+1)T^{\frac{3}{4}}) + \tilde{O}(d^{\frac{3}{2}}CT^{\frac{1}{4}}).$$

Remark 2. *In the non-corrupted setting or when the attack budget $C \leq \tilde{O}(1)$, the regret upper bound of our algorithm is still sub-linear. In contrast, classic algorithms such as LinUCB and LinTS fail completely under such attack budget and suffers from linear regret Garcelon et al. (2020).*

Remark 3. *Our proposed algorithm is the first one in existing works that can deal with adversarial attacks under linear contextual bandit environment with sub-linear regret. Our algorithm can improve the robustness under either the attacks on rewards or context. To the best of our knowledge, there is no existing work that improves the robustness under attack of context. Note*

that our proposed algorithm does not need to know anything about the attack budget, whether the attack is for rewards or for context in order to achieve this regret upper bound for finite-arm problems.

6 EXPERIMENTAL RESULTS

In this section, we show by experiments on both synthetic and real datasets that our proposed algorithm is robust to adversarial corruptions. To the best of our knowledge, there is no existing robust algorithm that works for stochastic contextual bandit with changing feature vectors in the general environment. We compare our proposed algorithm with LinUCB, LinTS and Greedy algorithm.

For all the experiments below, we set the number of total rounds $T = 10^6$. At each round, there are $K = 20$ arms and the dimension of features is $d = 10$. We set the total attack budget as $C = 100$. We choose the hyper-parameters exactly as Theorem 2 suggests. Suggested by the explicit formulae of regret bounds combined from Theorem 1 and Lemma 3, we set epoch length $H = \frac{\beta_T \gamma_{T+1}}{\sqrt{e-1}} \sqrt{T}$. Similar to Garcelon et al. (2020), we set error probability $\delta = 0.01$ and regularization parameter $\lambda = 0.1$. If arm a is pulled at round t , the true sample reward is drawn from $\mathcal{N}(x_{t,a}^T \theta, 0.01)$. An appropriate attack (See Section 6.1 for details of attacks) is assigned to the reward or the context and only the corrupted observation is shown to the player. Results reported below are averaged over 10 independently repeated experiments. We run the experiments on three datasets listed below:

1). Simulation: We simulate all the feature vectors $x_{t,a}$ and the model parameter θ from $\text{Uniform}(-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}})$ to make sure $|x_{t,a}^T \theta| \leq 1$. Mean rewards $\mu_{t,a} = x_{t,a}^T \theta$ are then transformed by $\mu_{t,a} \leftarrow \frac{x_{t,a}^T \theta + 1}{2}$ to make sure the final mean rewards $\mu_{t,a} \in [0, 1]$.

2). MovieLens 100K dataset: This dataset has 100K ratings of 943 users on 1,682 movies. We use LIBPMF (Yu et al., 2012, 2014) to perform matrix factorization with $d = 10$ on the ratings data and get the features matrices for both users and movies. The movies are then treated as arms and for each repeated experiment, we randomly select $K = 20$ movies. The model parameter θ is defined as the average of randomly selected 100 users' feature vectors. The mean reward is then defined as $\mu_{t,a} = x_{t,a}^T \theta$ and shifted to $[0, 1]$.

3). Netflix dataset: Netflix dataset contains ratings on 17,770 items from 480,189 users. The preprocessing of this data is almost the same as the MovieLens dataset, except now that we use only the items having more




Figure 1: Comparison of LinUCB, LinTS, Greedy, RobustBandit and BOB-No-Restart algorithms under adversarial attacks on rewards and context.

than 10,000 ratings and this leaves us with 2042 items.

6.1 Type of Attacks

Details of the attacks considered in the experiments are presented below. For all attacks, the attacker cannot attack anymore if the attack budget C is exhausted.

1). **Garcelon**: This attack is based on the reward attack in Garcelon et al. (2020). Whenever the pulled arm is not among the attacker’s target arms, the attacker modifies the reward into a random noise $\mathcal{N}(0, 0.01)$.

2). **Oracle**: This attack is a simple modification from Jun et al. (2018). It pushes the reward of an arm to some margin ϵ_0 below the reward of the worst arm and does not attack if such condition is already attained. We use $\epsilon_0 = 0.01$ in the experiments.

For Garcelon and Oracle attacks in the above, the attacker attacks the context or rewards only when the pulled arm is among the top $N = 0.5K$ arms. This means that the attacker is malicious in the sense that it will attack when the pulled arm is among the best half arms.

6.2 BOB-No-Restart: An empirically good algorithm based on the proposed method

Although the restarting technique gives theoretically sub-linear regret bounds and it is necessary for the analysis of EXP3 layer to work, we found that restarting LinUCB usually abandons useful information in the past epochs. Empirically, a two-layer bandit struc-

ture without restarting performs exceptionally well (See Figure 1) and we call this algorithm Bandit-Over-Bandit-No-Restart (BOB-No-Restart). To be more specific, BOB-No-Restart still has two layers and the top layer still guides the selection of C' to help enlarge the exploration parameter. However, instead of choosing C' every epoch, the EXP3 layer now chooses C' every round using the immediate observed reward from last round. The chosen C' now serves directly to the bottom LinUCB layer every round. We present more details of BOB-No-Restart in Algorithm 3 in Appendix.

6.3 Results

We show the plots of averaged cumulative regret against number of rounds in Figure 1. From left to right in Figure 1, the plots are for simulations, Movielens and Netflix respectively. From the 1st row to the 2nd row, the attack types are Garcelon and Oracle respectively. From the plots, we can see that our proposed RobustBandit algorithm consistently outperforms LinUCB, LinTS and Greedy algorithms under adversarial attacks and improves the robustness. Moreover, BOB-No-Restart algorithm constantly improves robustness significantly compared with all algorithms.

7 CONCLUSION and FUTURE WORK

In this work, we study the robustness of stochastic linear contextual bandit problems. We propose a first robust bandit algorithm for the contextual setting un-

der adversarial attacks which can achieve sub-linear regrets. Our algorithm is also the first one that can deal with attacked context. Our algorithm does not need to know the attack budget or format in order to make decisions under finite-arm problems, while the attacker can be omniscient and fully adaptive after observing the player’s decisions. The proposed algorithm has regret upper bound $\tilde{O}(d^{\frac{1}{2}}(C+1)T^{\frac{3}{4}}) + \tilde{O}(d^{\frac{3}{2}}CT^{\frac{1}{4}})$. Extensive experiments show that our algorithm improves the robustness of bandit problems under attacks.

Future Work: First of all, as proved by Garcelon et al. (2020), if an algorithm has $f(T)$ regret under non-corrupted cases, the algorithm must suffer from linear regret when $C = \Omega(f(T))$. This shows that the regret lower bound does not solely relies on C or T . Instead, the regret lower bound relies on interactions of both C and T . We think a possible future direction is to derive the optimal balance between C and T in regret lower bounds for linear contextual bandit algorithm under attacks. Secondly, an optimal algorithm that matches this lower bound is of substantial interest. Lastly, the BOB-No-Restart algorithm has exceptionally good performance in experiments although it does not have theoretical guarantees. It is also interesting to investigate the theoretical behavior of this algorithm.

References

- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1): 48–77, 2002.
- Ilija Bogunovic, Andreas Krause, and Jonathan Scarlett. Corruption-tolerant gaussian process bandit optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 1071–1081. PMLR, 2020.
- 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.
- Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Learning to optimize under non-stationarity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1079–1087, 2019.
- Qin Ding, Cho-Jui Hsieh, and James Sharpnack. An efficient algorithm for generalized linear bandit: Online stochastic gradient descent and thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 1585–1593. PMLR, 2021.
- Evrard Garcelon, Baptiste Roziere, Laurent Meunier, Jean Tarbouriech, Olivier Teytaud, Alessandro Lazaric, and Matteo Pirodda. Adversarial attacks on linear contextual bandits. In *Advances in Neural Information Processing Systems*, volume 33, pages 14362–14373, 2020.
- Anupam Gupta, Tomer Koren, and Kunal Talwar. Better algorithms for stochastic bandits with adversarial corruptions. In *Conference on Learning Theory*, pages 1562–1578. PMLR, 2019.
- Kwang-Sung Jun, Lihong Li, Yuzhe Ma, and Jerry Zhu. Adversarial attacks on stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 3640–3649, 2018.
- Sayash Kapoor, Kumar Kshitij Patel, and Purushottam Kar. Corruption-tolerant bandit learning. *Machine Learning*, 108(4):687–715, 2019.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- Yingkai Li, Edmund Y Lou, and Liren Shan. Stochastic linear optimization with adversarial corruption. *arXiv preprint arXiv:1909.02109*, 2019.
- Fang Liu and Ness Shroff. Data poisoning attacks on stochastic bandits. In *International Conference on Machine Learning*, pages 4042–4050. PMLR, 2019.
- Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 114–122, 2018.
- Yuzhe Ma, Kwang-Sung Jun, Lihong Li, and Xiaojin Zhu. Data poisoning attacks in contextual bandits. In *International Conference on Decision and Game Theory for Security*, pages 186–204. Springer, 2018.
- Eric M Schwartz, Eric T Bradlow, and Peter S Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017.
- Michael Woodroffe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.

Weiqiang Wu, Jing Yang, and Cong Shen. Stochastic linear contextual bandits with diverse contexts. In *International Conference on Artificial Intelligence and Statistics*, pages 2392–2401. PMLR, 2020.

Jianyi Yang and Shaolei Ren. Robust bandit learning with imperfect context. *arXiv preprint arXiv:2102.05018*, 2021.

Hsiang-Fu Yu, Cho-Jui Hsieh, Si Si, and Inderjit S. Dhillon. Scalable coordinate descent approaches to parallel matrix factorization for recommender systems. In *IEEE International Conference of Data Mining*, 2012.

Hsiang-Fu Yu, Cho-Jui Hsieh, Si Si, and Inderjit S Dhillon. Parallel matrix factorization for recommender systems. *Knowledge and Information Systems*, 41(3):793–819, 2014.

8 SUPPLEMENTARY MATERIAL

8.1 Proof of Lemma 1

Proof. We prove the result holds for attack of rewards and context respectively below.

Attack rewards: Note that in this case, $\tilde{X}_t = X_t$ and $\tilde{V}_t = V_t$, so we have from Equation 3 that

$$\|\hat{\theta}_t - \theta\|_{V_t} \leq \|\hat{\theta}_t^0 - \theta\|_{V_t} + \|V_t^{-1} \sum_{s=1}^{t-1} c_s X_s\|_{V_t} \quad (11)$$

From the result in Abbasi-Yadkori et al. (2011), $\|\hat{\theta}_t^0 - \theta\|_{V_t} \leq \beta_t$ with probability at least $1 - \delta$. For the second term in the above inequality,

$$\|V_t^{-1} \sum_{s=1}^{t-1} c_s X_s\|_{V_t} \leq \sum_{s=1}^{t-1} |c_s| \sqrt{X_s^T V_t^{-1} X_s} \leq \sqrt{\sum_{s=1}^{t-1} |c_s|^2} \sqrt{\sum_{s=1}^{t-1} X_s^T V_s^{-1} X_s} \leq \gamma_t C.$$

Attack context: Denote $c_s^* = \tilde{X}_s - X_s$. Then because of the bounded attack budget, $\sum_{t=1}^T \|c_s^*\| \leq C$. For the ridge regression estimator under corrupted context, note that $\tilde{Y}_t = Y_t$ in this case, we have

$$\begin{aligned} \hat{\theta}_t - \theta &= \tilde{V}_t^{-1} \sum_{s=1}^{t-1} \tilde{X}_s Y_s - \theta \\ &= \tilde{V}_t^{-1} \sum_{s=1}^{t-1} (X_s + c_s^*) (X_s^T \theta + \epsilon_s) - \tilde{V}_t^{-1} (\lambda I_d + \sum_{s=1}^{t-1} \tilde{X}_s \tilde{X}_s^T) \theta \\ &= -\tilde{V}_t^{-1} \sum_{s=1}^{t-1} (X_s + c_s^*) c_s^{*T} \theta + \tilde{V}_t^{-1} \sum_{s=1}^{t-1} \tilde{X}_s \epsilon_s - \lambda \tilde{V}_t^{-1} \theta \end{aligned}$$

For the first term in the above,

$$\begin{aligned} \|\tilde{V}_t^{-1} \sum_{s=1}^{t-1} (X_s + c_s^*) c_s^{*T} \theta\|_{\tilde{V}_t} &= \left\| \sum_{s=1}^{t-1} \tilde{X}_s c_s^{*T} \theta \right\|_{\tilde{V}_t^{-1}} \\ &\leq \sum_{s=1}^{t-1} \sqrt{(c_s^{*T} \theta)^2 \tilde{X}_s^T \tilde{V}_t^{-1} \tilde{X}_s} \leq \sum_{s=1}^{t-1} \|c_s^*\| \|\theta\| \sqrt{\tilde{X}_s^T \tilde{V}_t^{-1} \tilde{X}_s} \\ &\leq \sqrt{\sum_{s=1}^{t-1} \|c_s^*\|^2} \sqrt{\sum_{s=1}^{t-1} \tilde{X}_s^T \tilde{V}_t^{-1} \tilde{X}_s} \leq C \gamma_t. \end{aligned}$$

For the second term, $\|\tilde{V}_t^{-1} \sum_{s=1}^{t-1} \tilde{X}_s \epsilon_s\|_{\tilde{V}_t} = \|\sum_{s=1}^{t-1} \tilde{X}_s \epsilon_s\|_{\tilde{V}_t^{-1}}$. Using Theorem 1 in Abbasi-Yadkori et al. (2011), we have with probability at least $1 - \delta$,

$$\|\tilde{V}_t^{-1} \sum_{s=1}^{t-1} \tilde{X}_s \epsilon_s\|_{\tilde{V}_t} \leq \sqrt{2\sigma^2 \log \left(\frac{\det(\tilde{V}_t)^{1/2} \det(\lambda I_d)^{-1/2}}{\delta} \right)}.$$

Since $\text{tr}(\tilde{V}_t) = \lambda d + \sum_{s=1}^{t-1} \|\tilde{X}_s\|^2 \leq d\lambda + t$ and $\det(\tilde{V}_t) \leq (\frac{1}{d} \text{tr}(\tilde{V}_t))^d$, we have $\|\tilde{V}_t^{-1} \sum_{s=1}^{t-1} \tilde{X}_s \epsilon_s\|_{\tilde{V}_t} \leq \sigma \sqrt{d \log \left(\frac{1+t/\lambda}{\delta} \right)}$.

For the third term, $\|\lambda \tilde{V}_t^{-1} \theta\|_{\tilde{V}_t} \leq \sqrt{\lambda}$. So $\|\hat{\theta}_t - \theta\|_{\tilde{V}_t} \leq \beta_t + \gamma_t C$.

In conclusion, no matter it is under the attack of rewards or context, $\|\hat{\theta}_t - \theta\|_{\tilde{V}_t} \leq \beta_t + \gamma_t C$. The lemma follows by noticing that for all $x \in \mathbb{R}^d$,

$$|x^T (\hat{\theta}_t - \theta)| \leq \|\hat{\theta}_t - \theta\|_{\tilde{V}_t} \|x\|_{\tilde{V}_t^{-1}}.$$

□

8.2 Proof of Theorem 1

Proof. Denote $\tilde{r}_t = (\tilde{x}_{t,*} - \tilde{X}_t)^T \theta$, we first bound \tilde{r}_t and then relate this to the true single-round regret $r_t = (x_{t,*} - X_t)^T \theta$. According to Lemma 1,

$$\begin{aligned} \tilde{r}_t &= (\tilde{x}_{t,*} - \tilde{X}_t)^T \theta \leq (\tilde{x}_{t,*} - \tilde{X}_t)^T \hat{\theta}_t + (\beta_t + \gamma_t C) \left(\|\tilde{x}_{t,*}\|_{\tilde{V}_t^{-1}} + \|\tilde{X}_t\|_{\tilde{V}_t^{-1}} \right) \\ &\leq (\beta_t + \gamma_t C') \left(\|\tilde{X}_t\|_{\tilde{V}_t^{-1}} - \|\tilde{x}_{t,*}\|_{\tilde{V}_t^{-1}} \right) + (\beta_t + \gamma_t C) \left(\|\tilde{x}_{t,*}\|_{\tilde{V}_t^{-1}} + \|\tilde{X}_t\|_{\tilde{V}_t^{-1}} \right) \\ &\leq [2\beta_t + \gamma_t(C' + C)] \|\tilde{X}_t\|_{\tilde{V}_t^{-1}} + \gamma_t(C - C') \|\tilde{x}_{t,*}\|_{\tilde{V}_t^{-1}} \\ &\leq (2\beta_t + \gamma_t(C + C')) \|\tilde{X}_t\|_{\tilde{V}_t^{-1}} \quad \text{since } C \leq C' \end{aligned}$$

So

$$\begin{aligned} \tilde{R}(T) &:= \sum_{t=1}^T \tilde{r}_t \leq (2\beta_T + \gamma_T(C + C')) \sum_{t=1}^T \|\tilde{X}_t\|_{\tilde{V}_t^{-1}} \\ &\leq (2\beta_T + \gamma_T(C + C')) \gamma_{T+1} \sqrt{T}. \end{aligned}$$

Under attack of rewards, $\tilde{r}_t = r_t$ since $\tilde{x}_{t,a} = x_{t,a}$, so $R(T) = \tilde{R}(T)$. While under attack of context,

$$\begin{aligned} r_t - \tilde{r}_t &= [(x_{t,*} - \tilde{x}_{t,*} - (X_t - \tilde{X}_t))^T \theta] \\ &\leq (\|x_{t,*} - \tilde{x}_{t,*}\| + \|X_t - \tilde{X}_t\|) \|\theta\| \leq |c_t| \end{aligned}$$

So $R(T) = \sum_{t=1}^T r_t \leq \tilde{R}(T) + C$ under attack of context. In conclusion, with probability at least $1 - \delta$,

$$R(T) \leq (2\beta_T + \gamma_T(C + C')) \gamma_{T+1} \sqrt{T} \quad (\text{attack rewards}) \quad (12)$$

$$R(T) \leq (2\beta_T + \gamma_T(C + C')) \gamma_{T+1} \sqrt{T} + C \quad (\text{attack context}) \quad (13)$$

From Lemma 11 of Abbasi-Yadkori et al. (2011), we have that when $\lambda \geq \max(\max \|\tilde{X}_t\|^2, 1)$,

$$\gamma_t^2 \leq 2 \log \frac{\det(\tilde{V}_t)}{\det(\lambda I_d)} \leq 2d \log \frac{\frac{1}{d} \text{tr}(\tilde{V}_t)}{\lambda} \leq 2d \log \left(1 + \frac{t-1}{d\lambda} \right). \quad (14)$$

Therefore, $\beta_T, \gamma_T \leq O(\sqrt{d \log T})$, which ends the proof. \square

8.3 Proof of Lemma 3

Proof. From Corollary 3.2 in Auer et al. (2002), the regret of EXP3 algorithm is upper bounded by $2\sqrt{e-1}Q\sqrt{LK' \log K'}$, where Q is the maximum absolute sum of rewards in any epoch, L is the number of rounds and K' is the number of arms. In our case, $Q \leq H$, $L = \lceil \frac{T}{H} \rceil$ and $K' = |J|$. So we have

$$\text{Quantity (B)} \leq 2\sqrt{e-1}H \sqrt{\lceil \frac{T}{H} \rceil |J| \log |J|} \leq \tilde{O}(\sqrt{TH}).$$

\square

8.4 LinUCB algorithm for solving non-corrupted linear contextual bandit problem

Algorithm 2 LinUCB Algorithm

Input: T , λ , sub-Gaussian parameter σ , error probability δ .

- 1: Initialize $V_t = \lambda I_d$, $\hat{\theta}_t = 0$ and $b_t = 0$.
 - 2: **for** $t = 1$ **to** T **do**
 - 3: Observe context $x_{t,a}$.
 - 4: Calculate exploration parameter $\beta_t = \sigma \sqrt{d \log \left(\frac{1+t/\lambda}{\delta} \right)} + \sqrt{\lambda}$.
 - 5: Pull arm $a_t = \operatorname{argmax}_a x_{t,a}^T \hat{\theta}_t + \beta_t \|x_{t,a}\|_{V_t^{-1}}$.
 - 6: Observe reward Y_t .
 - 7: Update LinUCB components $V_{t+1} = V_t + X_t X_t^T$, $b_{t+1} = b_t + X_t Y_t$ and $\hat{\theta}_{t+1} = V_{t+1}^{-1} b_{t+1}$
 - 8: **end for**
-

8.5 BOB-No-Restart Algorithm

Algorithm 3 BOB-No-Restart Algorithm

Input: T , β_t as in Equation 1.

- 1: Initialize $\hat{\theta}_t = 0$, $\tilde{V}_t = \lambda I_d$ and candidate attack budget set $J = \{0\} \cup \{2^j\}_{j=0}^{\lceil \log_2 2KT \rceil}$.
- 2: Initialize exponential weights $w_j(1) = 1$ for $j = 1, \dots, |J|$.
- 3: Initialize the exploration parameter for EXP3 as $\alpha = \min \left\{ 1, \sqrt{\frac{|J| \log |J|}{(e-1)T}} \right\}$.
- 4: **for** $t = 1$ **to** T **do**
- 5: Observe (attacked) contextual feature vectors $\tilde{x}_{t,a}$.
- 6: Update probability distribution for pulling candidates in J as

$$p_j(t) = \frac{\alpha}{|J|} + (1 - \alpha) \frac{w_j(t)}{\sum_{i=1}^{|J|} w_i(t)}, \forall j = 1, \dots, |J|.$$

- 7: Draw $j_t \leftarrow j \in [|J|]$ with probability $p_j(t)$.
- 8: Pull arm according to the following equation

$$a_t = \operatorname{argmax}_{a=1, \dots, K} \tilde{x}_{t,a}^T \hat{\theta}_t + (\beta_t + \gamma_t J_{j_t}) \|\tilde{x}_{t,a}\|_{\tilde{V}_t^{-1}}$$

- 9: Observe reward \tilde{Y}_t .
- 10: Update LinUCB components, i.e., $\tilde{V}_{t+1} = \tilde{V}_t + \tilde{X}_t \tilde{X}_t^T$ and $\hat{\theta}_{t+1} = \tilde{V}_{t+1}^{-1} \sum_{s=1}^t \tilde{X}_s \tilde{Y}_s$.
- 11: Update EXP3 components: $\hat{y}_t(j) \leftarrow 0$ for all $j \neq j_t$, $\hat{y}_t(j) = \tilde{Y}_t / p_j(t)$ if $j = j_t$ and

$$w_j(t+1) = w_j(t) \times \exp \left(\frac{\alpha}{|J|} \hat{y}_t(j) \right).$$

- 12: **end for**
-