A Simple Approach for Non-stationary Linear Bandits

Peng Zhao, Lijun Zhang, Yuan Jiang, Zhi-Hua Zhou

National Key Laboratory for Novel Software Technology
Nanjing University, Nanjing 210023, China

Abstract

This paper investigates the problem of non-stationary linear bandits, where the unknown regression parameter is evolving over time. Existing studies develop various algorithms and show that they enjoy an $\widetilde{\mathcal{O}}(T^{2/3}P_T^{1/3})$ dynamic regret, where T is the time horizon and P_T is the path-length that measures the fluctuation of the evolving unknown parameter. In this paper, we discover that a serious technical flaw makes their results ungrounded, and then present a fix, which gives an $\widetilde{\mathcal{O}}(T^{3/4}P_T^{1/4})$ dynamic regret without modifying original algorithms. Furthermore, we demonstrate that instead of using sophisticated mechanisms, such as sliding window or weighted penalty, a simple restarted strategy is sufficient to attain the same regret guarantee. Specifically, we design an UCB-type algorithm to balance exploitation and exploration, and restart it periodically to handle the drift of unknown parameters. Our approach enjoys an $\widetilde{\mathcal{O}}(T^{3/4}P_T^{1/4})$ dynamic regret. Note that to achieve this bound, the algorithm requires an oracle knowledge of the path-length P_T. Combining the bandits-over-bandits mechanism by treating our algorithm as the base learner, we can further achieve the same regret bound in a parameter-free way. Empirical studies also validate the effectiveness of our approach.

1 Introduction

Multi-Armed Bandits (MAB) [Robbins, 1952] models the sequential decision-making with partial information, where the player requires to choose one of the K slot machines at each iteration in order to maximize the cumulative reward. MAB is a paradigmatic instance of the exploration versus exploitation trade-offs, which is fundamental in many areas of artificial intelligence, such as reinforcement learning [Sutton and Barto, 2018] and evolutionary algorithms [Crepinsek et al., 2013].

In many real-world decision-making problems, each arm is usually associated with certain side information. Therefore, researchers start to formulate structured bandits in which the reward distributions of each arm are connected by a common but unknown parameter. Particularly, Stochastic Linear Bandits (SLB) has received much attention [Auer, 2002, Dani et al., 2007, Chu et al., 2011, Abbasi-Yadkori et al., 2011, Li et al., 2019]. In SLB, at iteration t, the player makes a decision X_t from a feasible set $\mathcal{X} \subseteq \mathbb{R}^d$, and then observes the reward r_t satisfying

$$\mathbb{E}[r_t|X_t] = X_t^\top \theta_*, \tag{1}$$

where θ_* is an unknown regression parameter. The goal of the player is to minimize the (pseudo) regret,

$$\text{Regret}_T = T \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \theta_* - \sum_{t=1}^{T} X_t^\top \theta_*. \tag{2}$$

The stochastic linear bandits problem is well-studied in literatures. By exploiting the tool of upper confidence bounds, various approaches demonstrate an $\widetilde{\mathcal{O}}(d\sqrt{T})$ regret [Dani et al., 2007, Abbasi-Yadkori et al., 2011],[1] which matches the $\Omega(d\sqrt{T})$ lower bound established by Dani et al. [2007], up to $\log T$ factors.

However, the observation model (1) assumes that the regression parameter θ_* is fixed, which is unfortunately hard to satisfy in real-life applications, because data are usually collected in non-stationary environments. For instance, in recommender systems the regression parameter models customers' interests, which could vary over time when customers look through product pages. Therefore, it is crucial to facilitate stochastic linear bandits with capability of handling non-stationarity.

To address above issue, Cheung et al. [2019a] proposed the *non-stationary* linear bandits model, which assumes the reward r_t satisfies

$$\mathbb{E}[r_t | X_t] = X_t^\top \theta_t,$$

where θ_t is the unknown regression parameter at iteration t. Different from the standard SLB setting in (1), non-stationary linear bandits allow the unknown parameter to vary over time, whose evolution is often called the path-length defined as $P_T = \sum_{t=2}^{T} \|\theta_{t-1} - \theta_t\|_2$, which naturally measures the non-stationarity of environments. The player's goal is to minimize the following (pseudo) *dynamic regret*,

$$\text{D-Regret}_T = \sum_{t=1}^{T} \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \theta_t - \sum_{t=1}^{T} X_t^\top \theta_t, \tag{3}$$

namely, the cumulative regret against the optimal strategy that has full information of unknown parameters.

Recently, Cheung et al. [2019a] first established an $\Omega(d^{2/3} T^{2/3} P_T^{1/3})$ minimax lower bound for the non-stationary linear bandits problem. On the upper bound side, Cheung et al. [2019a] proposed an online UCB-type algorithm called WindowUCB, which is based on the sliding window least square estimator to track the evolving parameters. Subsequently, Russac et al. [2019] developed the WeightUCB algorithm, which adopted the weighted least square estimator for parameter estimation. The authors prove an $\widetilde{\mathcal{O}}(d^{2/3} T^{2/3} P_T^{1/3})$ dynamic regret guarantee for their algorithms, matching the aforementioned lower bound up to $\log T$ factors. However, we exhibit that a serious technical flaw makes their arguments and regret guarantees ungrounded. We revisit the analysis and demonstrate that it is actually impossible to upper bound the crucial term in their argument by the desired quantity as they expected. Further, we present a fix. Without modifying original algorithms, we prove that their algorithms attain an $\widetilde{\mathcal{O}}(d^{7/8} T^{3/4} P_T^{1/4})$ dynamic regret. This is the first contribution of this paper.

---

[1]We adopt the notation of $\widetilde{\mathcal{O}}(\cdot)$ to suppress logarithmic factors in the time horizon T.

Furthermore, although these two strategies [Cheung et al., 2019a, Russac et al., 2019] attain nice dynamic regret guarantees (after fixing the technical flaws), their algorithms and analyses are fairly complicated. Instead, we discover that a quite simple algorithm based on the *restarted strategy* (simply running an UCB-style algorithm and restarting it periodically), surprisingly, achieves the same dynamic regret guarantee and is more efficient. This is the second contribution of this paper. Indeed, our proposed algorithm enjoys the following three advantages compared with previous studies.

- The proposed algorithm for non-stationary linear bandits is very simple and thus easy to analyze its dynamic regret, only exploiting the standard self-normalized concentration inequality for classical stochastic linear bandits.

- Compared with WindowUCB, the sliding window least square based approach [Cheung et al., 2019a], our approach supports online update and enjoys a one-pass manner *without* storing historical data. Meanwhile, WindowUCB demands an $\mathcal{O}(w)$ memory where w is the window length; by contrast, our approach only requires a *constant* memory.

- Compared with WeightUCB, the weighted least square based approach [Russac et al., 2019], our approach and analysis are much simpler, without involving other complicated deviation results. Additionally, WeightUCB maintains and manipulates the covariance matrix and its variant, and thus takes a longer running time.

Overall, our approach is more friendly to the resource-constrained learning scenarios due to its simplicity. In the following, we start with a brief review of related work in Section 2. Then, Section 3 presents our proposed approach and the theoretical results. Section 4 provides the analysis. We further supply the empirical studies in Section 5 and finally conclude the paper in Section 6. Appendix A supplements technical lemmas.

2 Related Work

Online learning in non-stationary environments has drawn considerable attention recently, in both full-information and bandits settings. We focus on related work in the bandits setting.

Non-stationary multi-armed bandits problem with abrupt changes was first studied by Auer [2002]. Denoted by K the number of arms and by L the number of distribution changes, Auer [2002] proposed EXP3.S, a variant of EXP3, which achieves an $\widetilde{\mathcal{O}}(\sqrt{KLT})$ regret bound when L is known. The rate is minimax optimal up to $\log T$ factors. Later studies demonstrated that $\widetilde{\mathcal{O}}(\sqrt{KLT})$ regret is attainable by sliding window and weighted penalty strategies [Garivier and Moulines, 2011], as well as the restarted strategy [Allesiardo et al., 2017]. All these algorithms require the number of changes L as the input parameter, which is undesired in practice. Recently, Auer et al. [2019] achieved a near-optimal rate $\widetilde{\mathcal{O}}(\sqrt{KLT})$ without knowing prior knowledge of L. On the other hand, Besbes et al. [2019] studied the non-stationary MAB with slowly changing distributions, and proved an $\widetilde{\mathcal{O}}((K \log K)^{1/3} V_T^{1/3} T^{2/3})$ dynamic regret, where $V_T = \sum_{t=2}^{T} \|\boldsymbol{\mu}_t - \boldsymbol{\mu}_{t-1}\|_{\infty}$ is the total variation of changes in reward distributions.

Non-stationary linear bandits problem was first studied by Cheung et al. [2019a]. The authors established an $\Omega(d^{2/3}T^{2/3}P_T^{1/3})$ minimax lower bound, and then proposed the WindowUCB algorithm based on the sliding window least square, achieving an $\widetilde{\mathcal{O}}(d^{7/8}T^{3/4}P_T^{1/4})$ dynamic regret (after fixing the technical gap). Nevertheless, to implement the sliding window least square, WindowUCB needs to store historical data in a buffer. A natural replacement is the weighted least square, which supports online update and enjoys both nice empirical performance and sound theoretical guarantee [Guo et al., 1993, Zhao et al., 2021]. Based on the idea, Russac et al. [2019] proposed the WeightUCB algorithm and proved that the approach attains the same dynamic regret. Nevertheless, both algorithmic design and regret analysis of WeightUCB are fairly complicated. Besides, WeightUCB needs to maintain and manipulate covariance matrix and its variant (in the same scale), which leads to an evidently longer running time. Finally, both WindowUCB and WeightUCB require the unknown quantity P_T as an input. To avoid the limitation, Cheung et al. [2019a] developed the bandits-over-bandits mechanism as a meta algorithm and finally obtained an $\widetilde{\mathcal{O}}(d^{7/8}T^{3/4}P_T^{1/4})$ parameter-free dynamic regret guarantee.

In this work, we first revisit the analysis of two existing algorithms designed for non-stationary linear bandits in the literature [Cheung et al., 2019a, Russac et al., 2019]. We demonstrate that there exists a technical flaw in the analysis, making the claimed dynamic regret guarantee ungrounded. We present a new analysis to fix the technical gap. Next, we propose a simple algorithm based on the restarted strategy for non-stationary linear bandits and show that the simple algorithm can achieve the same dynamic regret guarantee as existing methods. We note that using the restarted strategy for non-stationary environments is not new, which has been applied in various scenarios, including non-stationary online convex optimization [Besbes et al., 2015], MAB with abrupt changes [Allesiardo et al., 2017], and MAB with gradual changes [Besbes et al., 2019]. However, to the best of our knowledge, our work is the first time to apply the restarted strategy to non-stationary linear bandits.

3 Our Results

We first introduce the formal problem setup and then present our approach.

3.1 Problem Setup

Setting. In non-stationary (infinite-armed) linear bandits, at each iteration t, let $X_t \in \mathcal{X} \subseteq \mathbb{R}^d$ denote the contextual information of the chosen arm and r_t denote its associated reward, and the model is assumed to be linearly parameterized, i.e.,

$$r_t = X_t^\top \theta_t + \eta_t, \tag{4}$$

where $\theta_t \in \mathbb{R}^d$ is the unknown parameter and η_t is the noise satisfying certain tail condition specified below. As mentioned earlier, to guide the algorithmic design of non-stationary linear bandits, it is natural to employ the following (pseudo) *dynamic regret*

as the performance measure:

$$\text{D-Regret}_T = \sum_{t=1}^{T} \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \theta_t - \sum_{t=1}^{T} X_t^\top \theta_t,$$

which is the cumulative regret against the optimal strategy that has full information of unknown parameters.

Assumptions. We assume the noise η_t be conditionally R-sub-Gaussian with a fixed constant $R > 0$. That is, $\mathbb{E}[\eta_t \mid X_{1:t}, \eta_{1:t-1}] = 0$, and for any $\lambda \in \mathbb{R}$,

$$\mathbb{E}[\exp(\lambda \eta_t) \mid X_{1:t}, \eta_{1:t-1}] \leq \exp\left(\frac{\lambda^2 R^2}{2}\right),$$

The feasible set and unknown parameters are assumed to be bounded, i.e., $\forall \mathbf{x} \in \mathcal{X}$, $\|\mathbf{x}\|_2 \leq L$, and $\|\theta_t\|_2 \leq S$ holds for all $t \in [T]$. For convenience, we further assume $\langle \mathbf{x}, \theta_t \rangle \leq 1$, but we will keep the dependence in L and S for better comprehension of the results.

3.2 RestartUCB Algorithm

RestartUCB algorithm has two key ingredients: upper confidence bounds for trading off the exploration and exploitation, and the restarted strategy for handling the non-stationarity of environments. Specifically, our proposed RestartUCB algorithm proceeds in epochs. At each iteration, we first estimate the unknown regression parameter from historical data within the epoch, and then construct upper confidence bounds of the expected reward for selecting the arm. Finally, we periodically restart the algorithm to be resilient to the drift of underlying parameter θ_t.

In the following, we first specify the estimator used in the RestartUCB algorithm, then investigate its estimate error to construct upper confidence bounds, and finally describe the restarted strategy.

Estimator. At iteration t, we adopt the regularized least square estimator by only exploiting data in the current epoch. More precisely, the estimator $\widehat{\theta}_t$ is the solution of the following problem:

$$\min_{\theta} \ \lambda \|\theta\|_2^2 + \sum_{s=t_0}^{t-1} (X_s^\top \theta - r_s)^2, \tag{5}$$

where t_0 is the starting point of the current epoch, and $\lambda > 0$ is the regularization coefficient. Clearly, $\widehat{\theta}_t$ admits a closed-form solution as

$$\widehat{\theta}_t = V_{t-1}^{-1} \left(\sum_{s=t_0}^{t-1} r_s X_s \right), \tag{6}$$

where $V_{t-1} = \lambda I + \sum_{s=t_0}^{t-1} X_s X_s^\top$. We remark that the estimator (6) (essentially, both the terms of V_{t-1} and $\sum_{s=t_0}^{t-1} r_s X_s$) can be updated online *without* storing historical data in the memory owing to the restarted strategy. Furthermore, it is known that (5) can be *exactly* solved by the recursive least square algorithm, whose solution is provably

equivalent to the closed-form expression (6). This feature can further accelerate our approach in that it saves the computation of the inverse of covariance matrix V_{t-1}, which is arguably the most time-consuming step at each iteration.

By contrast, Cheung et al. [2019a] adopted the following sliding window least square estimator,

$$\widehat{\theta}_t^{\mathrm{sw}} = (V_{t-1}^{\mathrm{sw}})^{-1} \left(\sum_{s=1 \vee (t-w)}^{t-1} r_s X_s \right), \tag{7}$$

where $V_{t-1}^{\mathrm{sw}} = \lambda I + \sum_{s=1 \vee (t-w)}^{t-1} X_s X_s^{\top}$ is the covariance matrix formed by historical data in the sliding window and $w > 0$ is the window length. For online update, WindowUCB will remove the oldest data item in the window and then add the new item. So it requires to store the nearest w data items in the memory for future update, resulting in an $\mathcal{O}(w)$ space complexity which cannot be regarded as a constant because the setting of w depends on the time horizon T.

Upper Confidence Bounds. Based on the estimator $\widehat{\theta}_t$ in (6), we further construct upper confidence bounds for the expected reward. To this end, it is required to investigate the estimate error. Inspired by the analysis of WindowUCB [Cheung et al., 2019a], we have the following result.

Lemma 1. *For any $t \in [T]$ and $\delta \in (0, 1)$, with probability at least $1 - \delta$, the following holds for all $\mathbf{x} \in \mathcal{X}$,*

$$|\mathbf{x}^{\top}(\theta_t - \widehat{\theta}_t)| \leq L^2 \sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 + \beta_t \|\mathbf{x}\|_{V_{t-1}^{-1}}, \tag{8}$$

where $H > 0$ is the restarting period (or epoch size), and β_t is the radius of confidence region,

$$\beta_t = \sqrt{\lambda} S + R \sqrt{2 \log \frac{1}{\delta} + d \log \left(1 + \frac{(t - t_0) L^2}{\lambda d} \right)}. \tag{9}$$

Remark 1. The analysis of estimate error serves as the foundation of designing the UCB-type algorithms. In fact, the pioneering study [Cheung et al., 2019a] has studied the estimate error of sliding window least square estimator for non-stationary linear bandits, however, the technical reasoning suffers from some gaps and makes the estimate error bound and the claimed $\widetilde{\mathcal{O}}(T^{2/3} P_T^{1/3})$ dynamic regret guarantee ungrounded. The flaw appears in a key technical lemma [Cheung et al., 2019b, Lemma 3], and is unfortunately inherited by the later studies including WeightUCB [Russac et al., 2019, Theorem 2], the early version of this paper [Zhao et al., 2020b, Lemma 3], and perturbation based method [Kim and Tewari, 2020, Theorem 7]. In this version, we correct the previous results, at the price of another $\sqrt{dH}$ factor appearing in front of the path-length term comparing to the original result (see Lemma 1 in the early version of our work [Zhao et al., 2020b]). The additional $\sqrt{dH}$ factor in the estimation error will lead to an $\widetilde{\mathcal{O}}(T^{3/4} P_T^{1/4})$ dynamic regret, which is slightly worse than the original $\widetilde{\mathcal{O}}(T^{2/3} P_T^{1/3})$ rate. We present more technical discussions in Section 4.3.

Algorithm 1 RestartUCB

Input: time horizon T, restarting period H, confidence δ, regularizer λ, scaling parameters S and L

1: Set epoch counter $j = 1$
2: **while** $j \leq \lceil T/H \rceil$ **do**
3:      Set $\tau = (j-1)H$
4:      Initialize $X_\tau \in \mathcal{X}$
5:      $V_\tau = \lambda I_d$
6:      **for** $t = \tau + 1, \ldots, \tau + H - 1$ **do**
7:          Compute $\widehat{\theta}_t = V_{t-1}^{-1} S_{t-1}$
8:          Compute β_t by (9) with $t_0 = \tau$
9:          Select $X_t = \arg\max_{\mathbf{x} \in \mathcal{X}} \{ \langle \mathbf{x}, \widehat{\theta}_t \rangle + \beta_t \|\mathbf{x}\|_{V_{t-1}^{-1}} \}$
10:          Receive the reward r_t
11:          Update $V_t = V_{t-1} + X_t X_t^\top$ and $S_t = S_{t-1} + r_t X_t$
12:      **end for**
13:      Set $j = j + 1$
14: **end while**

The estimate error (8) essentially suggests an upper confidence bound of the expected reward $\mathbf{x}^\top \theta_t$. Hence, we adopt the principle of *optimism in the face of uncertainty* [Auer, 2002] and choose the arm that maximizes its upper confidence bound,

$$
\begin{aligned}
X_t &= \arg\max_{\mathbf{x} \in \mathcal{X}} \left\{ \langle \mathbf{x}, \widehat{\theta}_t \rangle + L^2 \sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 + \beta_t \|\mathbf{x}\|_{V_t^{-1}} \right\} \\
&= \arg\max_{\mathbf{x} \in \mathcal{X}} \left\{ \langle \mathbf{x}, \widehat{\theta}_t \rangle + \beta_t \|\mathbf{x}\|_{V_t^{-1}} \right\}.
\end{aligned}
\tag{10}
$$

So at iteration t, the algorithm first solves the estimator based on (6), then obtains the confidence radius β_t by (9), and finally pulls the arm X_t according to the selection criteria (10).

Restarted Strategy. To handle the changes of unknown regression parameters, RestartUCB algorithm proceeds in epochs and restarts the procedure every H iterations, as illustrated in Figure 1. We call the variable H as the restarting period or epoch size, which is the key parameter to trade off between the stability and non-stationarity. In each epoch, RestartUCB performs the UCB-style algorithm as described in the last part. We summarize overall procedures in Algorithm 1.
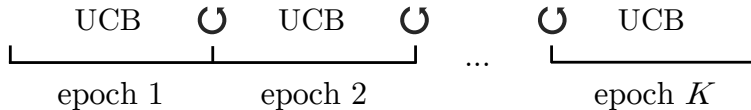

Figure 1: Illustration of RestartUCB algorithm.

We will show in the next subsection that a fixed length is sufficient to achieve the same theoretical guarantees as previous works [Cheung et al., 2019a, Russac et al.,

2019]. Nevertheless, since the regret guarantee is not optimal, it would be interesting to see whether an adaptive epoch length with a certain statistical detection will give an improved regret guarantee.

3.3 Theoretical Guarantees

We show that notwithstanding its simplicity RestartUCB algorithm enjoys the same dynamic regret guarantee as the existing methods for non-stationary linear bandits, including WindowUCB [Cheung et al., 2019a] and WeightUCB [Russac et al., 2019].

In the following, we first analyze the regret within each epoch (Theorem 1), and then sum over epochs to obtain the guarantee of the whole time horizon (Theorem 2).

Theorem 1. *For each epoch $\mathcal{E}$ whose size is H and any $\delta \in (0,1)$, with probability at least $1 - 2\delta$, the dynamic regret within the epoch is upper bounded by*

$$\text{D-Regret}(\mathcal{E}) \triangleq \sum_{t \in \mathcal{E}} \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \theta_t - \sum_{t \in \mathcal{E}} X_t^\top \theta_t$$

$$\leq 2L^2 \sqrt{\frac{d}{\lambda}} \cdot H^{\frac{3}{2}} \mathcal{P}(\mathcal{E}) + 2\beta_H \sqrt{2dH \log\left(1 + \frac{L^2 H}{\lambda d}\right)},$$

where $\beta_H = \sqrt{\lambda} S + R \sqrt{2 \log(1/\delta) + d \log\left(1 + \frac{HL^2}{\lambda d}\right)}$ is the confidence radius of the epoch, and $\mathcal{P}(\mathcal{E})$ denotes the path-length within epoch $\mathcal{E}$, i.e., $\mathcal{P}(\mathcal{E}) = \sum_{t \in \mathcal{E}} \|\theta_{t-1} - \theta_t\|_2$.

By summing the dynamic regret over epochs, we can therefore obtain dynamic regret over of the whole time horizon.

Theorem 2. *With probability at least $1 - 1/T$, the dynamic regret of RestartUCB (Algorithm 1) over the whole time horizon is upper bounded by*

$$\text{D-Regret}_T = \sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \theta_t - \sum_{t=1}^T X_t^\top \theta_t \leq \widetilde{\mathcal{O}}\left(d^{\frac{1}{2}} H^{\frac{3}{2}} P_T + dT/\sqrt{H}\right), \qquad (11)$$

where $P_T = \sum_{t=2}^T \|\theta_{t-1} - \theta_t\|_2$ is the path-length, and H is the restarting period.
Furthermore, by setting the restarting period optimally as

$$H = \min\{H^*, T\} = \min\left\{\lfloor d^{\frac{1}{4}} T^{\frac{1}{2}} P_T^{-\frac{1}{2}} \rfloor, T\right\}, \qquad (12)$$

RestartUCB achieves the following dynamic regret,

$$\text{D-Regret}_T \leq \begin{cases} \widetilde{\mathcal{O}}\left(d^{\frac{7}{8}} T^{\frac{3}{4}} P_T^{\frac{1}{4}}\right) & \text{when } P_T \geq \sqrt{d}/T, \\ \widetilde{\mathcal{O}}(d\sqrt{T}) & \text{when } P_T < \sqrt{d}/T. \end{cases} \qquad (13)$$

Remark 2. As shown in Theorem 2, the setting of optimal restarting period H^* in (12) requires the prior information of path-length P_T, which is generally unavailable. We will discuss how to remove the undesired dependence in the next subsection.
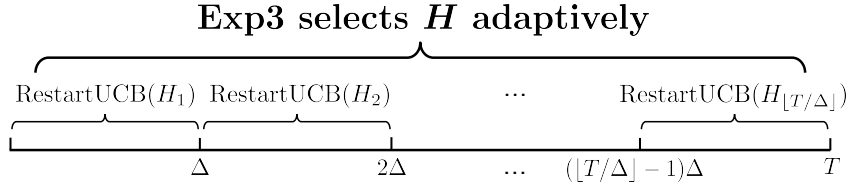
Figure 2: Illustration of Bandits-over-Bandits mechanism [Cheung et al., 2019b] with application to our proposed RestartUCB algorithm. The overall algorithm is in two-layer meta-base structure: the algorithm treats RestartUCB as the base-learner to handle non-stationary linear bandits with a given restarting period, and employs Exp3 as the meta-learner to adaptively choose the optimal restarting period H.

3.4 Adapting to Unknown Non-stationarity

As mentioned earlier, the restarting period H plays a key role in dealing with the non-stationarity of environments. Intuitively, a small restarting period should be employed when the environments change very dramatically, and a large one should be used when the environments are relatively stable. Our theoretical result also validates the intuition. As one can see in Theorem 2, the restarting period H trades off between the path-length term P_T and the total horizon T, which necessitates an appropriate balance. Nevertheless, the optimal configuration of restarting period, as shown in (12), requires the prior knowledge of path-length P_T, which essentially measures the non-stationarity of underlying environments and is thus generally unavailable.

To compensate the lack of this information of environmental non-stationarity, we design an online ensemble method [Zhou, 2012] by employing the meta-base framework, which is recently used in full-information non-stationarity online learning [Daniely et al., 2015, Jun et al., 2017, Zhang et al., 2018, Zheng et al., 2019, Zhang et al., 2020a,b, Zhao et al., 2020c] and non-stationary bandit online learning [Agarwal et al., 2017, Cheung et al., 2019a, Zhao et al., 2020a]. In this paper, to deal with the issue for non-stationary linear bandits, we employ the *Bandits-over-Bandits* (BOB) mechanism, proposed by Cheung et al. [2019a] in designing parameter-free algorithm for non-stationary linear bandits based on sliding window least square estimator.

In the following, we describe how to apply the BOB mechanism to eliminate the requirement of the unknown path-length in RestartUCB. The essential idea is to use the RestartUCB algorithm as the base-learner to handle non-stationary linear bandits with a given restarting period, and on top of that we will employ a second-layer bandits algorithm as the meta-learner to adaptively learn the optimal restarting period. We name the RestartUCB algorithm with Bandits-over-Bandits mechanism as "RestartUCB-BOB". Figure 2 illustrates the meta-base structure of RestartUCB-BOB algorithm. To be more concrete, although the exact value of the optimal restarting period H^* (or equivalently, the path-length P_T) is unknown, we can make some guess of its possible value, as the P_T is always bounded. Then, we can use a certain meta-algorithm to adaptively track the best restarting period. To achieve this goal, RestartUCB-BOB first requires to examine the performance of base-learner with different restarting period. Therefore, RestartUCB-BOB will perform in several episodes, and in each episode

RestartUCB-BOB employs the base-algorithm RestartUCB with a particular restarting period and receives the returned cumulative reward over the episode as the reward feedback. We denote by $\Delta \in [T]$ the episode length. The restarting period will be adaptively adjusted by employing EXP3 [Auer et al., 2002] as the meta-algorithm. In the configuration of RestartUCB-BOB, the episode length is set $\Delta = \lceil d\sqrt{T} \rceil$. Further, the pool of candidate restarting periods $\mathcal{H}$ is configured as follows:

$$\mathcal{H} = \left\{ H_i = \lfloor d^{\frac{1}{4}} S^{-\frac{1}{2}} \cdot 2^{i-1} \rfloor \mid i \in [N] \right\}, \tag{14}$$

where $N = \lceil \frac{1}{2} \log_2(ST) \rceil + 1$ is the number of candidate restarting periods and recall that S is the upper bound of the norm of underlying regression parameters as specified in Section 3.1. Let $H_{\min}$ $(H_{\max})$ be the minimal (maximal) restarting period in the pool $\mathcal{H}$, then it is evident to verify that

$$H_{\min} = \lfloor d^{\frac{1}{4}} S^{-\frac{1}{2}} \rfloor, H_{\max} = \lfloor d^{\frac{1}{4}} \sqrt{T} \rfloor \leq \Delta. \tag{15}$$

To conclude, RestartUCB-BOB is in a two-layer meta-base structure and will perform in episodes. In each episode, the base-learner is RestartUCB associated with a particular restarting period in the candidate pool determined by the meta-learner EXP3; besides, the cumulative reward of the base-learner within the episode is fed to the meta-learner as the feedback to adaptively choose a better restarting period. We refer the reader to Section 7 of Cheung et al. [2019b] for more descriptions of algorithmic details.

The following theorem presents the dynamic regret guarantee for RestartUCB-BOB. Note that the algorithm does not require the prior knowledge of the path-length P_T.

Theorem 3. RESTARTUCB *together with Bandits-over-Bandits mechanism satisfies*

$$\text{D-Regret}_T = \sum_{t=1}^{T} \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^{\top} \theta_t - \sum_{t=1}^{T} X_t^{\top} \theta_t \leq \widetilde{\mathcal{O}}\big(d^{\frac{7}{8}} T^{\frac{3}{4}} P_T^{\frac{1}{4}}\big), \tag{16}$$

without requiring the path-length P_T ahead of time.

Remark 3. From the theorem, we can observe that RestartUCB-BOB enjoys the same dynamic regret bound as RestartUCB with an oracle tuning (13) shown in Theorem 2, while RestartUCB-BOB now requires no prior knowledge on the environmental non-stationarity measure P_T. Nevertheless, the attained dynamic regret upper bound still exhibits a certain gap to the $\Omega(d^{2/3} T^{2/3} P_T^{1/3})$ minimax lower bound of non-stationary linear bandits [Cheung et al., 2019a]. Therefore, it remains open on how to obtain rate-optimal and parameter-free dynamic regret. Indeed, even with an oracle tuning, RestartUCB still cannot achieve optimal dynamic regret. We are not sure whether this is due to the limitation of the regret analysis or the algorithm itself. Finally, we note that recent studies achieve near-optimal dynamic regret without prior information for multi-armed bandits [Auer et al., 2018, 2019] and contextual bandits [Chen et al., 2019] by means of change detection. These studies could be useful in designing parameter-free algorithms for non-stationary linear bandits, which will be investigated in the future.

4 Analysis

In this section, we provide proofs of theoretical results presented in the previous section.

4.1 Proof of Lemma 1

Proof. From the model assumption (4) and the estimator (6), we can verify that the estimate error can be decomposed as,

$$\widehat{\theta}_t - \theta_t = V_{t-1}^{-1}\left(\sum_{s=t_0}^{t-1} X_s X_s^\top (\theta_s - \theta_t) + \sum_{s=t_0}^{t-1} \eta_s X_s - \lambda\theta_t \right).$$

Therefore, by Cauchy-Schwarz inequality, we know that for any $\mathbf{x} \in \mathcal{X}$,

$$|\mathbf{x}^\top(\widehat{\theta}_t - \theta_t)| \le \|\mathbf{x}\|_2 \cdot A_t + \|\mathbf{x}\|_{V_{t-1}^{-1}} \cdot B_t, \tag{17}$$

where

$$A_t = \left\| V_{t-1}^{-1}\left(\sum_{s=t_0}^{t-1} X_s X_s^\top (\theta_s - \theta_t) \right) \right\|_2, \text{ and } B_t = \left\| \sum_{s=t_0}^{t-1} \eta_s X_s - \lambda\theta_t \right\|_{V_{t-1}^{-1}}.$$

We will give upper bounds for these two terms separately, as summarized in the following two lemmas.

Lemma 2. *For any $t \in [T]$, we have*

$$\left\| V_{t-1}^{-1}\left(\sum_{s=t_0}^{t-1} X_s X_s^\top (\theta_s - \theta_t) \right) \right\|_2 \le L\sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2. \tag{18}$$

Lemma 3. *For any $t \in [T]$, we have*

$$\left\| \sum_{s=t_0}^{t-1} \eta_s X_s - \lambda\theta_t \right\|_{V_{t-1}^{-1}} \le \sqrt{\lambda}S + R\sqrt{2\log\frac{1}{\delta} + d\log\left(1 + \frac{(t-t_0)L^2}{\lambda d}\right)}, \tag{19}$$

where $\beta_t \triangleq \sqrt{\lambda}S + R\sqrt{2\log\frac{1}{\delta} + d\log\left(1 + \frac{(t-t_0)L^2}{\lambda d}\right)}$ is the confidence radius used in RestartUCB.

Based on the inequality (17), Lemma 2, Lemma 3, and the boundedness of the feasible set, we have for any $\mathbf{x} \in \mathcal{X}$,

$$|\mathbf{x}^\top(\theta_t - \widehat{\theta}_t)| \le L^2\sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 + \beta_t \|\mathbf{x}\|_{V_{t-1}^{-1}},$$

which competes the proof. $\qquad\square$

We proceed to prove Lemma 2 and Lemma 3. It is noteworthy mentioning that previous works in non-stationary linear bandits [Cheung et al., 2019b, Russac et al., 2019, Zhao et al., 2020b] also need to upper bound some quantities similar to A_t, while their results are general invalid due to a serious technical flaw that will be explicitly stated in Section 4.3. Lemma 2 serves as the key component to fix existing results.

Proof of Lemma 2. Notice that

$$\left\| V_{t-1}^{-1} \Big(\sum_{s=t_0}^{t-1} X_s X_s^\top (\theta_s - \theta_t) \Big) \right\|_2 = \left\| V_{t-1}^{-1} \Big(\sum_{s=t_0}^{t-1} X_s X_s^\top \Big(\sum_{p=s}^{t-1} (\theta_p - \theta_{p+1}) \Big) \Big) \right\|_2$$

$$= \left\| V_{t-1}^{-1} \Big(\sum_{p=t_0}^{t-1} \Big(\sum_{s=t_0}^{p} X_s X_s^\top (\theta_p - \theta_{p+1}) \Big) \Big) \right\|_2$$

$$\leq \sum_{p=t_0}^{t-1} \left\| V_{t-1}^{-1} \Big(\sum_{s=t_0}^{p} X_s X_s^\top \Big)(\theta_p - \theta_{p+1}) \right\|_2$$

$$\leq \sum_{p=t_0}^{t-1} \left\| V_{t-1}^{-1} \Big(\sum_{s=t_0}^{p} X_s X_s^\top \Big) \right\|_2 \|\theta_p - \theta_{p+1}\|_2.$$

We now derive the upper bound for the term $\|V_{t-1}^{-1}(\sum_{s=t_0}^{p} X_s X_s^\top)\|_2$. Denote by $\mathcal{S}(1) = \{\mathbf{x} \mid \|\mathbf{x}\|_2 = 1\}$ the unit sphere.

$$\left\| V_{t-1}^{-1} \Big(\sum_{s=t_0}^{p} X_s X_s^\top \Big) \right\|_2 = \sup_{\mathbf{z} \in \mathcal{S}(1)} \sup_{\widetilde{\mathbf{z}} \in \mathcal{S}(1)} \left| \mathbf{z}^\top V_{t-1}^{-1} \Big(\sum_{s=t_0}^{p} X_s X_s^\top \Big) \widetilde{\mathbf{z}} \right|$$

$$= \left| \mathbf{z}_*^\top V_{t-1}^{-1} \Big(\sum_{s=t_0}^{p} X_s X_s^\top \Big) \widetilde{\mathbf{z}}_* \right|$$

$$\leq \|\mathbf{z}_*\|_{V_{t-1}^{-1}} \left\| \sum_{s=t_0}^{p} X_s (X_s^\top \widetilde{\mathbf{z}}_*) \right\|_{V_{t-1}^{-1}}$$

$$\leq \|\mathbf{z}_*\|_{V_{t-1}^{-1}} \left\| \sum_{s=t_0}^{p} X_s \|X_s\|_2 \|\widetilde{\mathbf{z}}_*\|_2 \right\|_{V_{t-1}^{-1}}$$

$$\leq \frac{L}{\sqrt{\lambda}} \left\| \sum_{s=t_0}^{p} X_s \right\|_{V_{t-1}^{-1}}$$

$$\leq \frac{L}{\sqrt{\lambda}} \sum_{s=t_0}^{p} \|X_s\|_{V_{t-1}^{-1}}$$

$$\leq \frac{L}{\sqrt{\lambda}} \sqrt{H} \sqrt{\sum_{s=t_0}^{p} \|X_s\|_{V_{t-1}^{-1}}^2}$$

$$\leq L \sqrt{\frac{dH}{\lambda}}.$$

In above, the first equation makes use of the property of the matrix 2-norm: for a matrix $M \in \mathbb{R}^{m \times n}$, $\|M\|_2 = \sup_{\|\mathbf{x}\|_2=1} \sup_{\|\mathbf{y}\|_2=1} |\mathbf{x}^\top M \mathbf{y}|$, whose proof can be found from the book [Meyer, 2000, Chapter 5, Eq. (5.2.9)] and also Lemma 6 for self-containedness. Further, $(\mathbf{z}_*, \widetilde{\mathbf{z}}_*)$ denotes the optimizer of the right hand side optimization problem in the first line. In the proof, we use the fact that for any $\mathbf{x}$, we have $\|\mathbf{x}\|_{V_{t-1}^{-1}} \leq \|\mathbf{x}\|_2/\sqrt{\lambda}$

as $V_{t-1} \succeq \lambda I$. The second last step holds by the Cauchy-Schwarz inequality. Besides, the last step follows from the fact: for any $p \in \{t_0, \ldots, t-1\}$,

$$
\begin{aligned}
\sum_{s=t_0}^{p} \|X_s\|_{V_{t-1}^{-1}}^2 &= \sum_{s=t_0}^{p} \text{Tr}(X_s^\top V_{t-1}^{-1} X_s) \\
&= \text{Tr}\left(V_{t-1}^{-1} \sum_{s=t_0}^{p} X_s X_s^\top \right) \\
&\leq \text{Tr}\left(V_{t-1}^{-1} \sum_{s=t_0}^{p} X_s X_s^\top \right) + \sum_{s=p+1}^{t-1} X_s^\top V_{t-1}^{-1} X_s + \lambda \sum_{i=1}^{d} \mathbf{e}_i^\top V_{t-1}^{-1} \mathbf{e}_i \\
&= \text{Tr}\left(V_{t-1}^{-1} \sum_{s=t_0}^{p} X_s X_s^\top \right) + \text{Tr}\left(V_{t-1}^{-1} \sum_{s=p+1}^{t-1} X_s X_s^\top \right) + \text{Tr}\left(V_{t-1}^{-1} \lambda \sum_{i=1}^{d} \mathbf{e}_i \mathbf{e}_i^\top \right) \\
&= \text{Tr}(I_d) = d.
\end{aligned}
$$

Hence, we complete the proof. $\qquad\square$

Proof of Lemma 3. From the self-normalized concentration inequality [Abbasi-Yadkori et al., 2011, Theorem 1], restated in Theorem 5 of Section A, we know that

$$
\left\| \sum_{s=t_0}^{t-1} \eta_s X_s \right\|_{V_{t-1}^{-1}} \overset{(38)}{\leq} \sqrt{2R^2 \log\left(\frac{\det(V_{t-1})^{1/2} \det(\lambda I)^{-1/2}}{\delta} \right)}
$$

$$
\leq R\sqrt{2\log\frac{1}{\delta} + d\log\left(1 + \frac{(t-t_0)L^2}{d}\right)},
$$

where the last inequality is obtained from the analysis of the determinant, as shown in the proof of Lemma 4. Meanwhile, since $V_{t-1} \succeq \lambda I$, we know that

$$
\|\lambda\theta_t\|_{V_{t-1}^{-1}}^2 \leq 1/\lambda_{\min}(V_{t-1})\|\lambda\theta_t\|_2^2 \leq \frac{1}{\lambda}\|\lambda\theta_t\|_2^2 \leq \lambda S^2.
$$

Therefore, the upper bound of B_t can be immediately obtained by combining the above inequalities. $\qquad\square$

4.2 Proof of Theorems 1 and 2

Proof of Theorem 1. Due to Lemma 1 and the fact that $X_t^*, X_t \in \mathcal{X}$, each of the following holds with probability at least $1 - \delta$,

$$
\langle X_t^*, \theta_t \rangle \leq \langle X_t^*, \widehat{\theta}_t \rangle + L^2 \sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 + \beta_t \|X_t^*\|_{V_{t-1}^{-1}},
$$

$$
\langle X_t, \theta_t \rangle \geq \langle X_t, \widehat{\theta}_t \rangle - L^2 \sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 - \beta_t \|X_t\|_{V_{t-1}^{-1}}.
$$

13

By the union bound, the following holds with probability at least $1 - 2\delta$,

$$\langle X_t^*, \theta_t \rangle - \langle X_t, \theta_t \rangle$$

$$\leq \langle X_t^*, \widehat{\theta}_t \rangle - \langle X_t, \widehat{\theta}_t \rangle + 2L^2\sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 + \beta_t(\|X_t^*\|_{V_{t-1}^{-1}} + \|X_t\|_{V_{t-1}^{-1}})$$

$$\leq 2L^2\sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 + 2\beta_t\|X_t\|_{V_{t-1}^{-1}},$$

where the last step comes from the following implication of the arm selection criteria (10) such that $\langle X_t^*, \widehat{\theta}_t \rangle + \beta_t\|X_t^*\|_{V_{t-1}^{-1}} \leq \langle X_t, \widehat{\theta}_t \rangle + \beta_t\|X_t\|_{V_{t-1}^{-1}}$.

Hence, dynamic regret within epoch $\mathcal{E}$ is bounded by,

$$\text{D-Regret}(\mathcal{E}) \leq \sum_{t \in \mathcal{E}} 2L^2\sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2 + 2\beta_t\|X_t\|_{V_{t-1}^{-1}}$$

$$\leq 2L^2\sqrt{\frac{d}{\lambda}}H^{\frac{3}{2}}\mathcal{P}(\mathcal{E}) + 2\beta_H\sqrt{2dH\log\left(1 + \frac{L^2H}{\lambda d}\right)},$$

where the last inequality holds due to the standard elliptical potential lemma (Lemma 4), whose statement and proof are presented in Section A. □

Proof of Theorem 2. By taking the union bound over the dynamic regret of all $\lceil T/H \rceil$ epochs, we know that the following holds with probability at least $1 - 2/T$,

$$\text{D-Regret}_T = \sum_{s=1}^{\lceil T/H \rceil} \text{D-Regret}(\mathcal{E}_s) \leq 2L^2\sqrt{\frac{d}{\lambda}}H^{\frac{3}{2}}P_T + 2T\widetilde{\beta}_H\sqrt{\frac{2d}{H}\log\left(1 + \frac{L^2H}{\lambda d}\right)},$$

where $\widetilde{\beta}_H = \sqrt{\lambda}S + R\sqrt{2\log(T\lceil\frac{T}{H}\rceil) + d\log\left(1 + \frac{HL^2}{\lambda d}\right)}$. Ignoring logarithmic factors in time horizon T, we finally obtain that

$$\text{D-Regret}_T \leq \widetilde{\mathcal{O}}\left(d^{\frac{1}{2}}H^{\frac{3}{2}}P_T + dT/\sqrt{H}\right).$$

When $P_T < \sqrt{d}/T$ (which corresponds a small amount of non-stationarity), we simply set the restarting period as T and achieve an $\widetilde{\mathcal{O}}(d\sqrt{T})$ regret bound. Note that under such a configuration, our algorithm actually performs no restart and thereby recovers the standard LinUCB algorithm for the stationary stochastic linear bandits [Abbasi-Yadkori et al., 2011]. Besides, when coming to the non-degenerated case of $P_T \geq \sqrt{d}/T$, we set the restarting period optimally as $H = \lfloor d^{1/4}T^{1/2}P_T^{-1/2} \rfloor$ and attain an $\widetilde{\mathcal{O}}(d^{\frac{7}{8}}T^{\frac{3}{4}}P_T^{\frac{1}{4}})$ dynamic regret bound. This ends the proof. □

4.3 Revisiting Existing Results

Previous studies show an $\widetilde{\mathcal{O}}(T^{2/3}P_T^{1/3})$ dynamic regret for non-stationary linear bandits, however, the technical reasoning suffers from some gaps and makes the overall regret

guarantee ungrounded. In the following, we first spot the flaws of their original proofs and then discuss the key component of our new analysis.

Indeed, the flaw appears in a key technical lemma for regret analysis of WindowUCB, the pioneering study of non-stationary linear bandits [Cheung et al., 2019b, Lemma 3]. The flaw is unfortunately inherited by the later studies, including WeightUCB [Russac et al., 2019, Theorem 2], the early version of this paper [Zhao et al., 2020b, Lemma 3], and perturbation based method [Kim and Tewari, 2020, Theorem 7]. To be more concrete, Lemma 3 of Cheung et al. [2019b] (also see Lemma 3 of Zhao et al. [2020b]) claims that for any $t \in [T]$,

$$\left\| V_{t-1}^{-1} \left(\sum_{s=t_0}^{t-1} X_s X_s^\top (\theta_s - \theta_t) \right) \right\|_2 \leq \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2. \tag{20}$$

Actually, the quantity $\|V_{t-1}^{-1}(\sum_{s=t_0}^{t-1} X_s X_s^\top (\theta_s - \theta_t))\|_2$ is of great importance for the regret analysis of non-stationary linear bandits algorithms, because it will be finally converted to the path-length of unknown regression parameters. Our Lemma 2 gives an upper bound of $L\sqrt{\frac{dH}{\lambda}} \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2$, which has a worse dependence in terms of dimension d and restarting period H compared to (20). However, we will demonstrate that the proof of the above claim (20) suffers serious technical flaws, which makes the result ungrounded. We restate their proof [Cheung et al., 2019b, Appendix B] as follows:

$$\left\| V_{t-1}^{-1} \left(\sum_{s=t_0}^{t-1} X_s X_s^\top (\theta_s - \theta_t) \right) \right\|_2$$

$$= \left\| V_{t-1}^{-1} \left(\sum_{s=t_0}^{t-1} X_s X_s^\top \left(\sum_{p=s}^{t-1} (\theta_p - \theta_{p+1}) \right) \right) \right\|_2$$

$$= \left\| V_{t-1}^{-1} \left(\sum_{p=t_0}^{t-1} \left(\sum_{s=t_0}^{p} X_s X_s^\top (\theta_p - \theta_{p+1}) \right) \right) \right\|_2$$

$$\leq \sum_{p=t_0}^{t-1} \left\| V_{t-1}^{-1} \left(\sum_{s=t_0}^{p} X_s X_s^\top \right) (\theta_p - \theta_{p+1}) \right\|_2$$

$$\leq \sum_{p=t_0}^{t-1} \sigma_{\max} \left(V_{t-1}^{-1} \left(\sum_{s=t_0}^{p} X_s X_s^\top \right) \right) \|\theta_p - \theta_{p+1}\|_2$$

$$\leq \sum_{p=t_0}^{t-1} \|\theta_p - \theta_{p+1}\|_2, \tag{21}$$

where $\sigma_{\max}(\cdot)$ is the largest singular value. The key is the last step (21) but its proof is questionable: they need to show the following results holds universally for all $p \in \{t_0, \ldots, t-1\}$,

$$\sigma_{\max} \left(V_{t-1}^{-1} \left(\sum_{s=t_0}^{p} X_s X_s^\top \right) \right) \leq 1. \tag{22}$$

To this end, denoted by $A = \sum_{s=t_0}^{p} X_s X_s^\top$, the authors show that $V_{t-1}^{-1} A$ shares the same characteristics polynomial with $V_{t-1}^{-1/2} A V_{t-1}^{-1/2}$, namely, $\det(\eta I - V_{t-1}^{-1} A) = \det(\eta I -$

$V_{t-1}^{-1/2}AV_{t-1}^{-1/2})$ holds for any η. Since $V_{t-1}^{-1/2}AV_{t-1}^{-1/2}$ is clearly symmetric positive semi-definite, they claim that

$$\mathbf{z}^\top V_{t-1}^{-1}A\mathbf{z} \geq 0 \tag{23}$$

also holds for $\mathbf{z} \in \mathcal{S}(1) = \{\mathbf{x} \mid \|\mathbf{x}\|_2 = 1\}$, which is crucial for their remaining proof.

$$\sigma_{\max}\left(V_{t-1}^{-1}\left(\sum_{s=t_0}^{p} X_s X_s^\top\right)\right) = \sup_{\mathbf{z}\in\mathcal{S}(1)} \mathbf{z}^\top V_{t-1}^{-1}\left(\sum_{s=t_0}^{p} X_s X_s^\top\right)\mathbf{z} \tag{24}$$

$$\overset{(23)}{\leq} \sup_{\mathbf{z}\in\mathcal{S}(1)} \left\{ \mathbf{z}^\top V_{t-1}^{-1}\left(\sum_{s=t_0}^{p} X_s X_s^\top\right)\mathbf{z} + \mathbf{z}^\top V_{t-1}^{-1}\left(\sum_{s=p+1}^{t-1} X_s X_s^\top\right)\mathbf{z} + \lambda\mathbf{z}^\top V_{t-1}^{-1}\mathbf{z} \right\} \tag{25}$$

$$= \sup_{\mathbf{z}\in\mathcal{S}(1)} \mathbf{z}^\top V_{t-1}^{-1}V_{t-1}\mathbf{z} = 1.$$

However, we identify that there are two issues in the above arguments. First, the step in (24) doubtful. For a matrix $M \in \mathbb{R}^{m\times n}$, we have $\|M\|_2 = \sup_{\|\mathbf{x}\|_2=1}\sup_{\|\mathbf{y}\|_2=1}|\mathbf{y}^\top M\mathbf{x}|$ (see for Lemma 6), while it is not warranted that $\|M\|_2 = \sup_{\|\mathbf{z}\|_2=1}|\mathbf{z}^\top M\mathbf{z}|$ which is seemingly important for the following arguments. Regardless of this first issue, the second issue about the claim (23) and the result in (25). We discover that the claim (23) is even more severe. We discover that the claim (23) is ungrounded (at least its current proof cannot stand for the correctness). The big caveat is that $V_{t-1}^{-1}A \in \mathbb{R}^{d\times d}$ is *not* guaranteed to be symmetric. The logic behind the claim is that, suppose $P, Q \in \mathbb{R}^{d\times d}$ are with the same characteristics polynomial, i.e., $\det(\eta I - Q) = \det(\eta I - P)$ holds for any η, and meanwhile P is symmetric positive semi-definite (which guarantees $\mathbf{z}^\top P\mathbf{z} \geq 0$ for any $\mathbf{z} \in \mathbb{R}^d$), then we can also have $\mathbf{z}^\top Q\mathbf{z} \geq 0$ for any $\mathbf{z} \in \mathbb{R}^d$. Unfortunately, the reasoning is not correct, and we give a simple counterexample. Let P be the 2-dim identity matrix $[1,0;0,1]$, and $Q = [1,-10;0,1]$ is an *asymmetric* matrix, then clearly $\det(\eta I - P) = \det(\eta I - Q) = (\eta - 1)^2$ is true for any η; however, $\mathbf{z}^\top Q\mathbf{z} \geq 0$ does not hold in general, for example, $\mathbf{z}^\top Q\mathbf{z} = -8 < 0$ when $\mathbf{z} = (1,1)^\top$.

Fixing the gap. The term in the left hand of (20) is crucial for the dynamic regret analysis of non-stationary linear bandits, which is expected to be converted to the path-length indicating the non-stationarity of environments. In the proof of Cheung et al. [2019a], as shown in (21) and (22), the authors aim to upper bound the crucial quantity $\sigma_{\max}\left(V_{t-1}^{-1}\left(\sum_{s=t_0}^{p} X_s X_s^\top\right)\right)$ by some constant. However, the technical reasoning is wrong. We avoid this issue and provide a new analysis as exhibited in the proof of Lemma 2, which serves as the key component in our fix. Following our analysis of Lemma 2, it is not hard to give a similar estimator error analysis for the sliding window least square estimator [Cheung et al., 2019a] and the weighted least square estimator [Russac et al., 2019], which will fix their results from $\widetilde{\mathcal{O}}(d^{2/3}T^{2/3}P_T^{1/3})$ to $\widetilde{\mathcal{O}}(d^{7/8}T^{3/4}P_T^{1/4})$. Also see related discussions in Remark 1.

Impossibility result. In the following, we further prove that the desirable claim in (22) is actually impossible. Specifically, we construct a hard problem instance to show that the key quantity $\sigma_{\max}\left(V_{t-1}^{-1}\left(\sum_{s=t_0}^{p} X_s X_s^\top\right)\right)$ cannot be universally upper bounded by any constant without square-root dependence on H. For notational convenience, we focus on the first restarting epoch, so the starting index $t_0 = 1$.

Theorem 4. *Let $L = 1$ and $\lambda = 1$. We construct the feature as*

$$X_1 = \ldots = X_p = \left[\frac{1}{\sqrt{p}}, \frac{\sqrt{p-1}}{\sqrt{p}}\right]^\top, \quad and$$

$$X_{p+1} = \ldots = X_H = \left[\frac{1}{\sqrt{H-p}}, \frac{\sqrt{H-p-1}}{\sqrt{H-p}}\right]^\top. \tag{26}$$

Denote by $A = \sum_{s=1}^p X_s X_s^\top$ and $B = \sum_{s=p+1}^H X_s X_s^\top$, then the covariance matrix is $V_{t-1} = A + B + I_d$. Under such cases, considering the checkpoint of $p = \lfloor H/3 \rfloor$, we have

$$\|V_{t-1}^{-1} A\|_2 = \sigma_{\max}(V_{t-1}^{-1} A) \geq 0.0564 \sqrt{H}. \tag{27}$$

Proof. For simplicity of notation, let $y = \sqrt{p-1}$ and $z = \sqrt{H-p-1}$. By the constructed example in (26), we have

$$A = \sum_{s=1}^p X_s X_s^\top = \begin{bmatrix} 1 & y \\ y & y^2 \end{bmatrix} \text{ and } B = \sum_{s=p+1}^H X_s X_s^\top = \begin{bmatrix} 1 & z \\ z & z^2 \end{bmatrix}.$$

For convenience, we will write the covariance matrix V_t simply V when no confusion can arise. So the concerned matrix $V^{-1} A$ can be calculated as

$$V^{-1} A = \begin{bmatrix} 2 + \lambda & y + z \\ y + z & y^2 + z^2 + \lambda \end{bmatrix}^{-1} \begin{bmatrix} 1 & y \\ y & y^2 \end{bmatrix}$$

$$= \frac{1}{(2+\lambda)(y^2 + z^2 + \lambda) - (y+z)^2} \cdot \begin{bmatrix} y^2 + z^2 + \lambda & -(y+z) \\ -(y+z) & 2 + \lambda \end{bmatrix} \begin{bmatrix} 1 & y \\ y & y^2 \end{bmatrix}$$

$$= \frac{1}{(1+\lambda)(y^2 + z^2) - 2yz + (2+\lambda)\lambda} \cdot \begin{bmatrix} z^2 - yz + \lambda & yz^2 - y^2 z + \lambda y \\ (1+\lambda)y - z & (1+\lambda)y^2 - yz \end{bmatrix}.$$

Denote by $s = (1+\lambda)(y^2 + z^2) - 2yz + (2+\lambda)\lambda$, $\alpha = z^2 - yz + \lambda$, and $\beta = (1+\lambda)y - z$, we then have

$$V^{-1} A (V^{-1} A)^\top = \frac{1 + y^2}{s^2} \begin{bmatrix} \alpha^2 & \alpha\beta \\ \alpha\beta & \beta^2 \end{bmatrix}.$$

The eigenvalues (we denote them by $\bar{\lambda}$, to distinguish the notation with the regularizer coefficient λ) of matrix $[\alpha^2, \alpha\beta; \alpha\beta, \beta^2]$ should satisfy $(\alpha^2 - \bar{\lambda})(\beta^2 - \bar{\lambda}) - \alpha^2 \beta^2 = 0$. By solving the equation, we can obtain that

$$\bar{\lambda}_{\max} = \alpha^2 + \beta^2 = (z^2 - yz + \lambda)^2 + ((1+\lambda)y - z)^2 \geq (z^2 - yz + \lambda)^2.$$

When $\lambda = 1$ and $p = aH$ (here we assume aH is an integer for simplicity), we have

$$\bar{\lambda}_{\max} \geq (z^2 - yz + \lambda)^2$$

$$= \left((1-a)p - 1 - \sqrt{p-1}\sqrt{(1-a)p - 1} + 1\right)^2$$

$$\geq \left((1-a)H - \sqrt{a(1-a)}H\right)^2$$

17

$$= (1-a)(\sqrt{1-a} - \sqrt{a})^2 H^2. \tag{28}$$

Note that the condition of $a \in (0, 1/2)$ is required to make the second inequality hold. On the other hand, we have

$$\frac{1+y^2}{s^2} = \frac{p}{(2(y^2+z^2) - 2yz + 3)^2} \geq \frac{p}{(2(y^2+z^2)+4)^2} = \frac{a}{4H}. \tag{29}$$

Combining (28) and (29), we have

$$\sigma_{\max}(V^{-1}A) = \sqrt{\lambda_{\max}(V^{-1}A(V^{-1}A)^\top)} \geq \sqrt{\bar{\lambda}_{\max} \cdot \frac{1+y^2}{s^2}} \geq \sqrt{\frac{a'}{4}} \cdot \sqrt{H},$$

where $a' = (1-a)a(\sqrt{1-a} - \sqrt{a})^2$ is a universal constant. When choosing $a = 1/3$ as selected in the main paper, $a' = 0.0127$ and the lower bound is $\sigma_{\max}(V^{-1}A) \geq 0.0564\sqrt{H}$. □

Moreover, we report some numerical results for validation: when $H = 3000$, $\sigma_{\max}(V_{t-1}^{-1}A) = 5.852$ and the theoretical lower bound is $0.0564\sqrt{H} = 3.087$; when $H = 30000$, $\sigma_{\max}(V_{t-1}^{-1}A) = 18.474$ and the theoretical lower bound is $0.0564\sqrt{H} = 9.763$.

4.4   Proof of Theorem 3

Proof. We begin with the following decomposition of the dynamic regret.

$$\sum_{t=1}^T \langle X_t^*, \theta_t \rangle - \langle X_t, \theta_t \rangle = \underbrace{\sum_{t=1}^T \langle X_t^*, \theta_t \rangle - \sum_{i=1}^{\lceil T/\Delta \rceil} \sum_{t=(i-1)\Delta+1}^{i\Delta} \langle X_t(H^\dagger), \theta_t \rangle}_{\texttt{base-regret}}$$

$$+ \underbrace{\sum_{i=1}^{\lceil T/\Delta \rceil} \sum_{t=(i-1)\Delta+1}^{i\Delta} \langle X_t(H^\dagger), \theta_t \rangle - \langle X_t(H_i), \theta_t \rangle}_{\texttt{meta-regret}},$$

where $H^\dagger$ is the best restarting period to approximate the optimal restarting period H^* in the pool $\mathcal{H}$, and $H^* = \lfloor (dT/(P_T)^{2/3} \rfloor$. The first term is the dynamic regret of RestartUCB with the best restarting period in the candidate pool $\mathcal{H}$, and hence called base-regret. The second term is the regret overhead of meta-algorithm due to adaptive exploration of unknown optimal restarting period, and is thus called the meta-regret. We bound the two terms respectively.

We first consider the base-regret. Indeed, from the construction of candidate restarting periods pool $\mathcal{H}$, we confirm that there exists an restarting period $H^\dagger \in \mathcal{H}$ such that $H^\dagger \leq H^* \leq 2H^\dagger$. Therefore, employing the dynamic regret bound (11) in Theorem 2, we have the following upper bound for the base-regret:

$$\texttt{base-regret} \leq \sum_{i=1}^{\lceil T/\Delta \rceil} \widetilde{\mathcal{O}}\left(d^{\frac{1}{2}} H^{\dagger \frac{3}{2}} P_i + \frac{d\Delta}{\sqrt{H^\dagger}} \right) \tag{30}$$

18

$$= \widetilde{\mathcal{O}}\left(d^{\frac{1}{2}}H^{\dagger\frac{3}{2}}P_T + \frac{dT}{\sqrt{H^{\dagger}}}\right) \tag{31}$$

$$\leq \widetilde{\mathcal{O}}\left(d^{\frac{1}{2}}H^{*\frac{3}{2}}P_T + \frac{dT}{\sqrt{2H^{*}}}\right) \tag{32}$$

$$= \widetilde{\mathcal{O}}\left(d^{\frac{7}{8}}T^{\frac{3}{4}}P_T^{\frac{1}{4}}\right), \tag{33}$$

where (30) is due to Theorem 2 and P_i denotes the path-length in the i-th episode of the meta-learner's update. (31) follows by summing over all update episodes, and the inequality (32) holds since the optimal restarting period H^* is provably in the range of $[H_{\min}, H_{\max}]$ and satisfies $H^{\dagger} \leq H^* \leq 2H^{\dagger}$.

Next, we give an upper bound for the meta-regret. The analysis follows the proof argument in the sliding window based approach [Cheung et al., 2019b, Proposition 1]. Note that the definition of the meta-regret is defined over the *expected* reward (namely, $\mathbb{E}[r_t(X)] = X^{\top}\theta_t$), whereas the actual returned feedback is the noisy one (i.e., $r_t(X) = X^{\top}\theta_t + \eta_t$) which might be unbounded due to the additive sub-Gaussian noise. Fortunately, the light-tail property enables us to continue the use of adversarial MAB algorithms, e.g., Exp3 [Auer et al., 2002]. Specifically, by the concentration inequality endowed by the sub-Gaussian noise, we know that the received reward lies in the bounded region with high probability, which is presented in Lemma 7. Denote by $\mathcal{E}$ the event that Lemma 7 holds, and denote by $R_i \triangleq \sum_{t=(i-1)\Delta+1}^{i\Delta}\langle X_t(H^{\dagger}), \theta_t \rangle - \langle X_t(H_i), \theta_t \rangle$ the instantaneous regret of the meta learner. The meta-regret follows

$$\texttt{meta-regret} = \mathbb{E}\left[\sum_{i=1}^{\lceil T/\Delta \rceil} R_i\right]$$

$$= \mathbb{E}\left[\sum_{i=1}^{\lceil T/\Delta \rceil} R_i \,\Big|\, \mathcal{E}\right] \cdot \Pr[\mathcal{E}] + \mathbb{E}\left[\sum_{i=1}^{\lceil T/\Delta \rceil} R_i \,\Big|\, \overline{\mathcal{E}}\right] \cdot \Pr[\overline{\mathcal{E}}]$$

$$\leq \mathcal{O}\left(L_{\max}\sqrt{\frac{T}{\Delta}N}\right) \cdot \left(1 - \frac{2}{T}\right) + \mathcal{O}(T) \cdot \frac{2}{T}$$

$$= \mathcal{O}(\Delta NT) \leq \widetilde{\mathcal{O}}(d^{1/2}T^{3/4}), \tag{34}$$

where $L_{\max} \triangleq \max L_i$ for $i \in [\lceil T/\Delta \rceil]$ denotes the maximum cumulative loss in all episodes. The first equation is by definition, and the second one is by the law of total expectation. The next inequality follows from the following two aspects: the quantity is bounded according to the standard regret guarantee of Exp3 [Auer et al., 2002] when the event $\mathcal{E}$ holds; and it is trivially upper bounded when the event $\mathcal{E}$ does not happen. The failing probability is controlled by Lemma 7. The final equation is true by checking the parameters that the episode length is $\Delta = \lceil d\sqrt{T} \rceil$, and the number of candidate restarting periods N is of order $\mathcal{O}(\log T)$, hence be omitted in the $\widetilde{\mathcal{O}}(\cdot)$-notation.

Combining the upper bounds of base-regret (33) and meta-regret (34), we obtain that the expected dynamic regret of RestartUCB-BOB is bounded by $\widetilde{\mathcal{O}}(d^{\frac{7}{8}}T^{\frac{3}{4}}P_T^{\frac{1}{4}})$, which completes the proof of Theorem 3.

It is also worthy noting that the base-regret is actually with high-probability guarantees, while the meta-regret in our analysis only holds in expectation. Actually, we can

19

boost the result to the high-probability version by employing advanced meta-algorithms such as Exp3.IX [Neu, 2015] that can achieve a high-probability regret bound for adversarial MAB problems, as well as using the union bound in the analysis. □

5   Experiments

Despite the focus of this paper is on the theoretical aspect, we present empirical studies to further evaluate the proposed approach.

Contenders.   We study two kinds of non-stationary environments: the underlying parameter is *abruptly changing* or *gradually changing*. We will simulate both environments and details can be found in the next paragraph. We compare RestartUCB to (a) WindowUCB, based on the sliding window least square [Cheung et al., 2019a]; (b) WeightUCB, based on the weighted least square [Russac et al., 2019]; (c) StaticUCB, the algorithm designed for stationary linear bandits [Abbasi-Yadkori et al., 2011]. In the scenario of abrupt change, we additionally compare with OracleRestartUCB, which knows the exact information of change points a priori and restarts the algorithm when reaching a change point. Evidently, OracleRestartUCB is not a practical algorithm, which actually serves as the skyline of all the approaches.

Settings.   In abruptly-changing environments, the unknown regression parameter θ_t is periodically set as $[1, 0]$, $[-1, 0]$, $[0, 1]$, $[0, -1]$ in the first half of iterations, and $[1, 0]$ for the remaining iterations. In gradually-changing environments, the unknown regression parameter θ_t is moved from $[1, 0]$ to $[-1, 0]$ on the unit circle continuously. In both scenarios, we set $T = 50,000$ and number of arms $n = 20$. The feature is sampled from normal distribution $\mathcal{N}(0, 1)$ and rescaled such that $L = 1$. The random noise is generated according to $\mathcal{N}(0, 0.1)$. Since the path-length P_T is available in the synthetic datasets, we set the weight $\gamma = 1 - 1/\tau$ for WeightUCB, the window size $w = \tau$ for WindowUCB, and the restarting period $H = \tau$ for RestartUCB, here $\tau = 10 * \lfloor d^{1/4} T^{1/2} P_T^{-1/2} \rfloor$ is set as suggested by the theory. The simulation is repeated for 50 times, and we report the average and standard deviation.

Results.   Figure 3 shows performance comparisons of different approaches for non-stationary linear bandits. The performance is measured by the (pseudo-) dynamic regret, which is plotted the in y-axis in the logarithmic scale. In the *abruptly-changing environments*, OracleRestartUCB is definitely the best one as was expected since it knows exact information of change points a priori, and StaticUCB ranks the last as it does not take the non-stationarity issue into consideration. RestartUCB and WindowUCB have comparable performance, better than WeightUCB. Actually, RestartUCB is even slightly better than WindowUCB. We note that RestartUCB has an additional advantage over WindowUCB in terms of the computational issue: RestartUCB supports the one-pass update without storing historical data, whereas WindowUCB has to maintain a buffer and thus needs to scan data multiple times owing to the sliding window strategy. In the *gradually-changing environments*, WeightUCB ranks the first, followed by WindowUCB and RestartUCB. Nevertheless, as will be shown later, WeightUCB takes a significantly longer running time than our approach.

Figure 4 reports the running time including both mean and standard deviation. We

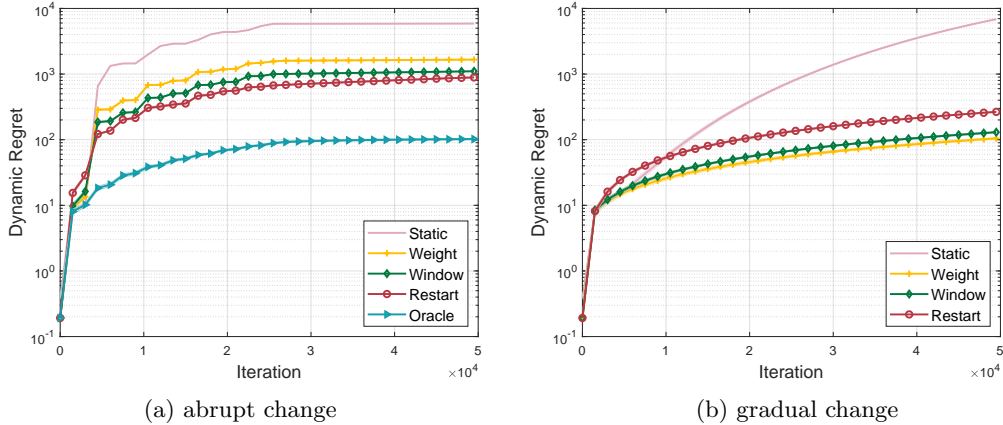| | (a) abrupt change | (b) gradual change |

Figure 3: Comparisons of different approaches in terms of dynamic regret. Note that the y-axis is plotted in the logarithmic scale.
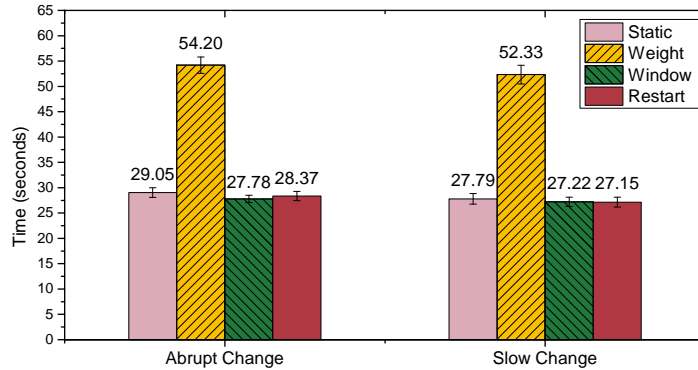

Figure 4: Comparisons of different approaches in terms of running time.

can see that the time costs of RestartUCB, WindowUCB and StaticUCB are almost the same. By contrast, WeightUCB requires a significantly longer running time, nearly twice the cost of other contenders. The reason lies in the fact that WeightUCB algorithm involves the computation of inverse of covariance matrix $V_t \in \mathbb{R}^{d \times d}$ and its variant $\widetilde{V}_t \in \mathbb{R}^{d \times d}$, while other three methods maintain and manipulate only one covariance matrix. It is worthy to note that our approach can be further accelerated by the recursive least square. This will save the inverse computation of the covariance matrix, which will be particularly desired in high-dimensional problems.

6   Conclusion

In this paper, we study the problem of non-stationary linear bandits, where the unknown regression parameter θ_t is changing over time. We propose a simple algorithm based on the restarted strategy, which enjoys strong theoretical guarantees notwithstanding its simplicity. Concretely, when the path-length of underlying parameters P_T is known, our proposed RestartUCB algorithm enjoys an $\widetilde{\mathcal{O}}(d^{7/8}T^{3/4}P_T^{1/4})$ dynamic regret, which

shares the same regret guarantees with previous methods developed in the literature yet is with more favorable computational advantage. In addition, we show that the same dynamic regret guarantee is attainable even when P_T is unknown by further using ResatartUCB as the base algorithm and combining the bandits-over-bandits mechanism as the meta scheduling. Empirical studies validate the efficacy of the proposed approach, particularly in the abruptly-changing environments.

The current upper bounds do not match the existing lower bound, even when the path-length term is known. In the future, we would like to investigate how to get rid of this regret gap and further study how to design algorithms for non-stationary linear bandits that achieve rate-optimal dynamic regret without prior information.

Acknowledgment

References

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems 24 (NIPS)*, pages 2312–2320, 2011.

Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E. Schapire. Corralling a band of bandit algorithms. In *Proceedings of the 30th Conference on Learning Theory (COLT)*, pages 12–38, 2017.

Robin Allesiardo, Raphaël Féraud, and Odalric-Ambrym Maillard. The non-stationary stochastic multi-armed bandit problem. *International Journal of Data Science and Analytics*, 3(4):267–283, 2017.

Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.

Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The non-stochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.

Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best arm with an unknown number of distribution changes. In *Proceedings of the 14th European Workshop on Reinforcement Learning European Workshop on Reinforcement Learning (EWRL)*, pages 393–400, 2018.

Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Proceedings of the 32nd Conference On Learning Theory (COLT)*, volume 99, pages 138–158, 2019.

Omar Besbes, Yonatan Gur, and Assaf J. Zeevi. Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244, 2015.

Omar Besbes, Yonatan Gur, and Assaf J. Zeevi. Optimal exploration–exploitation in a multi-armed bandit problem with non-stationary rewards. *Stochastic Systems*, 9(4): 319–337, 2019.

Yifang Chen, Chung-Wei Lee, Haipeng Luo, and Chen-Yu Wei. A new algorithm for non-stationary contextual bandits: Efficient, optimal, and parameter-free. In *Proceedings of the 32nd Conference On Learning Theory (COLT)*, pages 696–726, 2019.

Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Learning to optimize under non-stationarity. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1079–1087, 2019a.

Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Hedging the drift: Learning to optimize under non-stationarity. *ArXiv preprint*, arXiv:1903.01461, 2019b. URL https://arxiv.org/abs/1903.01461v2.

Wei Chu, Lihong Li, Lev Reyzin, and Robert E. Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 208–214, 2011.

Matej Crepinsek, Shih-Hsi Liu, and Marjan Mernik. Exploration and exploitation in evolutionary algorithms: A survey. *ACM Computing Surveys*, 45(3):35:1–35:33, 2013.

Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems 20 (NIPS)*, pages 345–352, 2007.

Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 1405–1411, 2015.

Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for switching bandit problems. In *Proceedings of the 22nd International Conference on Algorithmic Learning Theory (ALT)*, volume 6925, pages 174–188, 2011.

Lei Guo, Lennart Ljung, and Pierre Priouret. Performance analysis of the forgetting factor RLS algorithm. *International Journal of Adaptive Control and Signal Processing*, 7(6):525–538, 1993.

Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. Improved strongly adaptive online learning using coin betting. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 943–951, 2017.

Baekjin Kim and Ambuj Tewari. Randomized exploration for non-stationary stochastic linear bandits. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 71–80, 2020.

Yingkai Li, Yining Wang, and Yuan Zhou. Nearly minimax-optimal regret for linearly parameterized bandits. In *Proceedings of the 32nd Conference on Learning Theory (COLT)*, pages 2173–2174, 2019.

Carl D Meyer. *Matrix Analysis and Applied Linear Algebra*, volume 71. SIAM, 2000.

Gergely Neu. Explore no more: Improved high-probability regret bounds for nonstochastic bandits. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pages 3168–3176, 2015.

Philippe Rigollet and Jan-Christian Hútter. *High Dimensional Statistics*. Lecture Notes, 2019.

Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.

Yoan Russac, Claire Vernade, and Olivier Cappé. Weighted linear bandits for nonstationary environments. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 12040–12049, 2019.

Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.

Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 1330–1340, 2018.

Lijun Zhang, Shiyin Lu, and Tianbao Yang. Minimizing dynamic regret and adaptive regret simultaneously. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 309–319, 2020a.

Yu-Jie Zhang, Peng Zhao, and Zhi-Hua Zhou. A simple online algorithm for competing with dynamic comparators. In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 390–399, 2020b.

Peng Zhao, Guanghui Wang, Lijun Zhang, and Zhi-Hua Zhou. Bandit convex optimization in non-stationary environments. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1508–1518, 2020a.

Peng Zhao, Lijun Zhang, Yuan Jiang, and Zhi-Hua Zhou. A simple approach for nonstationary linear bandits. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 746–755, 2020b.

Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 12510–12520, 2020c.

Peng Zhao, Xinqiang Wang, Siyu Xie, Lei Guo, and Zhi-Hua Zhou. Distribution-free one-pass learning. *IEEE Transaction on Knowledge and Data Engineering*, 33:951–963, 2021.

Kai Zheng, Haipeng Luo, Ilias Diakonikolas, and Liwei Wang. Equipping experts/bandits with long-term memory. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 5927–5937, 2019.

Zhi-Hua Zhou. *Ensemble Methods: Foundations and Algorithms*. Chapman & Hall/CRC Press, 2012.

A Technical Lemmas

In this section, we provide several technical lemmas that frequently used in the proofs.

Theorem 5 (Self-Normalized Bound for Vector-Valued Martingales [Abbasi-Yadkori et al., 2011, Theorem 1]). *Let $\{F_t\}_{t=0}^{\infty}$ be a filtration. Let $\{\eta_t\}_{t=0}^{\infty}$ be a real-valued stochastic process such that η_t is F_t-measurable and conditionally R-sub-Gaussian for some $R > 0$, namely,*

$$\forall \lambda \in \mathbb{R}, \quad \mathbb{E}[\exp(\lambda \eta_t) \mid F_{t-1}] \leq \exp\left(\frac{\lambda^2 R^2}{2}\right). \tag{35}$$

Let $\{X_t\}_{t=1}^{\infty}$ be an $\mathbb{R}^d$-valued stochastic process such that X_t is F_{t-1}-measurable. Assume that V is a $d \times d$ positive definite matrix. For any $t \geq 0$, define

$$\bar{V}_t = V + \sum_{\tau=1}^{t} X_\tau X_\tau^\top, \quad S_t = \sum_{\tau=1}^{t} \eta_\tau X_\tau. \tag{36}$$

Then, for any $\delta > 0$, with probability at least $1 - \delta$, for all $t \geq 0$,

$$\|S_t\|_{\bar{V}_t^{-1}}^2 \leq 2R^2 \log\left(\frac{\det(\bar{V}_t)^{1/2} \det(V)^{-1/2}}{\delta}\right). \tag{37}$$

Lemma 4 (Elliptical Potential Lemma). *Suppose $U_0 = \lambda I$, $U_t = U_{t-1} + X_t X_t^\top$, and $\|X_t\|_2 \leq L$, then*

$$\sum_{t=1}^{T} \|U_{t-1}^{-\frac{1}{2}} X_t\|_2 \leq \sqrt{2dT \log\left(1 + \frac{L^2 T}{\lambda d}\right)}. \tag{38}$$

Proof. First, we have the following decomposition,

$$U_t = U_{t-1} + X_t X_t^\top = U_{t-1}^{\frac{1}{2}}(I + U_{t-1}^{-\frac{1}{2}} X_t X_t^\top U_{t-1}^{-\frac{1}{2}})U_{t-1}^{\frac{1}{2}}.$$

Taking the determinant on both sides, we get

$$\det(U_t) = \det(U_{t-1}) \det(I + U_{t-1}^{-\frac{1}{2}} X_t X_t^\top U_{t-1}^{-\frac{1}{2}}),$$

which in conjunction with Lemma 5 yields

$$\det(U_t) = \det(U_{t-1})(1 + \|U_{t-1}^{-\frac{1}{2}} X_t\|_2^2) \geq \det(U_{t-1}) \exp(\|U_{t-1}^{-\frac{1}{2}} X_t\|_2^2 / 2).$$

Note that in the first inequality, we utilize the fact that $1 + x \geq \exp(x/2)$ holds for any $x \in [0, 1]$. By taking advantage of the telescope structure, we have

$$\sum_{t=1}^{T} \|U_{t-1}^{-\frac{1}{2}} X_t\|_2^2 \leq 2 \log \frac{\det(U_T)}{\det(U_0)} \leq 2d \log \left(1 + \frac{L^2 T}{\lambda d}\right),$$

where the last inequality follows from the fact that $\text{Tr}(U_T) \leq \text{Tr}(U_0) + L^2 T = \lambda d + L^2 T$, and thus $\det(U_T) \leq (\lambda + L^2 T/d)^d$. Therefore, Cauchy-Schwarz inequality implies,

$$\sum_{t=1}^{T} \|U_{t-1}^{-\frac{1}{2}} X_t\|_2 \leq \sqrt{T \sum_{t=1}^{T} \|U_{t-1}^{-\frac{1}{2}} X_t\|_2^2} \leq \sqrt{2dT \log \left(1 + \frac{L^2 T}{\lambda d}\right)}.$$

$\square$

Lemma 5. *For any $\mathbf{v} \in \mathbb{R}^d$, we have*

$$\det(I + \mathbf{v}\mathbf{v}^\top) = 1 + \|\mathbf{v}\|_2^2.$$

Proof. Notice that

(i) $(I + \mathbf{v}\mathbf{v}^\top)\mathbf{v} = (1 + \|\mathbf{v}\|_2^2)\mathbf{v}$, therefore, $\mathbf{v}$ is its eigenvector with $(1 + \|\mathbf{v}\|_2^2)$ as the eigenvalue;

(ii) $(I + \mathbf{v}\mathbf{v}^\top)\mathbf{v}^\perp = \mathbf{v}^\perp$, therefore, $\mathbf{v}^\perp \perp \mathbf{v}$ is its eigenvector with 1 as the eigenvalue.

Consequently, $\det(I + \mathbf{v}\mathbf{v}^\top) = 1 + \|\mathbf{v}\|_2^2$. $\square$

Lemma 6 (Property 5.2.9 of Meyer [2000]). *For a real matrix $A \in \mathbb{R}^{m \times n}$, we have*

$$\|A\|_2 = \sup_{\|\mathbf{x}\|_2=1} \sup_{\|\mathbf{y}\|_2=1} |\mathbf{y}^\top A \mathbf{x}|.$$

Proof. The proof is from the solution manual of Meyer [2000]. Applying the Cauchy-Schwarz inequality yields $|\mathbf{y}^\top A \mathbf{x}| \leq \|\mathbf{y}\|_2 \|A\mathbf{x}\|_2$, which implies that

$$\sup_{\|\mathbf{x}\|_2=1} \sup_{\|\mathbf{y}\|_2=1} |\mathbf{y}^\top A \mathbf{x}| \leq \sup_{\|\mathbf{x}\|_2=1} \|A\mathbf{x}\|_2 = \|A\|_2.$$

Now show that equality is actually attained for some pair $\mathbf{x}$ and $\mathbf{y}$ on the unit 2-sphere. To do so, notice that when setting $\mathbf{x}_*$ is a vector of unit length such that

$$\|A\mathbf{x}_*\|_2 = \sup_{\|\mathbf{x}\|=1} \|A\mathbf{x}\|_2 = \|A\|_2,$$

and $\mathbf{y}_*$ is the vector such that

$$\mathbf{y}_* = \frac{A\mathbf{x}_*}{\|A\mathbf{x}_*\|_2} = \frac{A\mathbf{x}_*}{\|A\|_2},$$

then

$$\mathbf{y}_*^\top A \mathbf{x}_* = \frac{\mathbf{x}_*^\top A^\top A \mathbf{x}_*}{\|A\|_2} = \frac{\|A\mathbf{x}_*\|_2^2}{\|A\|_2} = \frac{\|A\|_2^2}{\|A\|_2} = \|A\|_2.$$

Hence we complete the proof. $\square$

Lemma 7. *Let $N = \lceil T/\Delta \rceil$. Denote by L_i the absolute value of cumulative rewards for episode i, i.e., $L_i \triangleq \sum_{t=(i-1)\Delta+1}^{i\Delta} r_t(X_t)$, then*

$$\Pr\left[\forall i \in [N], L_i \le LS\Delta + 2R\sqrt{\Delta \ln \frac{T}{\sqrt{\Delta}}}\right] \ge 1 - \frac{2}{T}. \tag{39}$$

Proof. For any episode i, the absolute sum of rewards can be written as

$$\left|\sum_{t=(i-1)\Delta+1}^{i\Delta} \langle X_t, \theta_t \rangle + \eta_t\right| \le \sum_{t=(i-1)\Delta+1}^{i\Delta} |\langle X_t, \theta_t \rangle| + \left|\sum_{t=(i-1)\Delta+1}^{i\Delta} \eta_t\right|$$

$$\le \Delta LS + \left|\sum_{t=(i-1)\Delta+1}^{i\Delta} \eta_t\right|,$$

where we have iteratively applied the triangle inequality as well as the fact that $|\langle X_t, \theta_t \rangle| \le LS$ for all t.

Further applying the standard concentration result of R-sub-Gaussian random variables [Rigollet and Hútter, 2019, Corollary 1.7], we get Now by property of the R-sub-Gaussian, it holds that

$$\Pr\left[\left|\frac{1}{\Delta}\sum_{t=(i-1)\Delta+1}^{i\Delta} \eta_t\right| \ge \varepsilon\right] \le 2\exp\left(-\frac{\Delta \varepsilon^2}{2R^2}\right),$$

which further implies that

$$\Pr\left[\left|\sum_{t=(i-1)\Delta+1}^{i\Delta} \eta_t\right| \ge 2R\sqrt{\Delta \ln \frac{T}{\sqrt{\Delta}}}\right] \le \frac{2\Delta}{T^2}.$$

So we can ensure a low failing probability, specifically, the probability of the event that the absolute value of the noise term η_t exceeds $2R\sqrt{\ln T}$ for a fixed t is at most $1/T^2$. By union bound, we have

$$\Pr\left[\exists i \in [N] : \left|\sum_{t=(i-1)\Delta+1}^{i\Delta} \eta_t\right| \ge 2R\sqrt{\Delta \ln \frac{T}{\sqrt{\Delta}}}\right]$$

$$\le \sum_{i=1}^{\lceil T/\Delta \rceil} \Pr\left[\left|\sum_{t=(i-1)\Delta+1}^{i\Delta} \eta_t\right| \ge 2R\sqrt{\Delta \ln \frac{T}{\sqrt{\Delta}}}\right] \le \frac{2}{T}.$$

Hence, we finish the proof. □

27